

RedactOR: An LLM-Powered Framework for Automatic Clinical Data De-Identification

Praphul Singh*, Charlotte Dzialo*, Jangwon Kim, Sumana Srivatsa,
Irfan Bulu, Sri Gadde, Krishnaram Kenthapadi

Oracle Health AI

{praphul.singh, charlotte.dzialo, jangwon.kim, sumana.srivatsa, irfan.bulu, sri.gadde, krishnaram.kenthapadi}@oracle.com

Abstract

Ensuring clinical data privacy while preserving utility is critical for AI-driven healthcare and data analytics. Existing de-identification (De-ID) methods, including rule-based techniques, deep learning models, and large language models (LLMs), often suffer from recall errors, limited generalization, and inefficiencies, limiting their real-world applicability. We propose a fully automated, multi-modal framework, RedactOR for de-identifying structured and unstructured electronic health records, including clinical audio records. Our framework employs cost-efficient De-ID strategies, including intelligent routing, hybrid rule and LLM based approaches, and a two-step audio redaction approach. We present a retrieval-based entity relexicalization approach to ensure consistent substitutions of protected entities, thereby enhancing data coherence for downstream applications. We discuss key design desiderata, de-identification and relexicalization methodology, and modular architecture of RedactOR and its integration with Oracle Health Clinical AI system. Evaluated on the i2b2 2014 De-ID dataset using standard metrics with strict recall, our approach achieves competitive performance while optimizing token usage to reduce LLM costs. Finally, we discuss key lessons and insights from deployment in real-world AI-driven healthcare data pipelines.

1 Introduction

The proliferation of AI-driven healthcare tools has heightened the need for robust de-identification (De-ID) systems to comply with privacy regulations such as HIPAA in the US and GDPR in the EU (Ahmed et al., 2020). Effective De-ID is critical for secure AI model training, evaluation, and debugging, data analytics, and clinical deployment (see §A.2). However, automating De-ID for electronic health records (EHRs) is challenging due

to data heterogeneity, schema variability, context-sensitive Protected Health Information (PHI) or Personally Identifiable Information (PII), and the multi-modal nature of healthcare data—text, images, and audio (Mohamed et al., 2023; Kayaalp, 2018).

Manual De-ID, though accurate, is impractical at scale given the data volume in clinical settings (Patterson et al., 2024). Automated approaches, including rule-based methods, BERT-based models, and LLMs (Meystre et al., 2010; Kovačević et al., 2024; Altalla’ et al., 2025), face limitations in generalization, contextual reasoning, and efficiency, particularly when trained on narrow datasets that do not reflect real-world EHR diversity (Liu et al., 2023). Since even a single leak of PHI/PII can have serious privacy implications, reliable, scalable De-ID remains a critical need.

Recent advancements address cost, scalability, and generalizability through techniques like prompt optimization, model quantization (Shekhar et al., 2024; Arefeen et al., 2024), and intelligent agent routing (Varangot-Reille et al., 2025), along with multi-modal De-ID for text and audio (Dhingra et al., 2024). Yet challenges persist – specifically with, maintaining high recall, ensuring consistent PHI/PII substitution (e.g., “Wilson” and “Dr. Adam Wilson” both mapped to “Chang” and “Dr. Kevin Chang” respectively), and evaluating privacy risks with stricter metrics beyond token-level scores.

We propose RedactOR, a fully automated, multi-modal framework for de-identifying structured and unstructured patient records, including clinical audio records. RedactOR combines LLM-based processing for unstructured text with rule-based handling of structured data to achieve low cost and latency, and extends text de-identification for audio redaction. Our framework includes a novel retrieval-based entity relexicalization component to ensure consistent PHI/PII replacement, enhancing coherence and privacy. We present key design

*These authors contributed equally to this work.

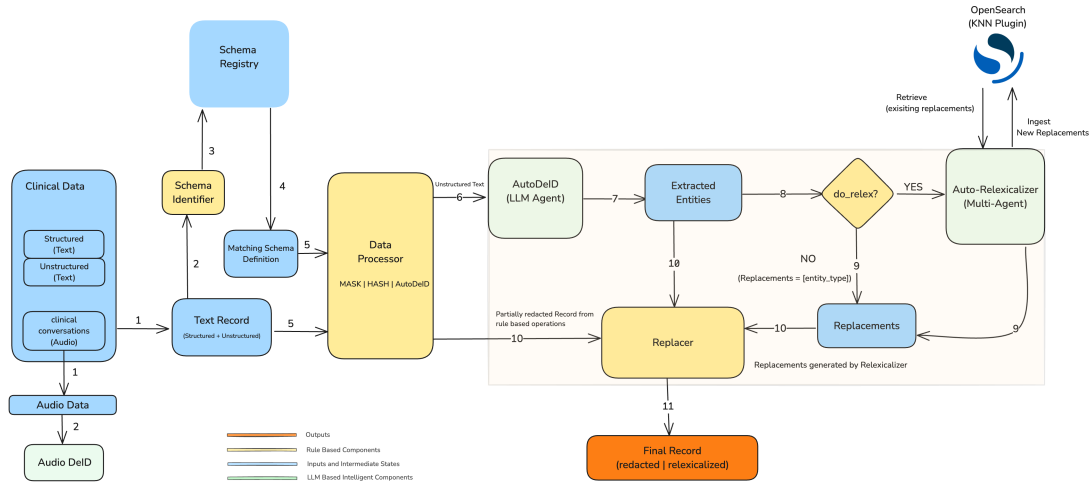


Figure 1: Architectural overview of RedactOR

requirements (§3.1), de-identification and relexicalization methodology and architecture (§3.2), and integration with Oracle Health Clinical AI system (§A.8). We demonstrate that our framework outperforms other LLM-based approaches and achieves performance comparable to specialized, closed-source solutions while remaining adaptable through prompt engineering – eliminating reliance on large annotated datasets (§4), and highlight lessons and insights from 12+ months of deployment in real-world clinical AI system (§5).

2 Related Work

Rule-based and ML-based De-ID. Rule-based systems rely on pattern matching, lexicons, and heuristics (Neamatullah et al., 2008; Meystre et al., 2010), offering simplicity and no training data requirements (Negash et al., 2023). However, rule creation is time-consuming and may lack robustness (Negash et al., 2023; Lee et al., 2017). Machine learning models, especially BiLSTM-based approaches (Ma and Hovy, 2016; Dernoncourt et al., 2017; Liu et al., 2017), improve generalization without manual rules but struggle to transfer across datasets (Stubbs et al., 2017; Yang et al., 2019). BERT-based models enhance De-ID (Meaney et al., 2022) but demand significant compute resources, hyperparameter tuning, and still exhibit gaps in handling certain PHI/PII types. **LLM-based De-ID.** LLMs offer flexible, zero/few-shot De-ID capabilities. Kim et al. used GPT-4 to augment training data, improving BERT model performance across datasets. Yashwanth and Shetkar showed fine-tuned LLMs outperform zero-shot models, particularly under format shifts. Altalla’

et al. found GPT-4 surpasses GPT-3.5 in De-ID accuracy and synthetic data generation. Similarly, Wiest et al. developed a custom open-source LLM-based Anonymizer pipeline benchmarking 8 LLMs to De-ID 250 German clinical letters. However, most studies lack evaluation on large real-world cross-dataset generalization.

Synthetic Data. The growing need for large datasets in medical research, alongside strict patient privacy rules, has led to increased interest in synthetic data. Synthetic data generation often involves differential privacy based approaches to protect patient privacy and generative adversarial networks (GAN) based methods for realistic data replication, and its utility depends on maintaining fidelity and minimizing biases to ensure reliable research and clinical decisions (Al Aziz et al., 2021). While synthetic data offers transformative potential for healthcare, careful consideration is needed to ensure its ethical and effective use in research and practice (Altalla’ et al., 2025).

Relexicalization. Replacing PHI/PII with realistic surrogates is underexplored. Many systems apply dummy replacements or simple rules (e.g., gender matching) (Sweeney, 1996; Alfalahi et al., 2012; Lison et al., 2021). Recent work (Vakili et al., 2024) demonstrated pseudonymizing BERT models for privacy-preserving data analysis, highlighting relexicalization’s value in maintaining utility while protecting privacy.

3 System Design and Architecture

3.1 Design Desiderata

RedactOR is designed around three core principles:

scalability, adaptability, and cost-efficiency. Scalability is achieved through an end-to-end automation pipeline, enabling efficient processing of structured and unstructured clinical data while minimizing computational overhead via intelligent routing and LLM-based De-ID strategies. Adaptability is ensured through a schema-agnostic processing architecture, facilitating seamless integration across heterogeneous EHR formats and multimodal data sources (text and audio) without the need for dataset-specific fine-tuning. Cost-efficiency is realized through token usage optimization in text-based De-ID and leveraging text-based entity extraction for audio, thereby eliminating the need for computationally expensive, audio-specific de-identification models. Additionally, retrieval-based re-lexicalization enhances contextual consistency in PHI/PII replacements, preserving both privacy and downstream utility, making the system highly effective for real-world AI-driven healthcare applications.

3.2 Architecture Overview

RedactOR consists of three main components: (i) Auto De-ID, Audio De-ID, and Auto Relexicalizer (see Figure 1).

First, Schema Identifier automatically identifies the appropriate schema from the Schema Registry based on the `dataType` parameter in each data instance (see §A.9.1) and forwards it to the Data Processor along with the corresponding text data. The data processor is designed to be agnostic to text data types, requiring only the schema (stored in the schema registry) with predefined rule flags for each field. Currently, a rule flag can be one of the following: (i) `passThrough` (rule-based) retains the field without any changes (used for non-PHI/PII data), (ii) `shouldMask` (rule-based) replaces PHI/PII fields with generic placeholders (e.g., [PERSON]), (iii) `shouldHash` (rule-based pseudo-anonymization) hashes identifiers to enable secure linkages across documents within the same domain, or (iv) `autoDeID` (LLM-based) applies LLM-based De-ID to the unstructured text fields. This schema-agnostic design is crucial for scaling our system to support health data De-ID tasks. Meanwhile, audio data is processed separately by the Audio De-ID component. To ensure that no PHI/PII field definition is missed, we enforce a human review of the schema before it is pushed into the schema registry.

Auto De-ID is an LLM component (§B.1) that

processes the context extracted by the data processor. It can support a dynamic list of entity types. We support 33 entity types in our production deployment as shown in §C. This context is split into chunks of a pre-defined size (ω), ensuring optimal model performance without exceeding LLM context length limits. Chunks are processed in parallel across a fixed number (p) of passes. ω and p are heuristically chosen hyperparameters. In each pass, the LLM extracts entities along with their surrounding context as position hints – that is, each extracted entity includes nearby words that uniquely identify its location in the text (e.g., “76 years old” instead of just “76”, or “Mr. John Smith, the patient” instead of just “John”). This context-aware extraction enables accurate entity matching and redaction without relying on potentially unreliable character position indices. In the first pass, the LLM detects as many entities as possible. In subsequent passes, entities already identified are masked in the text, prompting the model to focus on previously missed or hard-to-detect PHI/PII entities. Extracted entities from all passes are aggregated to form the final entity set.

The Auto Relexicalizer, a multi-agent component (see §B.2 and Figure 4), replaces redacted entities with contextually consistent and realistic alternatives. Relexicalization not only improves the usability of de-identified data but it also strengthens privacy by increasing the Hiding in Plain Sight (HIPS) factor (Carrell et al., 2020). Ensuring that replaced entities blend seamlessly with any remaining leaked PHI/PII makes re-identification attempts significantly more challenging. A combined example of Auto De-ID followed by Relexicalization is shown in §A.9. It employs multiple agents as follows:

- LLM-Based Entity Clustering: grouping extracted entities based on their context.
- Hybrid Retrieval (Vector Search + Filtering): retrieving pre-existing replacements.
- LLM-Based Validation: Determining the validity of the retrieved replacements.
- LLM-Based Generation: Generating new replacements for invalid retrievals.
- OpenSearch Indexing: Storing new replacements for future reuse.

Our work extends a recent work (Vakili et al., 2024) that presents an analysis of pseudo-anonymization. We offer an LLM-driven alternative for automated and scalable relexicalization. A regex-based replacer replaces extracted entities

with entity-type masks (e.g., [PERSON]) to redact PHI/PII in unstructured text fields. See §A.4.1 for an end-to-end example of Auto De-ID.

Our Audio De-ID feature performs a two-step redaction process to enhance privacy. First, it uses Automatic Speech Recognition (ASR) to detect timestamps of spoken words and applies LLM-based Auto De-ID to the transcript, adding an extra 100 – 200 msec of margin at token boundaries for improved protection. Next, it examines *unrecognized* voiced regions – identified by an aggressive Voice Activity Detection (VAD) – by analyzing their surrounding context words with an LLM (§B.3), evaluating the likelihood of these regions containing PHI/PII and selects the most likely ones. Finally, it mutes all time boundaries (including margins) for predicted PHI/PII tokens’ voiced regions. Testing on our internal data showed that the second step increased recall by approximately 10%. A brief end-to-end example illustrating this is presented in §A.5.1. In summary, our Audio De-ID component improves PHI/PII detection by addressing ASR misalignment and deletion errors by leveraging VAD and LLM based detection process in the second step.

By integrating Auto De-ID for unstructured text, Auto Relexicalizer for realistic entity replacement, and Audio De-ID for speech data, RedactOR provides a scalable, adaptable, and cost-effective De-ID pipeline that secures both text and audio data while preserving its utility.

See §A.4, §A.5, and §A.6 for detailed algorithmic descriptions of Auto De-ID, Auto Relexicalization, and Audio De-ID, respectively.

4 Experiments

We present the results of evaluating RedactOR against other LLM-based approaches and specialized, closed-source commercial solutions over a publicly available medical record dataset. For parity with other methods, we turn off the Auto Relexicalizer component. We set chunk size (ω) to 256 and number of passes (p) to 2.

4.1 Dataset

We evaluated using 2014 i2b2/UTHealth De-ID corpus (Stubbs and Uzuner, 2015) which is widely used in clinical De-ID research. This dataset comprises longitudinal clinical records for 296 patients (with 2-5 records per patient). The annotation scheme follows HIPAA guidelines and includes

additional indirect identifiers such as detailed date components (e.g., year), geographic information (states, countries), hospital names, clinician names, and patient professions. For our experiments, we randomly subsampled 100 clinical notes and evaluated on seven PHI/PII entity categories: AGE, CONTACT, DATE, ID, LOCATION, PERSON, and PROFESSION.

4.2 Comparative De-ID methods

We evaluated RedactOR against recent LLM-based methods, Yashwanth and Shettar (2024) (with their two prompt variants: ‘brief’ and ‘detailed’) and Altalla’ et al. (2025), as well as commercial De-ID APIs from AWS (Amazon Web Services) (AWS, 2025) and JSL (John Snow Labs) (Kocaman et al., 2023, 2025). For a fair comparison, we used GPT-4o (OpenAI, 2025) for all LLM-based methods. See §A.3 for additional details.

4.3 Evaluation Methods

We assessed De-ID performance using traditional metrics such as precision, recall and F1-score, as well as all-or-nothing recall, applied to (PERSON, AGE, CONTACT, ID, LOCATION). All-or-nothing recall (Scaiano et al., 2016) determines whether every instance of a given entity type or a document is correctly redacted. If any instance is missed, all-or-nothing recall is set to 0; otherwise, it is set to 1.

We evaluated all systems using a *stricter* methodology for true positive computation incorporating entity position matching. Position information is critical for data redaction in unstructured health records, as it often differentiates PHI from clinical information. For example, in the phrase ‘76 yrs old,’ the number ‘76’ represents age (PHI), whereas in ‘oxygen saturation rate is 76,’ it denotes a vital sign.

In the evaluation for each entity type, we compared our system with AWS and JSL using additional matching criteria, including entity-level text matching and label matching. The metrics were assessed at the PHI/PII entity level (multi-word spans) rather than individual tokens, as in (Yashwanth and Shettar, 2024), ensuring a fair comparison between entities with varying token counts. To account for minor variations (e.g., ‘Mrs. Mary Smith’ vs. ‘Mary Smith’), we applied the Levenshtein similarity with a heuristically determined threshold of 0.6. Yashwanth and Shettar (2024) and Altalla’ et al. (2025) are omitted in entity type spe-

Model	Precision	Recall	F1-score	All-Or-Nothing Recall
Y&S_Brief	0.5634	0.6580	0.6070	0.3700
Y&S_Detail	0.6178	0.8270	0.7072	0.5600
Altalla	0.9675	0.6715	0.7927	0.3600
RedactOR	0.9769	0.9525	0.9646	0.7900
AWS	0.9549	0.9425	0.9487	0.7500
JSL	0.9481	0.9865	0.9669	0.9000

Table 1: Performance of zero-shot GPT-4o and commercial De-ID systems on all PHI/PII entities. This evaluation does not consider the entity type constraint.

cific evaluation, because they provide only binary PHI/PII labels, but not entity types.

4.4 Comparison with LLM-Based De-ID Frameworks

As shown in Table 1, a comparison with other LLM-based methods (Yashwanth and Shettar, 2024; Altalla’ et al., 2025) indicates that RedactOR outperforms existing methods, achieving the highest F1-score of 0.9646 with a well-balanced precision and recall. The high recall can be attributed to RedactOR’s multi-chunk and multi-pass strategy, which systematically refines entity detection by iteratively masking extracted entities and forcing the model to focus on overlooked PHI. Similarly, in terms of precision, RedactOR outperforms Yashwanth and Shettar (2024), highlighting the effectiveness of its context-aware entity extraction. By leveraging contextual clues and maintaining intradocument consistency, RedactOR reduces false positives, whereas single-pass prompting methods tend to over-redact ambiguous terms.

While Yashwanth and Shettar (2024)’s Detailed version achieves higher recall than the Brief, it does so at the expense of precision. This highlights a fundamental trade-off in zero-shot protected terms extraction: models optimized for recall often over-mask non-protected terms, leading to reduced utility of the redacted text. RedactOR strikes a balance between recall and precision, making it more suitable for real-world clinical applications where both PHI/PII removal and utility are essential.

4.5 Comparisons with Commercial APIs

Unlike prior zero-shot LLM-based approaches, commercial De-ID APIs (e.g., AWS, JSL) are fine-tuned on proprietary clinical datasets. Table 1 shows that while RedactOR does not surpass JSL in recall, it achieves higher precision and a comparable F1-score. This suggests that while JSL bene-

fits from domain-specific fine-tuning, RedactOR’s context-aware extraction minimizes false positives, leading to more reliable entity masking.

Table 2 presents a breakdown of performance by entity types. RedactOR demonstrates high precision and strong recall across all entities, particularly excelling on CONTACT and PERSON entities. RedactOR achieves perfect recall on CONTACT entities, outperforming AWS and JSL, and shows competitive performance on PERSON and DATE entities. However, it underperforms on LOCATION and ID, presumably due to the structural variability of PHI in clinical texts. LOCATION entities, in particular, often appear within complex sentence structures, posing challenges to generic LLM-based masking. This suggests the need for instruction updates for these entities or a specialized LLM.

4.6 All-or-Nothing Recall Results

Table 3 shows performance by PHI/PII entity type with all-or-nothing recall, a stricter metric requiring both correct entity type and position alignment. These results highlight RedactOR’s strength in high-precision redaction for certain entities while emphasizing the advantage of domain-tuned models for broader recall coverage.

RedactOR outperforms other methods using the same underlying LLM (Table 1) and approaches the performance of specialized models like JSL. Specifically, it excels in CONTACT and ID compared to commercial methods. RedactOR’s multi-pass extraction and context-aware masking enhance the LLM’s effectiveness, showcasing its strength without specialized fine-tuning. However, JSL leads in most other entity types, achieving the highest overall recall. The recall gap – especially for complex entities like DATE and LOCATION – highlights the need for improved prompt instructions and possibly a specialized LLM for de-identification to match domain-specific systems.

4.7 Ablation Study with Open-Source Model

To demonstrate the adaptability of RedactOR to open-source models and evaluate the benefit of its multi-pass de-identification strategy, we conducted an ablation study using LLaMA-3.2-3B-Instruct (MetaAI, 2024) – a compact, publicly available LLM.

Figure 2 illustrates the effect of increasing the number of passes (from 1 to 4) on all-or-nothing recall across seven PHI/PII entity types. Notably, we

Entity Type	RedactOR			AWS			JSL		
	Recall	Precision	F1-score	Recall	Precision	F1-score	Recall	Precision	F1-score
AGE	0.8987	0.9930	0.9435	0.9684	0.9935	0.9808	0.9748	0.9688	0.9718
CONTACT	1.0000	1.0000	1.0000	1.0000	0.4545	0.6250	0.7879	1.0000	0.8814
DATE	0.9495	0.9988	0.9735	0.9202	0.9949	0.9561	0.9918	0.9883	0.9900
ID	0.9275	0.6667	0.7758	0.7917	0.6129	0.6909	0.8667	0.9123	0.8889
LOCATION	0.7855	1.0000	0.8799	0.7890	0.9820	0.8750	0.9469	0.9808	0.9636
PERSON	0.9595	0.9912	0.9751	0.9461	0.9461	0.9461	0.9572	0.9933	0.9749
PROFESSION	0.9167	1.0000	0.9565	0.9130	0.7778	0.8400	1.0000	0.9565	0.9778
All	0.9159	0.9790	0.9465	0.9042	0.9510	0.9270	0.9664	0.9839	0.9751

Table 2: Performance of De-ID systems for each entity type and all data.

observe consistent improvement across most entity types as the number of passes increases. The largest relative gains are observed between pass 1 and pass 2, especially for sparse or context-sensitive types such as ID, DATE, and LOCATION, which tend to be missed in early passes but are recovered in subsequent iterations.

For dominant or well-signaled types like PERSON, CONTACT, and PROFESSION, the recall is already high at pass 2, with marginal improvements beyond that point. By pass 3, the recall curve starts to saturate for most entity types, indicating diminishing returns on additional passes.

These trends suggest that the number of passes is a critical, model-dependent hyperparameter: lightweight models like LLaMA-3.2-3B benefit from 2–3 passes, while larger models may reach optimal performance sooner. RedactOR supports this flexibility by treating the pass count as a configurable parameter, allowing practitioners to trade off between computational cost and de-identification completeness depending on the capacity of the underlying LLM.

4.8 Qualitative Evaluation of Audio De-ID on Internal Data

To assess the impact of our two-step Audio De-ID process, we conducted a qualitative evaluation on an internal clinical audio dataset. The LLM-based

Entity Type	RedactOR	AWS	JSL
AGE	0.8904	0.9589	0.9452
CONTACT	1.0000	1.0000	0.7666
DATE	0.7300	0.6000	0.9500
ID	0.8958	0.8333	0.8541
LOCATION	0.6923	0.6026	0.8333
PERSON	0.8556	0.8041	0.8659
PROFESSION	0.9091	0.9090	1.0000
All	0.8214	0.7701	0.8906

Table 3: All-or-nothing recalls constrained on entity types for RedactOR and commercial models.

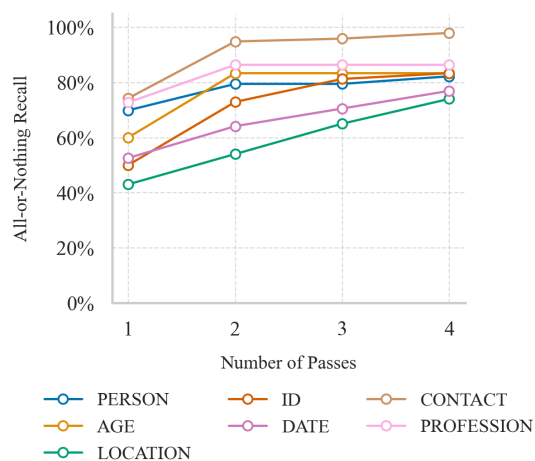


Figure 2: Entity-wise all-or-nothing recall for LLaMA-3.2-3B as the number of passes increases from 1 to 4. Most entities show the greatest gain between pass 1 and 2, with diminishing improvements thereafter.

timestamp detector identified and muted several additional audio segments that were missed by text de-identification on just the transcript. A subset of these newly muted segments included person names that had previously gone undetected, leading to a noticeable improvement in all-or-nothing recall on direct identifiers—approximately 12%. Another small portion involved relevant medical or personal context (e.g., complaints or medications), introducing minor precision trade-offs. The majority of muted segments, however, were non-informative, consisting of background noise or idle speech such as keyboard activity. Overall, over 84% of the additionally muted content was deemed to have no negative impact on clinical utility. These findings demonstrate that the second-pass audio detection enhances recall with minimal utility loss, validating its inclusion in real-world deployments.

5 Deployment Lessons and Insights

In the course of developing and deploying RedactOR, we realized that efficient scaling of de-identification is crucial to handle the large volume of healthcare data post-deployment, including audio files, SOAP notes, and longitudinal records with thousands of FHIR resources (Bender and Sartipi, 2013). Once the service is launched, a continuous influx of data follows, and as the products expand, the ability to optimize processing at scale becomes critical. Simple yet powerful reductions in computation and processing can significantly impact efficiency, cost, and system performance.

One key optimization we employed was reducing token usage. RedactOR extracts only PHI/PII entities and their position hints, minimizing the number of LLM’s output tokens by approximately 50%. This reduction not only lowers processing costs but also decreases latency, ensuring that De-ID remains accurate and efficient at scale.

Further, processing each FHIR resource individually causes delays and backlogs. Our schema agnostic approach allows us to batch lightweight resources (e.g., vitals, medications) in the schema processor by merging their schemas and free-text fields into a single composite schema. For example, if the batch size is n , n resource schemas could be combined into one, allowing all associated text to be de-identified in a single LLM request. The batch size can be chosen heuristically based on the LLM used, the context length it supports, the system prompt size, and the average number of tokens present in the unstructured texts of the resource schemas.

Finally, dynamic batching further enhances scalability by grouping incoming resources based on size and complexity. This approach enables large and diverse datasets to be processed in real time, preventing bottlenecks as data streams grow.

Our initial implementation focused on just de-identification but we decided to incorporate relexicalization after realizing that relexicalized data significantly enhances the ease of use by applied scientists as part of their model training, quality & bias evaluation, and debugging pipelines since this data has similar format and characteristics as the production data.

6 Conclusion And Future Directions

Motivated by the need for protecting patient privacy while enabling utility, we presented RedactOR

a multi-modal, scalable, flexible, and cost-efficient LLM-powered framework for clinical data de-identification, and demonstrated its efficacy in de-identifying 33 PHI/PII entities over the i2b2 dataset. We showed that our approach outperforms other LLM-based methods and achieves performance comparable to specialized, closed-source solutions. Further, RedactOR supports relexicalizing redacted entities with contextually consistent alternatives, enhancing data usability and strengthening privacy. By presenting the methodology, technical architecture, and lessons learned from over 12 months of production deployment as part of Oracle Health Clinical AI system, we hope that the insights and experience from our work are useful for researchers and practitioners working on clinical AI systems.

There are several avenues for future work. The variability in healthcare datasets across institutions affects generalizability, necessitating adaptive prompting techniques. While our method excels in detecting ID, PERSON, and DATE entities, it may require further refinement of entity-specific LLM instructions (e.g., for address- and occupation-related entities). More broadly, RedactOR’s generalizability can be enhanced across diverse institutional datasets by integrating domain-adaptive prompt enhancements.

Another direction is to investigate domain-adaptive VAD techniques in de-identification settings. Although we incorporate a VAD-based solution to mitigate ASR inaccuracies, our use of a simpler VAD algorithm introduces false positives, leading to over-redaction. Additionally, transcription variability due to noise and overlapping speech increases the risk of PHI/PII leakage. Integrating deep learning-based VAD models alongside domain-adaptive ASR techniques could enhance precision while maintaining recall, reducing unnecessary redactions without compromising PHI/PII protection.

Furthermore, we could design standardized benchmarks for evaluating relexicalization techniques, following ideas discussed in §A.7. An ideal dataset would include PHI/PII-tagged text, gold relexicalized outputs, and context to ensure consistency across documents and domains such as clinical notes, transcripts, and structured records. More broadly, a promising direction is to extend our framework to handle other modalities such as medical images and videos, and design corresponding end-to-end evaluation methodologies and benchmarks.

Acknowledgments

We would like to thank other members of Oracle Health AI for their collaboration while deploying RedactOR in production, and Weifeng Bao, Brent Beardsley, Michelle Chen, Dipankar Das, Long Duong, Raefer Gabriel, Kent Grueneich, Neil Hauge, Bhagya Hettige, Brad Jacobs, Mark Johnson, Shirley Liu, Cody Maheu, Atri Mandal, Virendra Marathe, Kiran Rama, Amitabh Saikia, Tushar Shandhilya, Gyan Shankar, and Vishal Vishnoi for insightful feedback and discussions.

References

- Tanbir Ahmed, Md Momin Al Aziz, and Noman Mohammed. 2020. De-identification of electronic health record using neural network. *Scientific reports*, 10(1):18600.
- Md Momin Al Aziz, Tanbir Ahmed, Tasnia Faequa, Xiaoqian Jiang, Yiyu Yao, and Noman Mohammed. 2021. Differentially private medical texts generation using generative neural networks. *ACM Transactions on Computing for Healthcare (HEALTH)*, 3(1):1–27.
- Alyaa Alfalahi, Sara Brissman, and Hercules Dalianis. 2012. Pseudonymisation of personal names and other PHIs in an annotated clinical Swedish corpus. In *Third Workshop on Building and Evaluating Resources for Biomedical Text Mining (BioTextM 2012) Held in Conjunction with LREC*, pages 49–54.
- Bayan Altalla', Sameera Abdalla, Ahmad Altamimi, Layla Bitar, Amal Al Omari, Ramiz Kardan, and Iyad Sultan. 2025. Evaluating GPT models for clinical note de-identification. *Scientific Reports*, 15(1):3852.
- Md Adnan Arefeen, Biplob Debnath, and Srimat Chakradhar. 2024. LeanContext: Cost-efficient domain-specific question answering using LLMs. *Natural Language Processing Journal*, 7:100065.
- AWS. 2025. Amazon Comprehend Medical. <https://aws.amazon.com/comprehend/medical/>. Accessed: March, 2025.
- Duane Bender and Kamran Sartipi. 2013. HL7 FHIR: An agile and RESTful approach to healthcare information exchange. In *Proceedings of the 26th IEEE international symposium on computer-based medical systems*, pages 326–331. IEEE.
- David S Carrell, Bradley A Malin, David J Cronkite, John S Aberdeen, Cheryl Clark, Muqun Li, Dikshya Bastakoty, Steve Nyemba, and Lynette Hirschman. 2020. Resilience of clinical text de-identified with “hiding in plain sight” to hostile reidentification attacks by human readers. *Journal of the American Medical Informatics Association*, 27(9):1374–1382.
- Franck Dernoncourt, Ji Young Lee, and Peter Szolovits. 2017. NeuroNER: An easy-to-use program for named-entity recognition based on neural networks. *arXiv preprint arXiv:1705.05487*.
- Priyanshu Dhingra, Satyam Agrawal, Chandra Sekar Veerappan, Thi Nga Ho, Eng Siong Chng, and Rong Tong. 2024. Speech de-identification data augmentation leveraging large language model. In *2024 International Conference on Asian Language Processing (IALP)*, pages 97–102. IEEE.
- Mehmet Kayaalp. 2018. Patient privacy in the era of big data. *Balkan medical journal*, 35(1):8–17.
- Woojin Kim, Sungeun Hahm, and Jaejin Lee. 2024. Generalizing clinical de-identification models by privacy-safe data augmentation using GPT-4. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 21204–21218.
- Veysel Kocaman, Muhammed Santas, Yigit Gul, Mehmet Butgul, and David Talby. 2025. Can zero-shot commercial APIs deliver regulatory-grade clinical text deidentification? In *ECIR Workshop on Narrative Extraction from Texts (Text2Story)*.
- Veysel Kocaman, D Talby, and H Ul Hak. 2023. RWD143 beyond accuracy: Automated de-identification of large real-world clinical text datasets. *Value in Health*, 26(12):S532.
- Aleksandar Kovačević, Bojana Bašaragin, Nikola Milošević, and Goran Nenadić. 2024. De-identification of clinical free text using natural language processing: A systematic review of current approaches. *Artificial intelligence in medicine*, page 102845.
- Hee-Jin Lee, Yonghui Wu, Yaoyun Zhang, Jun Xu, Hua Xu, and Kirk Roberts. 2017. A hybrid approach to automatic de-identification of psychiatric notes. *Journal of biomedical informatics*, 75:S19–S27.
- Pierre Lison, Ildikó Pilán, David Sánchez, Montserrat Batet, and Lilja Øvrelid. 2021. Anonymisation models for text data: State of the art, challenges and future directions. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4188–4203.
- Zengjian Liu, Buzhou Tang, Xiaolong Wang, and Qingcai Chen. 2017. De-identification of clinical notes via recurrent neural network and conditional random field. *Journal of biomedical informatics*, 75:S34–S42.
- Zhengliang Liu, Yue Huang, Xiaowei Yu, Lu Zhang, Zihao Wu, Chao Cao, Haixing Dai, Lin Zhao, Yiwei Li, Peng Shu, et al. 2023. DeID-GPT: Zero-shot medical text de-identification by GPT-4. *arXiv preprint arXiv:2303.11032*.

- Xuezhe Ma and Eduard Hovy. 2016. End-to-end sequence labeling via bi-directional lstm-cnns-crf. *arXiv preprint arXiv:1603.01354*.
- Christopher Meaney, Wali Hakimpour, Sumeet Kalia, and Rahim Moineddin. 2022. A comparative evaluation of transformer models for de-identification of clinical text data. *arXiv preprint arXiv:2204.07056*.
- MetaAI. 2024. [LLaMA 3.2](#). Accessed: 2025-05-20.
- Stephane M Meystre, F Jeffrey Friedlin, Brett R South, Shuying Shen, and Matthew H Samore. 2010. Automatic de-identification of textual documents in the electronic health record: A review of recent research. *BMC medical research methodology*, 10:1–16.
- Yahia Mohamed, Xing Song, Tamara M McMahon, Suman Sahil, Meredith Zozus, Zhan Wang, Greater Plains Collaborative, and Lemuel R Waitman. 2023. Electronic health record data quality variability across a multistate clinical research network. *Journal of Clinical and Translational Science*, 7(1):e130.
- Ishna Neamatullah, Margaret M Douglass, Li-Wei H Lehman, Andrew Reisner, Mauricio Villarroel, William J Long, Peter Szolovits, George B Moody, Roger G Mark, and Gari D Clifford. 2008. Automated de-identification of free-text medical records. *BMC medical informatics and decision making*, 8:1–17.
- Bekelu Negash, Alan Katz, Christine J Neilson, Moniruzzaman Moni, Marcello Nesca, Alexander Singer, and Jennifer E Enns. 2023. De-identification of free text data containing personal health information: A scoping review of reviews. *International Journal of Population Data Science*, 8(1):2153.
- OpenAI. 2025. GPT-4o documentation. <https://platform.openai.com/docs/models/gpt-4o>. Accessed: March, 2025.
- Brian W Patterson, Daniel J Hekman, Frank J Liao, Azita G Hamedani, Manish N Shah, and Majid Afshar. 2024. Call me Dr Ishmael: Trends in electronic health record notes available at emergency department visits and admissions. *JAMIA open*, 7(2):ooae039.
- Vivek Podder, Valerie Lew, and Sassan Ghassemzadeh. 2024. SOAP notes. *StatPearls [Internet]*.
- Martin Scaiano, Grant Middleton, Luk Arbuckle, Varada Kolhatkar, Liam Peyton, Moira Dowling, Debbie S Gipson, and Khaled El Emam. 2016. A unified framework for evaluating the risk of re-identification of text de-identification tools. *Journal of biomedical informatics*, 63:174–183.
- Shivanshu Shekhar, Tanishq Dubey, Koyel Mukherjee, Apoorv Saxena, Atharv Tyagi, and Nishanth Kotla. 2024. Towards optimizing the costs of LLM usage. *arXiv preprint arXiv:2402.01742*.
- Amber Stubbs, Michele Filannino, and Özlem Uzuner. 2017. De-identification of psychiatric intake records: Overview of 2016 CEGS N-GRID shared tasks Track 1. *Journal of biomedical informatics*, 75:S4–S18.
- Amber Stubbs and Özlem Uzuner. 2015. Annotating longitudinal clinical narratives for de-identification: The 2014 i2b2/UTHealth corpus. *Journal of biomedical informatics*, 58:S20–S29.
- Latanya Sweeney. 1996. Replacing personally-identifying information in medical records, the Scrub system. In *Proceedings of the AMIA annual fall symposium*, page 333.
- Thomas Vakili, Aron Henriksson, and Hercules Dalianis. 2024. End-to-end pseudonymization of fine-tuned clinical BERT models: Privacy preservation with maintained data utility. *BMC Medical Informatics and Decision Making*, 24(1):162.
- Clovis Varangot-Reille, Christophe Bouvard, Antoine Gourru, Mathieu Ciancone, Marion Schaeffer, and François Jacquenet. 2025. Doing more with less—implementing routing strategies in large language model-based systems: An extended survey. *arXiv preprint arXiv:2502.00409*.
- Isabella C Wiest, Marie-Elisabeth Leßmann, Fabian Wolf, Dyke Ferber, Marko Van Treeck, Jiefu Zhu, Matthias P Ebert, Christoph Benedikt Westphalen, Martin Wermke, and Jakob Nikolas Kather. 2025. Deidentifying medical documents with local, privacy-preserving large language models: The LLM-anonymizer. *NEJM AI*, 2(4):A1dbp2400537.
- Xi Yang, Tianchen Lyu, Qian Li, Chih-Yin Lee, Jiang Bian, William R Hogan, and Yonghui Wu. 2019. A study of deep learning methods for de-identification of clinical notes in cross-institute settings. *BMC medical informatics and decision making*, 19:1–9.
- YS Yashwanth and Rajashree Shettar. 2024. Zero and few short learning using large language models for de-identification of medical records. *IEEE Access*.

A Appendix

A.1 Ethics Statement

Despite the high performance of our De-ID system, there remains a non-zero risk that some PHI/PII might not be detected or removed. Consequently, any output produced by automated De-ID system should still be handled with the same security and privacy precautions as raw identifiable data. We underscore that users of our De-ID framework should apply rigorous privacy safeguards when handling the processed data, just as they would for original clinical records. For example, we restrict access to the de-identified data using encryption and access control mechanisms, and require scientists and engineers to go through appropriate privacy and

healthcare regulation related trainings before being granted access to the de-identified data.

Given the sensitivity of De-ID data pipelines, we do not release the prompt verbatim or source code of RedactOR to prevent potential privacy risks and attacks. However, to support transparency and reproducibility, we provide descriptions of individual components, including retrieval-based relexicalization, hybrid rule/LLM logic, intelligent routing, and a two-step audio redaction process, in the pipeline. This approach enables secure replication of our methodology while safeguarding patient confidentiality in our system.

A.2 Intended Use of De-Identified Data.

The de-identified data produced by RedactOR is intended for a variety of critical use cases within AI-driven healthcare systems. First, it serves as a valuable resource for *debugging production issues*, enabling engineers and data scientists to analyze system behavior and identify root causes of errors without compromising patient privacy. Second, the data supports *understanding production model behavior*, providing insights into model performance, biases, and failure modes, which guide iterative improvements and model refinements. Finally, the de-identified dataset can be leveraged for *training and evaluating downstream machine learning models*, including clinical documentation automation, clinical named entity recognition, and “needle in a haystack” tasks such as rare condition detection or retrieval of highly specific information from longitudinal records. Additionally, *de-identified data is essential for conducting R&D and facilitates collaboration with external researchers and clinicians*, enabling innovation while ensuring compliance with privacy regulations. These applications illustrate the dual importance of ensuring privacy while maintaining data utility for real-world healthcare advancements.

A.3 All Models

We compare our methods with others as follows:

1. [Yashwanth and Shettar \(2024\)](#): This study uses a zero-shot approach with two prompts – brief and detailed – applied with GPT-3.5 and GPT-4. The model returns “[Censored]” in lieu of explicit entity labels. In our experiments, we evaluate this approach using the GPT-4o to be a fair comparison with the Auto De-ID model.

2. [Altalla’ et al. \(2025\)](#): This study employs a zero-shot prompt with GPT-3 and GPT-4, where outputs are marked “[Redacted]” rather than providing explicit entity annotations. Although originally evaluated on a proprietary dataset, we adapt this baseline for the i2b2 corpus and evaluate it using the GPT-4o variant.
3. RedactOR (ours): The Auto De-ID model’s outputs are post-processed to align with our predefined PHI/PII entity categories. We use the following configuration parameters to ensure consistency across experiments: max_passes = 2, max_words = 256, and temperature = 0.
4. Commercial Cloud APIs: We assess two widely adopted commercial De-ID services that offer API-based solutions for PHI/PII extraction, namely, AWS (Amazon Web Services Medical Comprehend) and JSL (John Snow Labs). For these services, we standardize the output entity tags to align with our evaluation schema, ensuring fair comparison across systems.

These settings were chosen based on preliminary tuning to balance performance and computational efficiency.

A.4 Auto De-ID Algorithm

Algorithm 1 Auto De-ID Algorithm

Require: Text T , chunk size ω , prediction model M , entity types \mathcal{E} , passes p

Ensure: Redacted text \hat{T}

- 1: Split T into $m = \lceil |T|/\omega \rceil$ chunks.
 - 2: **for** each chunk c_i **do**
 - 3: **for** $j = 0$ to $p - 1$ **do**
 - 4: Extract entities $\mathcal{R}_i \leftarrow M(c_i, \mathcal{E})$
 - 5: Update fact dictionary $\mathcal{D} \leftarrow \mathcal{D} \cup \mathcal{R}_i$
 - 6: **end for**
 - 7: **end for**
 - 8: Replace detected entities in T with placeholders. **return** \hat{T}
-

The Auto De-ID algorithm processes text T by segmenting it into $m = \lceil |T|/\omega \rceil$ chunks. Each chunk undergoes p passes of entity extraction, where the set of identified entities is aggregated:

$$\mathcal{D} = \bigcup_{i=1}^m \bigcup_{j=1}^p \mathcal{R}_i$$

Choice of Chunk Size: The chunk size ω is chosen to balance entity extraction accuracy. Sending large texts in a single request increases the risk of missing entities, as the model may fail to attend to all parts equally. By processing smaller chunks, we reduce the chance of entity leakage and improve recall.

Multiple Passes Strategy: Using multiple passes (p) helps mitigate biases in the model’s attention. In the first pass, the model extracts easily detectable entities. In the subsequent passes, previously detected entities are masked using their entity type (e.g., [PERSON]), allowing the model to focus on overlooked entities. This iterative process enhances recall, especially for underrepresented entity types. These hyperparameters can be adjusted based on the capabilities of the chosen LLM.

Handling Context in Redaction: A naive approach to redaction might simply replace all extracted entity mentions in the text with their corresponding entity types (e.g., replacing “76” with [AGE]). However, this often leads to over-redaction when identical strings appear in different contexts. Consider the sentence:

“The patient is 76 years old and takes 76 mg of aspirin daily.”

In this case, only the first occurrence of “76” refers to the patient’s age and should be redacted as [AGE]. The second occurrence of “76” is part of a medication dosage and should not be redacted as age. If we blindly replace all instances of “76” with [AGE], we would incorrectly redact “76 mg,” resulting in a loss of valuable clinical information.

Our entity extraction method avoids this by instructing the LLM to extract entities along with sufficient context that signals their specific meaning and position in the text. In this example, “76 years old” would be extracted as an [AGE] entity, while “76 mg” would either be ignored or extracted as a separate [DOSAGE] entity. This ensures that only the appropriate mention is redacted.

Since obtaining exact character positions from LLMs is unreliable, context-based entity extraction allows us to align each detected entity with its precise occurrence in the text, ensuring accurate and minimal redaction.

Detected entities in T are replaced with placeholders to yield the final redacted text \hat{T} , ensuring sensitive data is masked correctly while preserving non-sensitive content.

A.4.1 Example Workflow

To illustrate Auto De-ID’s process, consider the following structured JSON input:

Field	Value
patient_name	Robert Johnson
patient_id	A12345
gender	male
medical_history	Robert Johnson, a patient aged 53 was admitted to Springfield General Hospital for chest pain. Dr. Mary Smith prescribed medication.

Step 0: Schema Identification The De-ID schema specifies rules for each field:

Field	De-ID Rule
patient_name	Mask as [PERSON]
patient_id	Hash
gender	Pass-through
medical_history	Auto De-ID (LLM-based)

Step 1: Schema Processing The structured fields are processed:

Field	Processed Value
patient_name	[PERSON]
patient_id	HASH(A12345)
gender	male

Step 2: Chunking The unstructured text is split into overlapping chunks:

Chunk	Text
C1	"Robert Johnson, a patient aged 53 was admitted to Springfield General Hospital for chest pain."
C2	"for chest pain. Dr. Mary Smith prescribed medication."

Step 3: LLM-Based Entity Extraction The model extracts PHI:

Entity Type	Extracted Entities
PERSON	Robert Johnson, Dr. Mary Smith
ORGANIZATION	Springfield General Hospital

Step 4: Multi-Pass Refinement Previously detected entities are masked in subsequent passes:

Pass	Input Chunk for 2nd Pass After Masking
2nd	"[PERSON], a patient aged 53 was admitted to [ORGANIZATION] for chest pain."
2nd	"for chest pain. [PERSON] prescribed medication."

Step 5: Final Entity Extraction Previously detected entities combined with the second pass extractions.

Entity Type	Extracted Entities
PERSON	Robert Johnson, Dr. Mary Smith
ORGANIZATION	Springfield General Hospital
AGE	53

Step 5: Final Redaction. The final de-identified record:

Field	Final Value
patient_name	[PERSON]
patient_id	HASH(A12345)
gender	male
medical_history	"[PERSON], a patient aged [AGE] was admitted to [ORGANIZATION] for chest pain. [PERSON] prescribed medication."

This final output ensures all PHI/PII is masked while maintaining text coherence.

A.5 Audio De-ID Algorithm

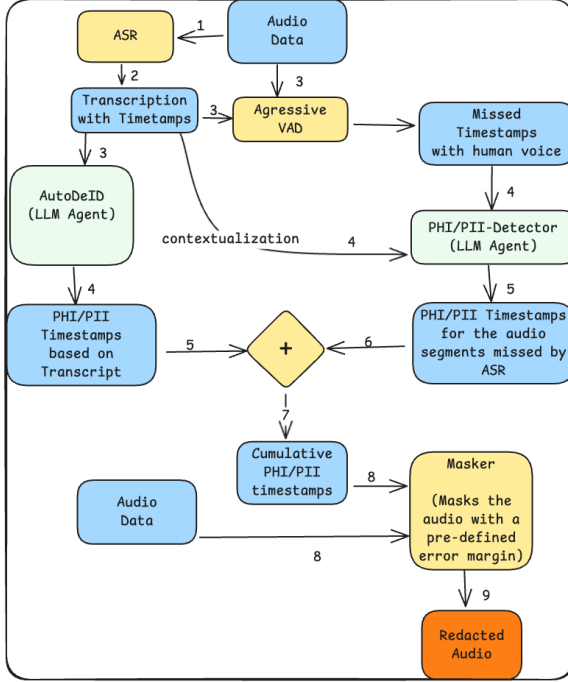


Figure 3: Audio De-Id Workflow Diagram

Algorithm 2 Audio De-ID Algorithm

Require: Audio A , ASR model M_{ASR} , VAD M_{VAD} , De-ID model M_{Deid} , PHI/PII Detector M_{LLM} , entity types \mathcal{E}

Ensure: Redacted Audio $A_{\hat{T}}$

- 1: Generate transcript $T \leftarrow M_{ASR}(A)$
- 2: Extract PHI/PII $\mathcal{D} \leftarrow M_{Deid}(T, \mathcal{E})$
- 3: Identify missing timestamps T_{missing} and detect human speech with M_{VAD}
- 4: **for** each human-voiced timestamp t_{human} **do**
- 5: Extract context, detect PHI/PII with M_{LLM} , and update \mathcal{D}
- 6: **end for**
- 7: Mute detected PHI/PII in A . **return** $A_{\hat{T}}$

Audio De-ID first converts speech to text using ASR:

$$T = M_{ASR}(A)$$

Entities are extracted from T to construct \mathcal{D} :

$$\mathcal{D} = M_{Deid}(T, \mathcal{E})$$

Gaps in ASR timestamps T_{missing} are analyzed with M_{VAD} to identify human speech regions, where PHI/PII detection is refined using an LLM-based model. The final redacted audio $A_{\hat{T}}$ is generated by muting PHI-containing segments.

A.5.1 Example Workflow

ASR-Generated Transcript:

“The patient visited Dr. Smith last week a follow-up in his clinic at Creekwood Hospital. They discussed medication changes and scheduled the next appointment for next month. The patient also mentioned feeling unwell over the weekend.”

Auto De-ID Detected PHI:

- “Dr. Smith” (00:04.23 - 00:04.80)
- “Creekwood Hospital” (00:12.57 - 00:13.20)

Identifying Missing Timestamps (Set Subtraction):

- (00:02.85 - 00:02.95) (Missed speech)
- (00:07.42 - 00:07.57) (Missed speech)
- (00:15.10 - 00:15.30) (Missed speech)

VAD Filtering (Keeping Only Human Speech Segments):

- (00:02.85 - 00:02.95) - Human voice detected ✓
- (00:07.42 - 00:07.57) - Human voice detected ✓
- (00:15.10 - 00:15.30) - Background noise, discarded ✗

Timestamp Adjustment for ASR Errors: To compensate for ASR errors, timestamps are adjusted with a safe margin of 300ms:

- **Before:** “Dr. Smith” (00:04.23 - 00:04.80)
- **After:** “Dr. Smith” (00:03.93 - 00:05.10)
- **Before:** “Creekwood Hospital” (00:12.57 - 00:13.20)

- **After:** “Creekwood Hospital” (00:12.27 - 00:13.50)

Reconstructing the Transcript: The transcript is updated by inserting missing timestamps:

```

“<human_timestamp_(00:02.85
- 00:02.95)> The patient
visited Dr. Smith last week
<human_timestamp_(00:07.42
- 00:07.57)> a follow-up in his
clinic at Creekwood Hospital.
They discussed medication
changes and scheduled the next
appointment for next month. The
patient also mentioned feeling
unwell over the weekend.”

```

This updated transcript is analyzed by an LLM to predict the likelihood of inserted timestamps containing PHI/PII. For example, suppose the LLM predicted the following:

- <human_timestamp_(00:02.85 - 00:02.95)> - NON-PHI/PII
- <human_timestamp_(00:07.42 - 00:07.57)> - PHI/PII

Then the final set of detected PHI/PII timestamps—extracted via Auto De-ID and combined by the LLM—guided additional timestamps is used for muting the corresponding sections in the final redacted audio. The final audio will correspond to the following:

“The patient visited [MUTED] last week [MUTED] a follow-up in his clinic at [MUTED]. They discussed medication changes and scheduled the next appointment for next month. The patient also mentioned feeling unwell over the weekend.”

A.6 Auto Relexicalization Algorithm

Auto Relexicalization clusters fact entities using M_{cluster} and retrieves candidate replacements from an index using vector search:

$$R_i = M_{\text{search}}(Q_i, I)$$

If the decision model M_{decision} rejects the match, a new replacement R_{new} is generated using a replacement model:

$$R_{\text{new}} = M_{\text{replace}}(Q_i, T)$$

Algorithm 3 Auto Relexicalization Algorithm

Require: Text T , fact dictionary \mathcal{D} , index I , clustering M_{cluster} , retrieval M_{search} , decision M_{decision} , replacement model M_{replace}

Ensure: Relexicalized text \hat{T}

- 1: Cluster entities: $C_{\mathcal{D}} \leftarrow M_{\text{cluster}}(T, \mathcal{D})$
 - 2: **for** each cluster C_i **do**
 - 3: Generate query Q_i and retrieve match $R_i \leftarrow M_{\text{search}}(Q_i, I)$
 - 4: **if** Valid replacement $M_{\text{decision}}(Q_i, R_i, T)$ **then**
 - 5: Use R_i
 - 6: **else**
 - 7: Generate new replacement $R_{\text{new}} \leftarrow M_{\text{replace}}(Q_i, T)$
 - 8: Ingest new replacement and original entity into index I
 - 9: Store R_{new} for final relexicalization
 - 10: **end if**
 - 11: **end for**
 - 12: Apply adjustments and replace entities. **return** \hat{T}
-

This new replacement, along with the original entity, is then ingested into the index I to ensure consistency across documents. The final text \hat{T} is formed after replacing entities accordingly.

A.7 Proposed Metrics for Relexicalization

We propose a set of evaluation metrics to assess the effectiveness of re-lexicalization in preserving entity roles, maintaining contextual coherence, ensuring replacement consistency, and minimizing unintended biases in clinical models.

Entity Preservation Rate evaluates whether the re-lexicalized entity retains its semantic role and contextual attributes. Higher scores indicate better preservation. For instance, in the sentence “*Dr. Emily Carter is a cardiologist at St. Mary’s Hospital,*” a poor re-lexicalization would be “*Alice is a teacher at Westwood Academy,*” which alters both the profession and institution type. A good re-lexicalization would be “*Dr. Kevin Chang is a cardiologist at Lincoln Medical Center,*” as it preserves the entity’s role and contextual relevance.

Contextual Coherence Score measures whether the re-lexicalized entity integrates naturally within the surrounding text without disrupting fluency or meaning. For example, in the original sentence “*John met his lawyer, Mr. Anderson, at the firm,*” a

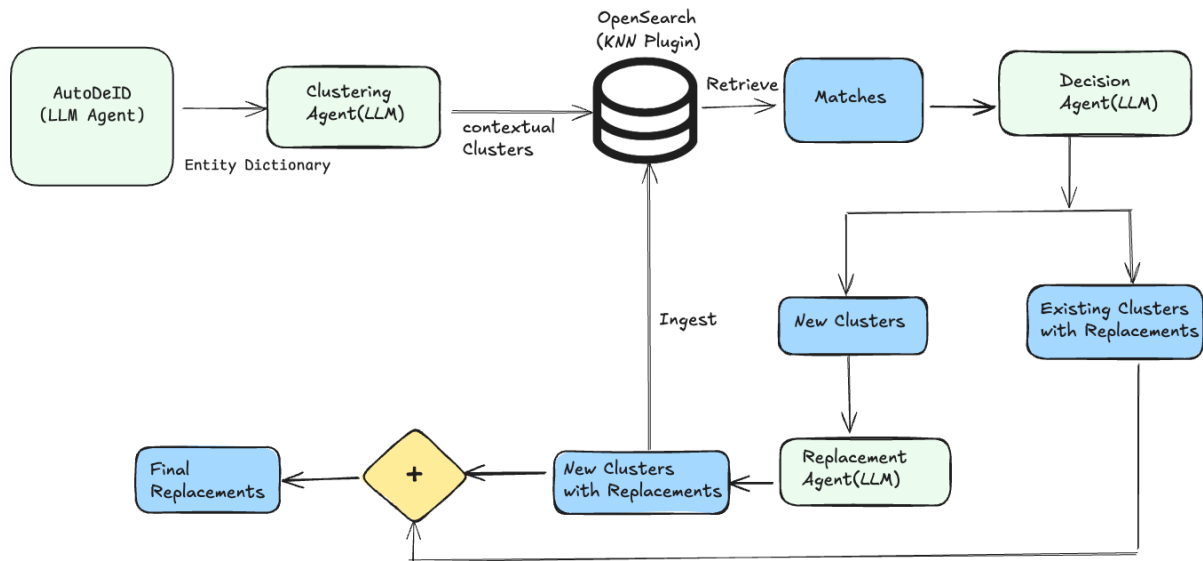


Figure 4: Relexicalization Workflow Diagram

poor substitution would be “*Harry met his lawyer, Pizza Hut, at the firm*”, introducing a semantic inconsistency. A more appropriate replacement would be “*Harry met his lawyer, Mr. Bennett, at the firm*”, maintaining contextual coherence.

Replacement Consistency Score ensures that an entity is consistently replaced across multiple documents, preserving identity coherence. For instance, in “*Dr. Emily Carter attended the surgery*”, a conflicting replacement in another document such as “*Dr. Jennifer Smith is a cardiologist*” introduces inconsistency. A high score indicates that the same entity is replaced uniformly across contexts.

Clinical Model Consistency assesses whether relexicalized data, when used in clinical decision-making models, avoids introducing biases related to race, ethnicity, region, or age group. If a model trained on real data produces a metric value X , a poor re-lexicalization may yield a metric of $X + \delta x$, where δx is significantly large, indicating a deviation from real-world behavior. An optimal re-lexicalization ensures that the metric shift remains marginal, preserving the integrity of the clinical model.

A.8 Integration of RedactOR with Oracle Health Clinical AI System

Our RedactOR framework is integrated into Oracle Health Clinical AI system to facilitate the privacy-preserving processing of longitudinal EHRs (including SOAP notes (Podder et al., 2024)), ambient intelligence data, and conversational AI outputs. The system operates autonomously, as the Production Environment is inaccessible to any user group, ensuring a fully automated pipeline. A dedicated worker module, running continuously, monitors the DataSink, retrieves unprocessed files, submits them to the RedactOR service for PHI/PII removal, and securely transfers the de-identified records to the Data Science Lab Environment. A one-way policy enforces strict data flow control, guaranteeing that only de-identified data is accessible within a secure research environment, where authorized users interact with it via an SSH-secured virtual machine, preserving both data integrity and analytical utility.

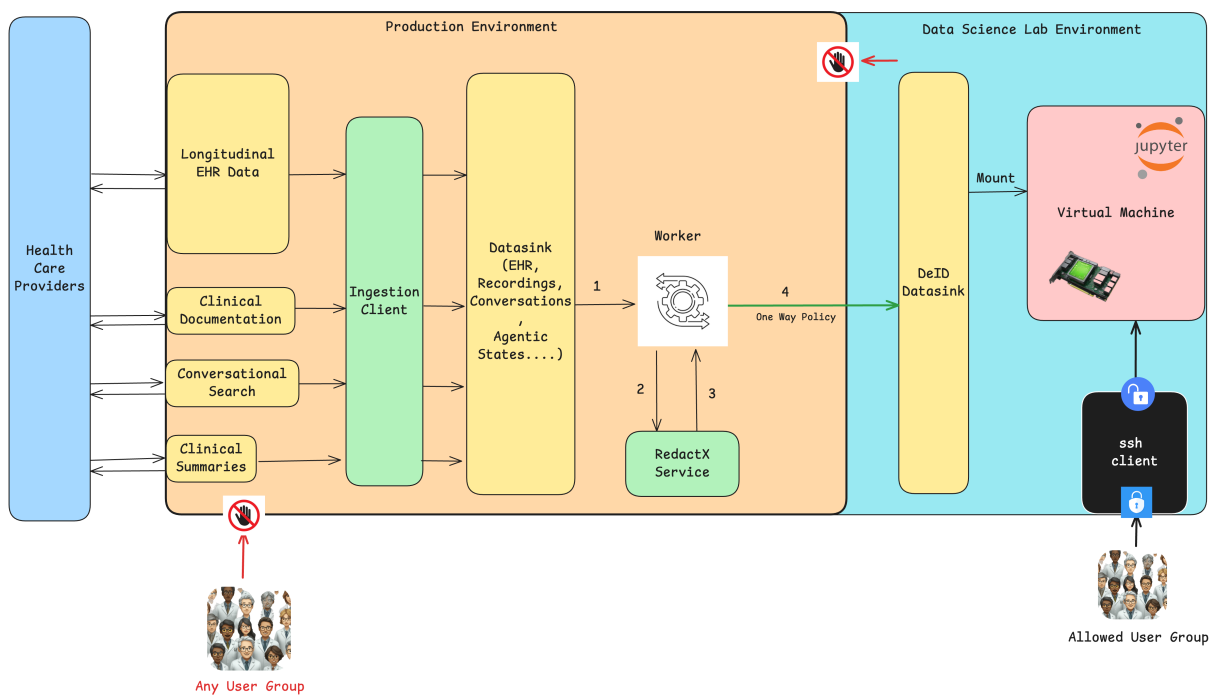


Figure 5: Diagram explaining the integration of our RedactOR framework with Oracle Health Clinical AI System

A.9 End to End Example Illustrating Text De-identification and Relexicalization using RedactOR

A.9.1 Schema for the Input Record

Listing 1: Sample Schema Definition

```
schema_definition = {
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "recordVersion": "1.0",
  "description": "Schema for a free text record",
  "dataType": "clinicalRecord"
  "properties": {
    "PatientId": {
      "type": "string",
      "description": "Patient ID.",
      "autoDeId": False,
      "shouldMask": False,
      "shouldHash": True,
      "entity_type": None
    },
    "MRN": {
      "type": "string",
      "description": "MRN of the Patient",
      "autoDeId": False,
      "shouldMask": False,
      "shouldHash": True,
      "entity_type": None
    },
    "AGE": {
      "type": "string",
      "description": "Age of the Patient",
      "autoDeId": False,
      "shouldMask": True,
      "shouldHash": False,
      "entity_type": "[AGE]"
    },
    "note": {
      "type": "string",
      "description": "Clinical Note",
      "autoDeId": True,
      "shouldMask": False,
      "shouldHash": False,
      "entity_type": None
    }
  }
}
```

A.9.2 Input Record

Listing 2: Sample record with PHI/PII

```
{
  "PatientId": "123456789",
  "MRN": "A987654321",
  "AGE": "45 years",
  "note": "John Doe, a 45-year-old male, presented to Stanford Medical Center on 03/16/2025 complaining of severe abdominal pain. He was referred by Dr. Emily Smith from Valley Health Clinic. His address is 123 Main St, Palo Alto, CA 94301. Contact number: (650) 555-1234. Past medical history includes hypertension and Type 2 diabetes. His insurance ID is INS-789456123. The patients wife, Jane Doe, can be reached at (650) 555-5678. A CT scan was performed and results were discussed with the patient at 2:00 PM. Follow-up scheduled on 03/22/2025 at 9:00 AM. Patient is employed as a software engineer at TechNova Corp. Social Security Number: 987-65-4321."
}
```


A.9.3 Entities Extracted by Auto De-ID

Listing 3: Extracted Entities

```
{
  "PERSON": [
    "Jane Doe",
    "Emily Smith",
    "John Doe"
  ],
  "ADDRESS": [
    "123 Main St, Palo Alto, CA 94301"
  ],
  "AGE": [
    "45 years"
    "45-year-old"
  ],
  "LOCATION": [
    "Palo Alto"
  ],
  "MARITAL_STATUS": [
    "wife"
  ],
  "PARENTHOOD": [],
  "OCCUPATION": [
    "software engineer"
  ],
  "BIRTH_DATE_TIME": [],
  "SSN_OR_TAXPAYER": [
    "987-65-4321"
  ],
  "EMAIL": [],
  "FIN": [
    "INS-789456123"
  ],
  "GUID": [
    "987-65-4321"
  ],
  "ORGANIZATION": [
    "TechNova Corp",
    "Stanford Medical Center",
    "Valley Health Clinic"
  ],
  "PHARMACY": [],
  "DIAGNOSTIC_LABS": []
}
```

A.9.4 DeID-Only Output

Listing 4: DeID-Only Output

```
{
  "PatientId": "HASHED_VALUE",
  "MRN": "HASHED_VALUE",
  "AGE": "[AGE]",
  "note": "[PERSON], a [AGE] male, presented to [ORGANIZATION] on 03/16/2025
    complaining of severe abdominal pain. He was referred by Dr. [PERSON] from [
    ORGANIZATION]. His address is [ADDRESS]. Contact number: [TELEPHONE_NUMBER].
    Past medical history includes hypertension and Type 2 diabetes. His insurance
    ID is [FIN]. The patients [MARITAL_STATUS], [PERSON], can be reached at [
    TELEPHONE_NUMBER]. A CT scan was performed and results were discussed with the
    patient at 2:00 PM. Follow-up scheduled on 03/22/2025 at 9:00 AM. Patient is
    employed as a [OCCUPATION] at [ORGANIZATION]. Social Security Number: [GUID]."
```

A.9.5 De-ID + Relexicalization Output

Listing 5: DeID+Relexicalization Output

```
{
  "PatientId": "HASHED_VALUE",
  "MRN": "HASHED_VALUE",
  "AGE": "mid-forties",
  "note": "Michael Johnson, a mid-forties male, presented to Harvard Medical Center
    on 03/16/2025 complaining of severe abdominal pain. He was referred by Dr.
    Sophia Brown from Green Valley Clinic. His address is 456 Elm St, Mountain
    View, CA 94041. Contact number: (123) 274-0846. Past medical history includes
    hypertension and Type 2 diabetes. His insurance ID is INS-123456789. The
    patients spouse, Alice Johnson, can be reached at (123) 274-6354. A CT scan
    was performed and results were discussed with the patient at 2:00 PM. Follow-
    up scheduled on 03/22/2025 at 9:00 AM. Patient is employed as a data scientist
    at Innovatech Inc.. Social Security Number: 123-45-6789."
```

B High Level Prompt Templates for Different LLM Components

Due to compliance, privacy, and business confidentiality considerations, we do not release the exact prompts used for each LLM component in RedactOR. Instead, to foster reproducibility and enable community adaptation, we provide high-level prompt templates that capture the *structure, intent, and output format* of each prompt while omitting sensitive implementation details.

These templates define:

- The **role of the LLM** in each component (e.g., entity extraction, clustering, relexicalization),
- The **core task description and guidelines** for execution,
- The **expected output schema** in JSON format for integration and evaluation,
- **Placeholders** for data inputs, reference context, and special parameters (e.g., shift values, entity-specific rules).

By offering these templates, we enable researchers and practitioners to develop specialized prompts tailored to their own datasets, privacy policies, and LLM configurations, while ensuring compatibility with our overall system architecture.

B.1 Auto De-ID LLM Component

```
prompt: |
{{role}} # High-level role of the LLM.
Describe that it acts as a De-Identification Specialist tasked with
extracting PHI/PII from clinical text while adhering to legal privacy regulations
(e.g., HIPAA). Include general expectations on accuracy and coverage.

{{entity_extraction_guidelines}} # Instructions on how to treat text
(e.g., case sensitivity, contractions), exclusions (e.g., medication names, diagnoses),
and how to handle special cases. Mention precision requirements for each entity type
and the need to distinguish similar types (e.g., ADDRESS vs LOCATION).

{{controls}} # List and description of specific PHI/PII categories to be extracted
(e.g., names, dates, contact info, IDs, financial details, technical identifiers,
demographic information). This section should align with the entity types and provide
guidance on inclusion/exclusion criteria.

{{context_awareness}} # Explain the need for context-sensitive entity extraction to
avoid over-redaction. Describe how identical strings may appear in different contexts
with different meanings and the importance of using surrounding context to correctly
identify which instance to redact. Highlight that the model must associate each extracted
entity with its specific textual occurrence based on context, not just string matching.

The output format should strictly just be a JSON dictionary with the entity mentioned
above as the key and its list of words/phrases found in the text as its value.

For eg., "ENTITY": ["A", "B", "C"]

You must not add any key which is not a part of the guidelines above. You must add all
the entity as the keys in the output even if the value list for that is empty.

The final output format must look like as follows. You must not produce anything
except the json output. Ensure the output can be parsed by Python json.loads

{
  <entity_type1>: <list_of_words_or_phrases_for_entity_type1>,
  <entity_type2>: <list_of_words_or_phrases_for_entity_type2>,
  ... and so on
}

{{self_checklist}} # List of validation checks the LLM must perform before returning output.
For example, ensure PERSON doesn't include pronouns, validate that BIRTH_DATE_TIME only includes
birth dates, and confirm all keys in output match allowed entity types.

Here is the input text:
{{input_text}} # Placeholder for the input clinical text to be de-identified.
This is the text the LLM will process.
```

B.2 Relexicalizer Components

```
clustering_prompt: |
  {{role}} # Describe the clustering task: grouping contextually identical entities from two medical documents.

  {{task_overview}} # Outline the goal: assign consistent identifiers to contextually similar entities across documents.

  {{guidelines}} # Provide detailed guidelines: consistency, identifier format, handling of subnames, JSON format compliance.

  {{example_input_output}} # Include sample input/output structure for clarity (optional for template use).

  {{input_dict_placeholder}} # Placeholder for the input dictionary of entities.

  {{reference_text_placeholder}} # Placeholder for the contextual text reference.

  Output Format:
  ```json
 {
 "ENTITY_TYPE": {
 "ENTITY_TYPE_1": ["entity_variant_1", "entity_variant_2"],
 "ENTITY_TYPE_2": ["entity_variant_3"]
 },
 ...
 }
}

query_prompt: |
 {{role}} # Describe task: generate semantic query strings for each entity cluster to retrieve similar entities.

 {{guidelines}} # Explain how to synthesize cluster information into a concise semantic query.

 {{example_input_output}} # Provide example input and expected query outputs for context (optional for template use).

 {{cluster_placeholder}} # Placeholder for the input clusters.

 {{reference_text_placeholder}} # Placeholder for reference context.

 Output Format:
  ```json
  {
    "ENTITY_TYPE": {
      "ENTITY_TYPE_1": "query_string_for_entity_1",
      "ENTITY_TYPE_2": "query_string_for_entity_2"
    },
    ...
  }
}

decision_prompt: |
  {{role}} # Describe task: evaluate semantic similarity of search results to query entities.

  {{guidelines}} # Provide detailed decision rules for matching: exact match, partial, cultural context, ambiguity.

  {{constraints}} # State output constraints: JSON format, result length consistency, no assumptions.

  {{example_input_output}} # Include examples of query, search results, context, and expected Y/N output.

  {{query_placeholder}} # Placeholder for input query.

  {{search_result_placeholder}} # Placeholder for search results.

  {{context_placeholder}} # Placeholder for reference context.

  Output Format:
  ```json
 {
 "result": ["Y", "N", ...]
 }
}

replacement_prompt: |
 {{role}} # Describe task: generate realistic replacement values for each cluster entity.

 {{guidelines}} # Explain general rules for replacement: alignment with entity attributes, format, and uniqueness.

 {{special_rules}} # Detailed per-entity replacement rules for privacy-preserving, contextually realistic generation.

 Output Format:
  ```json
  {
    "ENTITY_TYPE_1": {"replacement": "replacement_value_1", "type": "ENTITY_TYPE"},
    "ENTITY_TYPE_2": {"replacement": "replacement_value_2", "type": "ENTITY_TYPE"},
    ...
  }
  ...

  {{input_cluster_placeholder}} # Placeholder for the entity clusters to replace.
  {{existing_replacements_placeholder}} # Placeholder for existing replacements to avoid duplication.
  {{reference_text_placeholder}} # Placeholder for context to guide realistic replacements.
```

B.3 Audio PHI/PII Timestamps Detector

```
prompt: |
{{role}} # Describe task: infer most likely entity type
for missing audio segments using surrounding transcript context.

{{entity_definitions}} # Provide detailed descriptions of each possible
entity type (e.g., PERSON, AGE, ADDRESS, etc.) and how to identify them from context.

{{transcript_format}} # Describe how transcript data is presented
with timestamps and missing sections.

{{example_input_output}} # Provide an example transcript with missing
timestamps and corresponding expected entity predictions.

{{task_instruction}} # Instruct model to return a JSON dictionary where
keys are timestamps and values are predicted entity types or "UNKNOWN".

Output Format:
```json
{
 "TIMESTAMP1": "ENTITY_TYPE",
 "TIMESTAMP2": "ENTITY_TYPE",
 ...
}
```

{{partial_transcript_placeholder}} # Placeholder for the transcript
input with missing sections.
```

C Entity types supported by RedactOR in Production

- ADDRESS
- SSN_OR_TAXPAYER
- EMAIL
- PASSPORT_NUMBER_US
- TELEPHONE_NUMBER
- DRIVER_ID_US
- BANK_ACCOUNT_NUMBER
- BANK_SWIFT
- BANK_ROUTING
- CREDIT_DEBIT_NUMBER
- MEDICAL_RECORD_NUMBER
- HEALTH_PLAN_ID
- CERTIFICATE_NUMBER
- FIN
- VEHICLE_LICENSE_PLATE_US
- VEHICLE_IDENTIFIER_US
- GUID
- PERSON
- DIAGNOSTIC_LABS
- PHARMACY
- ORGANIZATION
- AGE
- LOCATION
- PARENTHOOD
- MARITAL_STATUS
- OCCUPATION
- RACE
- ETHNICITY
- BIRTH_DATE_TIME
- DEATH_DATE_TIME
- IP_ADDRESS
- URL
- MAC_ADDRESS