

Systematic Evaluation of the Quality of Synthetic Clinical Notes Rephrased by LLMs at Million-Note Scale

Jinghui Liu¹ Sarvesh Soni² Anthony Nguyen¹

¹Australian e-Health Research Centre, CSIRO, Australia

²National Library of Medicine, National Institutes of Health, USA

jinghui.liu@csiro.au

Abstract

Large language models (LLMs) can generate or synthesize clinical text for a wide range of applications, from improving clinical documentation to augmenting clinical text analytics. Yet evaluations typically focus on a narrow aspect – such as similarity or utility comparisons – even though these aspects are complementary and best viewed in parallel. In this study, we aim to conduct a systematic evaluation of LLM-generated clinical text, which includes intrinsic, extrinsic, and factuality evaluations of synthetic clinical notes rephrased from MIMIC databases at million-note scale. Our analysis demonstrates that synthetic notes preserve core clinical information and predictive utility for coarse-grained tasks despite substantial linguistic changes, but lose fine-grained details for task like ICD coding. We show this loss of detail can be substantially mitigated by rephrasing notes by chunks rather than by the whole note, but at the cost of reduced factual precision under incomplete context. Through fact-checking and error analysis, we further find that synthesis errors are dominated by misinterpretation of clinical context, alongside temporal confusion, measurement errors, and fabricated claims. Finally, we show that the synthetic notes – despite their task-agnostic nature – can effectively augment task-specific training for rare ICD codes.

1 Introduction

Large language models (LLMs) have demonstrated impressive capabilities in generating human-like text, including text written by domain experts (Shao et al., 2024). In healthcare, studies have shown that LLMs can generate clinical notes similar to those drafted by clinicians (Van Veen et al., 2024; Small et al., 2025; Kahl et al., 2026), driving interest in LLM-assisted clinical note-taking to reduce the significant burden of documentation (Roberts, 2024; Colicchio et al., 2019). The flexible mechanism of generative LLMs also naturally lends them

to synthetic data generation, where LLMs can be prompted to produce task-specific, synthetic clinical texts to help provide larger training samples for improving model performance and fairness in healthcare (Juwara et al., 2024; Ktena et al., 2024), especially since real-world, private patient records are often difficult to access.

Given the growing adoption of LLMs in clinical text generation, it is essential to measure the quality of the generated text. Studies on LLM-generated clinical documentation typically compare the generated notes with human-written reference notes via an array of textual similarity metrics (Croxford et al., 2025; Carandang et al., 2025), which measure the intrinsic quality of the LLM generation. For studies using LLMs for data synthesis, the evaluation is often performed based on a specific downstream task (Loni et al., 2025; Alshaikhdeeb et al., 2025), which focuses on the extrinsic utility of the generated content. While exceptions exist (Wang et al., 2024b), these different evaluation types are rarely presented jointly on a consistent set of synthetic clinical notes. Since existing metrics typically capture only a limited aspect of textual quality, a systematic evaluation combining multiple strategies, spanning multiple metrics and downstream tasks, is vital to understand how LLMs generate clinical text.

In this study, we aim to address this gap through a multi-angle evaluation of LLM-generated clinical notes, including 1) **similarity-based intrinsic evaluation**, 2) **utility-based extrinsic evaluation**, 3) **fairness evaluation in downstream applications**, 4) **fact-checking with LLM-as-judge**, and 5) **human analysis of LLM errors**. We apply these setups to clinical notes rephrased by LLMs, which we treat as synthetic notes.

Rephrasing can be viewed as the most basic form of text generation given that LLMs are provided with complete input contexts, and serves as a lean test case of whether LLMs can interpret clinical in-

formation in its original form. Furthermore, LLMs are considered particularly proficient in rephrasing noisy input into high-quality text, a technique now used widely in curating large-scale web corpora (Maini et al., 2025; Li et al., 2024). Intuitively, LLMs should be able to rephrase clinical notes with ease, but our findings reveal that significant information loss can occur, especially when rephrasing whole notes compared to smaller chunks — a finding demonstrated and cross-checked across our evaluation setups.

Additionally, we examine generation scenarios with varied prompts and repeated generation, and discuss their implications. Although the rephrased synthetic notes are task-agnostic, we show they can be valuable resources for data augmentation in a task-specific setting — demonstrated through ICD prediction (Ji et al., 2024; Gan et al., 2025). The benefits are particularly substantial for modeling rare ICD codes. These experiments shed light on the strengths and limitations of LLM-generated clinical notes.

2 Related Work

Evaluation is key to driving the development of models in NLP. In the clinical domain, proper evaluation is critical to ensure an NLP system is fit-for-purpose and can be applied in healthcare, where both accuracy and safety are critical. For LLM-generated text, existing works usually evaluate it by comparing to the human-written counterparts. For example, studies on generating clinical documentation using LLMs (Chung et al., 2025; Soni and Demner-Fushman, 2024) — such as drafting a discharge summary based on patient records (Schwieger et al., 2024) or writing a report from a medical image (Yu et al., 2023; Nicolson et al., 2024) — typically compare the generated notes with ground-truth notes written by human experts. This often involves multiple textual similarity metrics, such as ROUGE (Lin, 2004) and BERTScore (Zhang* et al., 2020). Crafting pairwise similarity metrics is a research field in itself (Croxford et al., 2025).

Another application of LLM-based text generation is clinical text synthesis (Mitra et al., 2025; Liu et al., 2024a; Li et al., 2023), which provides data augmentation for a given task. These studies typically focus on particular note types or labeled datasets, and perform evaluation based on the downstream task of interest (Tang et al., 2023; Soni

and Roberts, 2020), with task accuracy serving as the extrinsic measure. Unlike the aforementioned works, leveraging downstream tasks sidesteps the need for expert-written reference notes that are costly to produce at scale.

It is known that many existing metrics for text evaluation are imperfect (Sai et al., 2022) and only capture a specific aspect of textual features. Therefore, it is common to aggregate various metrics when comparing and ranking the performance of different generation models, and such aggregation tends to correlate better with human judgment (Xu et al., 2024; Van Veen et al., 2024). Furthermore, combining and analyzing different evaluation approaches can better illuminate generation quality by allowing results to be compared and cross-checked. In this study, we conduct such a systematic, multi-angle evaluation on a consistent set of synthetic notes generated at large scale — a combination that, to our knowledge, is still lacking in prior work.

Clinical text is noisy and varies in quality, a characteristic that resembles web data. Curating web data is an important research direction given its key role in training language models (Li et al., 2024), and this process can benefit from LLM-based synthesis. LLMs can be instructed to generate content with high “educational value” (Penedo et al., 2024; Gunasekar et al., 2023) or to produce improved versions of existing text snippets (Maini et al., 2024; Kang et al., 2025). The latter is also known as *rephrasing*. It has been found that even small-sized LLMs (<8B) are capable of producing high-quality rephrased corpora for improved and more efficient pretraining, and that larger LLMs do not necessarily improve downstream pretraining quality (Kang et al., 2025; Maini et al., 2024). This method has also been shown effective in the medical domain (Liu and Nguyen, 2024) for domain-specific pretraining compared to text synthesis from patient profiles alone (Kweon et al., 2024). Given its effectiveness, general applicability, and natural pairing with human-written notes for intrinsic evaluation, we adopt this method to generate synthetic clinical notes for our multi-angle evaluation.

3 Synthetic Clinical Notes by Rephrasing the MIMIC Databases

We generate synthetic clinical notes using LLMs by grounding on human-written notes, which are sourced from the MIMIC databases. We describe the synthesis process in the following.

3.1 Data Preparation

The MIMIC databases provide de-identified, real-world clinical data across various modalities. We considered all clinical notes from MIMIC-III (v1.4) (Johnson et al., 2016) and MIMIC-IV (v2.2) (Johnson et al., 2023). MIMIC-IV includes two note types (i.e., discharge summary and radiology report) while MIMIC-III provides a broader range of documentation charted during hospital encounters (e.g., physician and nursing notes). We align the format of the de-identification placeholders to underscores (“___”), consistent with MIMIC-IV. After deduplicating notes with the exact same content, the final pool of MIMIC notes comprises 4.6M clinical notes with 1.4B words.

3.2 Rephrasing Setup

Two approaches were considered in rephrasing clinical notes. The first is to rephrase *by the entire note*, preserving the original context. The second splits long notes into chunks and rephrases them *by chunk*. This is motivated by findings that LLMs may struggle to process information as the input context grows longer (Liu et al., 2024b; Hsieh et al., 2024), including during rephrasing (Maini et al., 2024). Each chunk contains approximately 150 words with sentence boundaries maintained, resulting in an average of 2.3 chunks per MIMIC note.

Small-scale generative LLMs are capable of effective rephrasing while maintaining efficiency (Kang et al., 2025; Maini et al., 2024). Following this approach, we use three ~8B models to rephrase MIMIC notes into “high-quality” text: Llama-3.1 (8B) (Dubey et al., 2024), Mistral-v0.3 (7B) (Jiang et al., 2023a), and Qwen-2.5 (7B) (Qwen et al., 2024). These efficient open-source LLMs also represent a realistic deployment choice for privacy- and resource-constrained clinical environments. Prompt templates, decoding parameters, and other inference details are provided in Appendix A. The complete process required about 1.8K GPU hours on NVIDIA H100 with vLLM (Kwon et al., 2023), producing a **9.1B-word** corpus, which we denote as MIMIC-EXT-SYNNOTES (details in Appendix Table 4).

4 Intrinsic Evaluation and Comparison with Human-written Notes

4.1 Experimental Setup

For intrinsic evaluation, we analyzed and compared surface-level and linguistic features between syn-

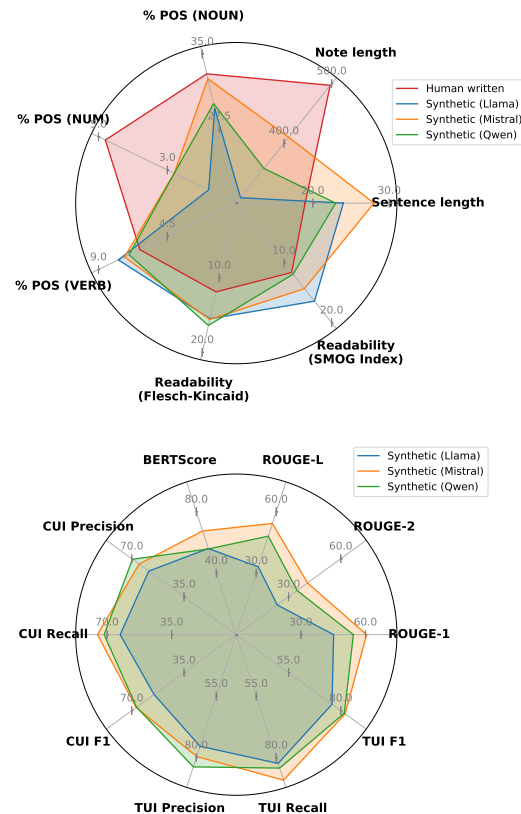


Figure 1: Linguistic features of synthetic notes and their textual similarities to human-written notes.

thetic and human-written notes. We computed sentence and note lengths, part-of-speech (POS) distributions (for nouns, verbs, and numerals), and two standard readability scores: Flesch-Kincaid Grade Level (Solnyshkina et al., 2017) and SMOG Index (Harry and Laughlin, 1969).

For pairwise metrics to capture the textual similarity between human-written and synthetic notes, we computed ROUGE (Lin, 2004) and BERTScore (Zhang* et al., 2020). Due to token length limitations in many pre-trained BERTScore base models (typically capped at 512 tokens), we opted for a DeBERTa-based model (He et al., 2023) with flexible input length and adaptation to natural language inference, which offers enhanced capability for semantic similarity comparison.

To quantify the semantic fidelity of LLM-generated notes in terms of medical concepts, we followed previous works (Gao et al., 2022; Xu et al., 2024) to extract UMLS Concept Unique Identifiers (CUIs) using QuickUMLS (Soldaini, 2016), and computed precision, recall, and F1 scores using the original notes as the reference, as proposed in MedCON (Yim et al., 2023). CUIs capture unique medical concepts from UMLS as a bridge




Figure 2: Comparing synthetic notes with human-written notes in downstream modeling. Each metric is calculated as the mean of five runs, and the error bars represent 95% confidence intervals.

for shared medical semantics. Following MedCON, we only considered clinically relevant semantic groups, such as anatomy, chemicals & drugs, and disorders. For a high-level comparison, we also mapped CUIs to their medical semantic types (Type Unique Identifiers, or TUIs) and computed the corresponding scores.

To facilitate efficient evaluation, we sampled 1K discharge summaries and 5K other note types from MIMIC-IV and MIMIC-III, respectively, as we observed that using more notes did not yield significantly different results.

4.2 Results

Figure 1 shows the comparison between human-written and LLM-generated notes. The synthetic notes are noticeably shorter in length but contain longer sentences with higher readability scores. They also show a **significant shift in lexical and semantic distributions**, as reflected by textual similarity scores with different degrees of variation across LLMs.

Analysis of TUI distributions (Table 1) reveals that the synthetic notes tend to omit concepts related to “*Body Space or Junction*” and “*Medical Device*” while introducing more from “*Natural Phenomenon or Process*” and “*Clinical Attribute*”, suggesting that synthetic text is **more descriptive but less dense with anatomical and specialist terminology**. This shows that LLMs may have a tendency to introduce more general terms and phrases in the generated text. The lists of frequent errors in semantic types made by each LLM are presented in Appendix Table 6.

	Semantic Type	%
Added	Natural Phenomenon or Process	16.5%
	Clinical Attribute	13.2%
	Phenomenon or Process	9.8%
	Pharmacologic Substance	9.5%
	Organ or Tissue Function	8.2%
Omitted	Body Space or Junction	12.2%
	Mental Process	11.8%
	Medical Device	11.6%
	Gene or Genome	11.2%
	Cell Component	8.8%

Table 1: Top 5 semantic types added (false positives) or omitted (false negatives) in synthetic notes, aggregated across three LLMs.

5 Extrinsic Evaluation with Healthcare Applications

5.1 Experimental Setup

We compared the utility of synthetic notes with human-written notes by training downstream clinical NLP models as a form of extrinsic evaluation. For a patient in a given task cohort, the input notes for the downstream model are either synthetic or human-written for both training and testing. In other words, we compared “*train synthetic, test synthetic*” (*TSTS*) and “*train real, test real*” (*TRTR*) by only varying the source of notes. The downstream performance is considered as a proxy to measure whether synthetic notes provide similar predictive utility to human-written notes. For this purpose, we did not consider *TSTR* (train synthetic, test real) or *TRTS* (train real, test synthetic), which can be confounded by factors such as writing styles and other features in the textual distribution.¹

¹Section 7.3 presents results for data augmentation, which aligns with *TSTR*.

We considered four clinical NLP tasks that represent key healthcare applications: risk stratification (in-hospital mortality and 30-day hospital readmission prediction) (Jiang et al., 2023b; Liang et al., 2025) and medical code assignment (DRG and ICD-10 prediction) (Wang et al., 2024a; Douglas et al., 2025). On the input side, they involve multiple note types authored by different caregivers at various stages of the care episode: in-hospital mortality (Jiang et al., 2023b; Si and Roberts, 2019) and DRG prediction (Wang et al., 2024a) tasks occur during the hospital episode and thus utilize all available clinical notes charted by the time of prediction (one day after ICU admission), such as nursing notes and radiology reports. The 30-day hospital readmission prediction (Boag et al., 2021; Huang et al., 2019) and ICD coding (Ji et al., 2024; Liu et al., 2022b) tasks occur after the end of hospital stay, and their input is based on the discharge summary alone.

On the output side, these tasks also involve varying complexities and requirements for clinical details: the first three tasks require the knowledge of overall patient status for a single-label classification, whereas the ICD coding task requires a higher level of attention to clinical details for the challenging multi-label classification with 8K ICD-10 codes (Edin et al., 2023). The four task cohorts are constructed from MIMIC with sample sizes ranging from 18K to 122K patients, and require processing of 546K notes. The statistics of the datasets are presented in Appendix Table 5.

For a given downstream task and a source of clinical notes, a classification model was trained, evaluated on the validation set for early stopping, and then tested on the held-out set to report final results. The only variation for a given task was the source of clinical notes, which are either human-written or synthetic, i.e., rephrased by one of the LLMs. Appendix B presents the details of the task cohorts and the training of downstream models.

5.2 Results

Figure 2 shows that synthetic notes exhibit comparable downstream performance to human-written notes on single-label tasks (i.e., all tasks except ICD coding), with results overlapping the 95% confidence intervals of using original notes. This suggests that **key clinical information is preserved in synthetic notes**, enabling accurate modeling and prediction for three prediction tasks. On the other hand, the ICD coding task (8K classes; multi-

label) that requires fine-grained details exhibits much lower performance (with a drop of 4–10% in micro F1 score), suggesting **loss of certain clinical details** during synthesis that are essential for coding. Crucially, rephrasing *by chunks* instead of *by notes* substantially closes this performance gap, indicating **LLMs’ difficulty in retaining and integrating information from long contexts** as a primary driver for detail loss.

5.3 Fairness Evaluation on Risk Stratification

Beyond predictive accuracy, we also assess fairness for two risk stratification tasks (in-hospital mortality and 30-day hospital readmission predictions). Following Brown et al. (2025), we split the patients into subgroups based on age (over and under 50), gender (female and male), and ethnicity (non-white and white), and computed the average absolute odds difference (AAOD) (Hardt et al., 2016) and statistical parity difference (SPD) (Besse et al., 2021) across these subgroups.

Figure 3 presents the results of the fairness evaluation, where we observe no significant changes in the fairness of the models trained on synthetic notes compared to those trained on human-written notes. This indicates that the LLMs did not introduce significant biases in their rephrasing that would affect downstream modeling of risk stratification.

6 Fact-checking with Error Analysis

We measure the truthfulness of the synthetic notes by fact-checking their claims against those made in the original notes. We then analyze the common error patterns exhibited by LLMs.

6.1 Experimental Setup

We follow Munnangi et al. (2025) to decompose synthetic notes into atomic facts (Min et al., 2023) and conduct fact-checking via entailment (Wang et al., 2024c). Following Xie et al. (2024), we calculate *precision* (the proportion of synthetic factual claims entailed by the human-written note) and *recall* (the proportion of reference factual claims entailed by those in the synthetic note).

For a pair of synthetic note N_S and human-written note N_H , they are decomposed into atomic factual claims C_S and C_H . Precision is defined as the proportion of facts in C_S that can be entailed by N_H :

$$Precision = \frac{1}{|C_S|} \sum_{c \in C_S} [N_H \models c] \quad (1)$$

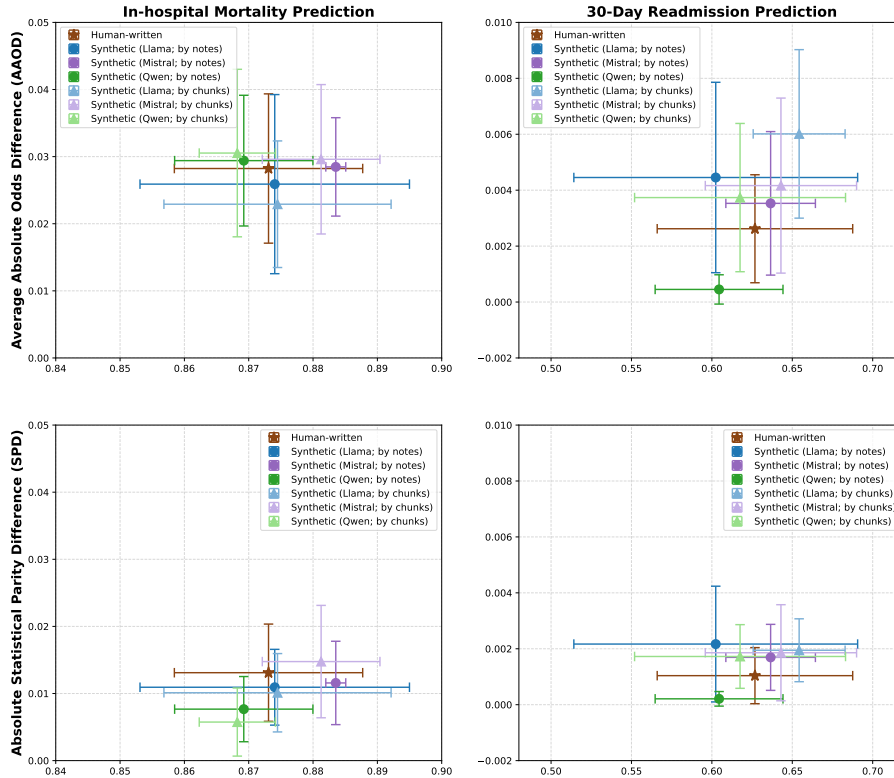


Figure 3: Fairness evaluation of the synthetic notes in downstream modeling in comparison with human-written notes. The x-axis presents model performance (AUROC score) and the y-axis presents model fairness (Average Absolute Odds Difference (AAOD) on the top panel and the absolute Statistical Parity Difference (SPD) on the bottom panel). For fairness metrics, the lower the better. The value of 0 for AAOD and SPD means no model discrimination between the examined subgroups. The results were averaged across age, gender, and ethnicity.

where $N_H \models c$ means note N_H entails fact c . Recall is defined as the proportion of facts in C_H that can be entailed by those from the synthetic notes C_S :

$$Recall = \frac{1}{|C_H|} \sum_{c \in C_H} [C_S \models c] \quad (2)$$

Note that this is different from Munnangi et al. (2025), which used sentences from the reference note instead of decomposed facts; we found the latter to be more interpretable. The final precision and recall are macro averages over the sampled datasets, and F1 is the harmonic mean of the final precision and recall scores. We use Gemini-2.5 to decompose facts and assess entailment due to its strong reasoning capabilities. Due to the cost of querying the closed-source model to provide robust assessment, we conducted this part of the experiments on a small set of sampled notes as exploratory analysis, which covered more than 10K facts. More details are described in Appendix C.

6.2 Results

Table 2 shows that synthetic notes generally exhibit **high precision in their factual claims, but considerably lower recall**, indicating missing details. This becomes more pronounced when rephrasing *long notes* (>2K words), with recall as low as 0.3. Meanwhile, rephrasing *by chunk* produces a substantially larger number of facts, improving recall and suggesting a preservation of clinical details, consistent with findings in Section 5.

Beyond aggregated metrics, one author with a background in medical informatics analyzed sample outputs and identified a taxonomy of common error types (see Appendix Table 7): **misinterpretation of clinical context is the dominant failure mode**, followed by measurement errors, temporal/recency confusion, and fabrication of claims when the documentation is subtle. LLMs sometimes make up names and dates even though these are removed from real notes, and they also struggle with shorthand notation, e.g., confusing blood pressure values for respiration rates.

Finally, the results highlight a trade-off inher-

	Fact Prec	Fact Rec	Fact F1	# Avg. Facts
Llama	0.903	0.587	0.712	68
<i>long notes</i>	0.904	0.309	0.460	101
<i>by chunk</i>	0.868 (↓)	0.794 (↑)	0.830 (↑)	297 (↑)
Mistral	0.972	0.812	0.885	105
<i>long notes</i>	0.959	0.719	0.822	186
<i>by chunk</i>	0.938 (↓)	0.779 (↑)	0.851 (↑)	251 (↑)

Table 2: Fact-checking results based on decomposed atomic facts from synthetic notes.

ent to chunk-level rephrasing: while it maintains more claims and improves recall, it also lowers precision and is thus accompanied by higher hallucination rates. In many of these cases, LLMs see incomplete context and may generate plausible but unsupported content whose reference context appears in another chunk.

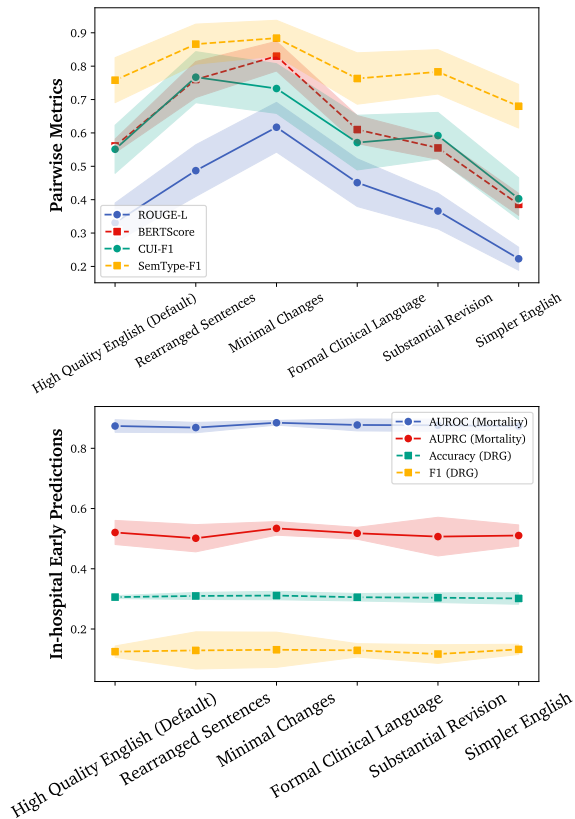


Figure 4: Impact of varied prompt instructions on intrinsic and extrinsic evaluations. Results are averaged over five random seeds.

7 Additional Analyses

We conduct additional analyses to examine the impact of prompting setups on the generated notes, including different instructions and iterative prompting. We also perform a data augmentation experi-

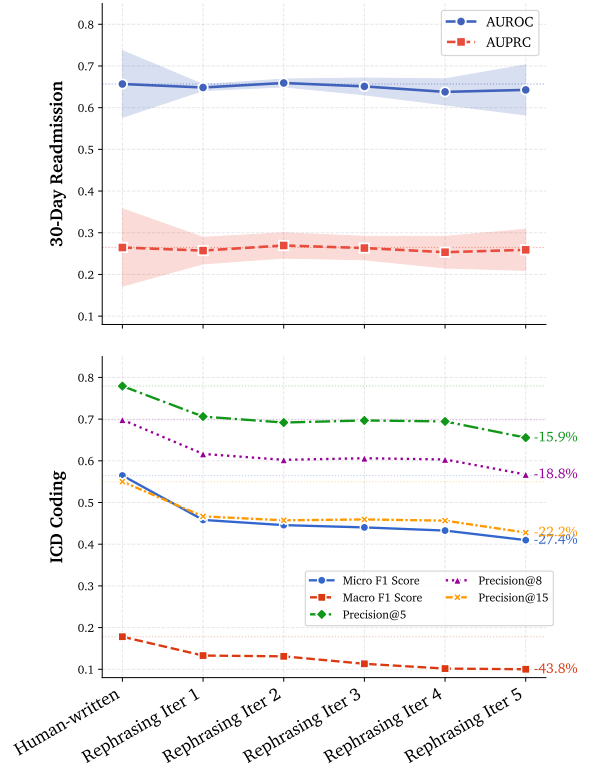


Figure 5: Impact of repeated synthesis on 30-day hospital readmission prediction across 5 iterations. Results are averaged over five random seeds.

ment on the ICD prediction task.

7.1 Impact of Varied Prompts

We examine alternative prompts (Appendix E) that vary in structure (e.g., “rearranging sentences”) and style (e.g., “simplified English”), and observe their impact on synthetic notes, reflected by both intrinsic and extrinsic measurements. As expected, prompt variation significantly affects textual similarity: both syntactic and semantic contents change dramatically from instruction to instruction (Figure 4; top). Interestingly, the default instruction to produce “high-quality” language has a similar effect to “substantial revision”, showing the complexity of handling clinical documentation from the LLM’s perspective.

On the other hand, the change of prompts has minimal impact on downstream modeling (Figure 4; bottom). This indicates that the synthetic notes preserve key clinical information related to patient severity despite differences in linguistic style, consistent with the findings in Section 5.

	All codes		Rare codes	
	micro F1	macro F1	micro F1	macro F1
Human-written notes	0.5644 (0.5622–0.5667)	0.1738 (0.1703–0.1772)	0.0359 (0.0173–0.0596)	0.0134 (0.0070–0.0199)
+ <i>synthetic notes</i>	0.5703 (0.5671–0.5736)	0.2141 (0.2007–0.2278)	0.1035 (0.0595–0.1485)	0.0448 (0.0258–0.0656)

Table 3: Augmentation results for ICD coding. Models are evaluated on the human-written notes in the test set by bootstrapping 1000 samples from multiple seeds. The rare codes are 391 ICD codes that appeared only 10 times in the cohort (the benchmark in Edin et al. (2023) excluded codes that appeared fewer than 10 times). The reported scores are mean values and 95% confidence intervals. All augmented results show significant improvements with p-value <0.01 with two-sided t-test.

7.2 Impact of Repeated Synthesis

We explore repeated synthesis, where the LLM iteratively rephrases its own outputs to examine semantic drift. Shumailov et al. (2024) demonstrates cases where recursive generations can lead to model collapse when training happens in a loop. This raises critical concerns over the quality of data for training models over time, as generated content gradually becomes an integral part of the data pool, especially after deploying LLMs in practice.

Figure 5 shows the results on two downstream tasks over five iterations of repeated synthesis, each based on the output from the previous round. The performance on the single-label readmission prediction task remains stable. For ICD coding, although the performance decreases over time and sometimes by a large margin, the decline is progressive rather than sudden, and we have yet to observe a “collapse” in downstream modeling, i.e., scoring zero accuracy. Note that in this setting, the “generator” (i.e., the LLM) was kept unchanged, and only the downstream models (i.e., task-specific “predictor”) were trained at each iteration instead of the “generator” (Shumailov et al., 2024). Moreover, the provision of a reference serves as a constraint to keep LLMs from generating completely irrelevant content. Nevertheless, the issue of LLMs missing fine-grained clinical details persists, as demonstrated by the declining ICD performance over repeated rephrasing.

7.3 Rephrased Notes for Augmentation

We examine the value of the synthetic notes in augmenting the training set for downstream modeling. Using LLMs to generate synthetic training data by following annotated examples has been studied for various clinical NLP tasks (Tang et al., 2023; Šuvalov et al., 2025). Unlike these studies, which craft application-specific prompts with demonstra-

tions (i.e., few-shot prompting), our synthesis does not consider any task-related instructions or labels in the prompt. In other words, we apply the **task-agnostic synthetic notes for task-specific data augmentation**.

We focus on the challenging ICD-10 coding task and augment the original training set with synthetic notes generated by Mistral (by chunks), and evaluate on the human-written notes. Note that this “*test real*” setting is different from Section 5, where synthetic notes were examined in a *TSTS* setting.

The results are presented in Table 3, where we see that augmentation with synthetic notes increases all F1 scores (significant improvements with p-value <0.01), especially for rare codes (micro F1: 0.03 → 0.10). This shows the additional potential of LLM-synthesized clinical text for data augmentation despite being task-agnostic.

8 Discussion & Conclusion

Existing studies on synthetic clinical text typically evaluate the generation from one angle and lack direct comparisons between these evaluation strategies on a consistent set of clinical notes. We consider these different angles complementary in essence and useful for providing a unified view of the synthetic text quality; thus, we included textual similarity comparison, downstream utility assessment, and factuality checking in this study. As the results show, the synthetic clinical text generated by LLMs has both potential and weaknesses, and its usage should be assessed according to the application context.

Two findings from our study warrant particular attention for practitioners using LLMs for clinical text generation. **The first is the long-context bottleneck.** The substantial ICD coding gap between note-level and chunk-level rephrasing suggests the use of chunking to improve generation

quality. Meanwhile, the removal of surrounding context by chunking could lead to higher rates of hallucination. This means *a trade-off needs to be made on this decision based on the generative application and tolerance of each error type*, and future work needs to analyze in more depth the decisions related to chunking, such as boundary definition and chunk length.

Second, the error analysis shows that LLM failures are not random but exhibit recurring patterns. This suggests the potential for *future work to develop specialized techniques to address these errors*, such as better prompting for temporal anchoring and explicit handling of abbreviations.

We used rephrasing as the generation method in this study, and we did not explicitly distinguish between “rephrasing” and “synthesizing”. Nevertheless, we acknowledge the nuanced difference between them in healthcare, especially the connotation of the term “synthesis” for privacy preservation. Grounding on real notes to create alternative versions of the clinical text — as done in “rephrasing” — forfeits privacy protection. Meanwhile, since the goal of the study is to measure text quality and to shed light on LLMs’ capabilities, we maintain this decision to allow evaluation in different setups, i.e., head-to-head comparison with human-written notes and on downstream tasks.

Finally, our experiments show that the rephrased notes can be instrumental in understanding the quality of LLM generation and valuable for data augmentation despite being task-agnostic. To support future research, we compile and share the rephrased notes in this study as MIMIC-EXT-SYNNOTES, a synthetic corpus with 9B tokens based on 4.6M MIMIC notes.²

Limitations

Choice of LLMs for generation. We only focused on efficient open-source LLMs in our experiments. While more recent models may handle long contexts better, recent work suggests the gap may be narrower than expected: Kang et al. (2025) showed that using larger and more capable models for rephrasing pretraining data does not necessarily translate into better pretraining performance. Our complementary results with Gemma and Mixtral (Appendix F) did not yield substantially different results, supporting this view. Therefore, we con-

sider the current model selection acceptable for fulfilling the study’s objectives. Nevertheless, extending the analysis to models at other scales — both smaller and larger — and considering frontier models in a privacy-preserved setting are important future work.

Methods for text synthesis and evaluation. As mentioned, given the scope of this study, we did not explore alternative approaches for generating synthetic clinical text, particularly reference-free methods. Since MIMIC contains abundant EHR data, future work could incorporate structured measurements in a cross-modal setting to enhance the quality and validity of synthetic text. For evaluation methodologies, we consider the use of models in the current study — such as QuickUMLS for CUI extraction and BERT with sliding window for classification — acceptable as they are consistent for synthetic and human-written notes, but future work needs to further conduct meta-evaluation of the validity and reliability of these methods.

Issues with fact-checking. Although we conducted error analysis using LLM-as-judge for fact-checking, we were unable to directly compare synthetic notes with human-written notes. Such comparison requires extensive time and clinical expertise to review long notes and interpret patient care context, which was infeasible for this study. Instead, we opted to rely on downstream task performance and prior methods from fact-checking (Xie et al., 2024) that demonstrate strong correlation with expert judgment. We also acknowledge the limitations of our current fact-checking setup, which is based on a single family of models and a small number of notes, as we presented it as an exploratory analysis. The choice of model for LLM-as-judge particularly needs further scrutiny, given that models can be sensitive in reported factuality scores (Munnangi et al., 2025). Future work needs to validate the robustness of alternative models — probably the cheaper, open-source options — for this critical task. Nevertheless, we kept the use of models and notes consistent throughout the experiment, so that the analysis of errors and the comparison between note-level and chunk-level rephrasing are still valid.

Ethical Considerations

The generation of synthetic text using LLMs warrants careful ethical considerations. Although our synthesis was based on actual MIMIC notes and

²The synthetic note corpus can be found at <https://github.com/JHLiu7/MIMIC-Ext-SynNotes>.

offered generally high precision of factual claims, we observed that synthetic notes still omit or introduce information, which can be further exacerbated by insufficient context. This highlights the need for cautious use and, in some cases, additional filtering to ensure reliability, depending on the use case. Furthermore, while we assessed potential bias in downstream models for risk prediction, other dimensions of bias might remain and need to be considered in future work.

We assessed the utility of synthetic notes in downstream modeling by comparing “train real, test real” (TRTR) and “train synthetic, test synthetic” (TSTS). Real-world signals (e.g., diagnosis codes) were utilized in both settings, so the TSTS models are applicable to “test real”. Meanwhile, careful consideration and transparency are essential if the synthetic notes were used to develop systems intended for implementation.

We followed the recommended practices for handling MIMIC notes. For rephrasing, we used only local, open-source models, and conducted all experiments in a secure environment to ensure data protection. For fact-checking with Gemini, we adhered to PhysioNet’s Responsible Use guidelines³ by accessing Gemini models via Vertex AI, ensuring compliance with HIPAA regulations.

Acknowledgments

We would like to thank our reviewers for their thoughtful, detailed, and constructive comments that helped to improve this paper. This research was supported in part by the Intramural Research Program of the National Institutes of Health (NIH). The contributions of the NIH author are considered Works of the United States Government. The findings and conclusions presented in this paper are those of the authors and do not necessarily reflect the views of the NIH or the U.S. Department of Health and Human Services. This work was also supported by CSIRO (Australia’s National Science Agency) through the Australian e-Health Research Centre.

References

Basel Alshaikhdeeb, Ahmed Abdelmonem Hemedan, Soumyabrata Ghosh, Irina Balaur, and Venkata Sathagopam. 2025. [Generation of synthetic clinical text: A systematic review](#). *arXiv [cs.CL]*.

³<https://physionet.org/news/post/gpt-responsible-use>

Philippe Besse, Eustasio del Barrio, Paula Gordaliza, Jean-Michel Loubes, and Laurent Risser. 2021. [A survey of bias in machine learning through the prism of statistical parity](#). *The American statistician*, pages 1–11.

William Boag, Olga Kovaleva, Thomas H McCoy, Jr, Anna Rumshisky, Peter Szolovits, and Roy H Perlis. 2021. [Hard for humans, hard for machines: predicting readmission after psychiatric hospitalization using narrative notes](#). *Translational psychiatry*, 11(1):32.

Katherine E Brown, Chao Yan, Zhuohang Li, Xinmeng Zhang, Benjamin X Collins, You Chen, Ellen Wright Clayton, Murat Kantarcioglu, Yevgeniy Vorobeychik, and Bradley A Malin. 2025. [Large language models are less effective at clinical prediction tasks than locally trained machine learning models](#). *Journal of the American Medical Informatics Association: JAMIA*, page ocaf038.

Kristine Ann M Carandang, Jasper Meynard Arana, Ethan Robert Casin, Christopher Monterola, Daniel Stanley Tan, Jesus Felix B Valenzuela, and Christian Alis. 2025. [Are LLMs reliable? an exploration of the reliability of large language models in clinical note generation](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 6: Industry Track)*, pages 1413–1422, Stroudsburg, PA, USA. Association for Computational Linguistics.

Philip Chung, Akshay Swaminathan, Alex J Goodell, Yeasul Kim, S Momsen Reincke, Lichy Han, Ben Deverett, Mohammad Amin Sadeghi, Abdel-Badih Ariss, Marc Ghanem, David Seong, Andrew A Lee, Caitlin E Coombes, Brad Bradshaw, Mahir A Sufian, Hyo Jung Hong, Teresa P Nguyen, Mohammad R Rasouli, Komal Kamra, and 10 others. 2025. [Verifying facts in patient care documents generated by large language models using electronic health records](#). *NEJM AI*, 3(1).

Tiago K Colicchio, James J Cimino, and Guilherme Del Fiol. 2019. [Unintended consequences of nationwide electronic health record adoption: Challenges and opportunities in the Post-Meaningful use era](#). *Journal of medical Internet research*, 21(6):e13313.

Emma Croxford, Yanjun Gao, Nicholas Pellegrino, Karen Wong, Graham Wills, Elliot First, Frank Liao, Cherodeep Goswami, Brian Patterson, and Majid Afshar. 2025. [Current and future state of evaluation of large language models for medical summarization tasks](#). *npj Health Systems*, 2(1):1–13.

James C Douglas, Yidong Gan, Ben Hachey, and Jonathan K Kummerfeld. 2025. [Less is more: Explainable and efficient ICD code prediction with clinical entities](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 30835–30847, Stroudsburg, PA, USA. Association for Computational Linguistics.

- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, and 516 others. 2024. [The llama 3 herd of models](#). *arXiv [cs.AI]*.
- Joakim Edin, Alexander Junge, Jakob D Havtorn, Lasse Borgholt, Maria Maistro, Tuukka Ruotsalo, and Lars Maaløe. 2023. [Automated medical coding on MIMIC-III and MIMIC-IV: A critical review and replicability study](#). In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '23*, pages 2572–2582, New York, NY, USA. Association for Computing Machinery.
- Yidong Gan, Maciej Rybinski, Ben Hachey, and Jonathan K Kummerfeld. 2025. [Aligning AI research with the needs of clinical coding workflows: Eight recommendations based on US data analysis and critical review](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 909–922, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Yanjun Gao, Dmitriy Dligach, Timothy Miller, Dongfang Xu, Matthew M M Churpek, and Majid Afshar. 2022. [Summarizing patients' problems from hospital progress notes using pre-trained sequence-to-sequence models](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 2979–2991.
- Suriya Gunasekar, Yi Zhang, Jyoti Aneja, Caio César Teodoro Mendes, Allie Del Giorno, Sivakanth Gopi, Mojan Javaheripi, Piero Kauffmann, Gustavo de Rosa, Olli Saarikivi, Adil Salim, Shital Shah, Harkirat Singh Behl, Xin Wang, Sébastien Bubeck, Ronen Eldan, Adam Tauman Kalai, Yin Tat Lee, and Yuanzhi Li. 2023. [Textbooks are all you need](#). *arXiv [cs.CL]*.
- Moritz Hardt, Eric Price, and Nathan Srebro. 2016. [Equality of opportunity in supervised learning](#). In *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS'16*, pages 3323–3331, Red Hook, NY, USA. Curran Associates Inc.
- G Harry and Mc Laughlin. 1969. [SMOG grading - a new readability formula](#). *The Journal of Reading*.
- Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2023. [DeBERTaV3: Improving DeBERTa using ELECTRA-style pre-training with gradient-disentangled embedding sharing](#). In *The Eleventh International Conference on Learning Representations*.
- Cheng-Ping Hsieh, Simeng Sun, Samuel Kriman, Shantanu Acharya, Dima Rekesh, Fei Jia, and Boris Ginsburg. 2024. [RULER: What's the real context size of your long-context language models?](#) In *First Conference on Language Modeling*.
- Chao-Wei Huang, Shang-Chi Tsai, and Yun-Nung Chen. 2022. [PLM-ICD: Automatic ICD coding with pre-trained language models](#). In *Proceedings of the 4th Clinical Natural Language Processing Workshop*, pages 10–20, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Kexin Huang, Jaan Altosaar, and Rajesh Ranganath. 2019. [ClinicalBERT: Modeling clinical notes and predicting hospital readmission](#). *arXiv [cs.CL]*.
- Shaoxiong Ji, Xiaobo Li, Wei Sun, Hang Dong, Ara Taalas, Yijia Zhang, Honghan Wu, Esa Pitkänen, and Pekka Marttinen. 2024. [A unified review of deep learning for automated medical coding](#). *ACM computing surveys*.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Léo Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023a. [Mistral 7B](#). *arXiv [cs.CL]*.
- Lavender Yao Jiang, Xujin Chris Liu, Nima Pour Nejatian, Mustafa Nasir-Moin, Duo Wang, Anas Abidin, Kevin Eaton, Howard Antony Riina, Ilya Laufer, Paawan Punjabi, Madeline Miceli, Nora C Kim, Cordelia Orillac, Zane Schnurman, Christopher Livia, Hannah Weiss, David Kurland, Sean Neifert, Yosef Dastagirzada, and 9 others. 2023b. [Health system-scale language models are all-purpose prediction engines](#). *Nature*, 619(7969):357–362.
- Alistair E W Johnson, Lucas Bulgarelli, Lu Shen, Alvin Gayles, Ayad Shammout, Steven Horng, Tom J Pollard, Benjamin Moody, Brian Gow, Li-Wei H Lehman, Leo A Celi, and Roger G Mark. 2023. [MIMIC-IV, a freely accessible electronic health record dataset](#). *Scientific data*, 10(1):1.
- Alistair E W Johnson, Tom J Pollard, Lu Shen, Li-Wei H Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. 2016. [MIMIC-III, a freely accessible critical care database](#). *Scientific data*, 3:160035.
- Lamin Juwara, Alaa El-Hussuna, and Khaled El Emam. 2024. [An evaluation of synthetic data augmentation for mitigating covariate bias in health data](#). *Patterns (New York, N.Y.)*, 5(4):100946.
- Nicolas M Kahl, Marshall J Frieden, Zach R Pope, Marlene M Millen, Vaishal M Tolia, Theodore C Chan, Christopher A Longhurst, Karandeep Singh, and Alan X You. 2026. [Evaluation of electronic health record-integrated artificial intelligence chart review](#). *Npj Health Systems*, 3(1):6.

- Feiyang Kang, Newsha Ardalani, Michael Kuchnik, Youssef Emad, Mostafa Elhoushi, Shubhabrata Sen Gupta, Shang-Wen Li, Ramya Raghavendra, Ruoxi Jia, and Carole-Jean Wu. 2025. [Demystifying synthetic data in LLM pre-training: A systematic study of scaling laws, benefits, and pitfalls](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 10750–10769, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Ira Ktena, Olivia Wiles, Isabela Albuquerque, Sylvestre-Alvise Rebuffi, Ryutaro Tanno, Abhijit Guha Roy, Shekoofeh Azizi, Danielle Belgrave, Pushmeet Kohli, Taylan Cemgil, Alan Karthikesalingam, and Sven Gowal. 2024. [Generative models improve fairness of medical classifiers under distribution shifts](#). *Nature Medicine*, 30(4):1166–1173.
- Sunjun Kweon, Junu Kim, Jiyoun Kim, Sujeong Im, Eunbyeol Cho, Seongsu Bae, Jungwoo Oh, Gyubok Lee, Jong Hak Moon, Seng Chan You, Seungjin Baek, Chang Hoon Han, Yoon Bin Jung, Yohan Jo, and Edward Choi. 2024. [Publicly shareable clinical large language model built on synthetic clinical notes](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 5148–5168.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. 2023. [Efficient memory management for large language model serving with PagedAttention](#). In *Proceedings of the 29th Symposium on Operating Systems Principles*, pages 611–626, New York, NY, USA. ACM.
- Patrick Lewis, Myle Ott, Jingfei Du, and Veselin Stoyanov. 2020. [Pretrained language models for biomedical and clinical tasks: Understanding and extending the State-of-the-Art](#). In *Proceedings of the 3rd Clinical Natural Language Processing Workshop*, pages 146–157, Online. Association for Computational Linguistics.
- Jeffrey Li, Alex Fang, Georgios Smyrnis, Maor Ivgi, Matt Jordan, Samir Yitzhak Gadre, Hritik Bansal, Etash Kumar Guha, Sedrick Keh, Kushal Arora, Saurabh Garg, Rui Xin, Niklas Muennighoff, Reinhard Heckel, Jean Mercat, Mayee F Chen, Suchin Gururangan, Mitchell Wortsman, Alon Albalak, and 40 others. 2024. [DataComp-LM: In search of the next generation of training sets for language models](#). In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Rumeng Li, Xun Wang, and Hong Yu. 2023. [Two directions for clinical data generation with large language models: Data-to-label and label-to-data](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 7129–7143, Singapore. Association for Computational Linguistics.
- Zihan Liang, Ziwen Pan, and Ruoxuan Xiong. 2025. [Causal representation learning from multimodal clinical records under non-random modality missingness](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 28779–28796, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Jinghui Liu, Daniel Capurro, Anthony Nguyen, and Karin Verspoor. 2022a. [“note bloat” impacts deep learning-based NLP models for clinical prediction tasks](#). *Journal of biomedical informatics*, 133:104149.
- Jinghui Liu, Bevan Koopman, Nathan J Brown, Kevin Chu, and Anthony Nguyen. 2024a. [Generating synthetic clinical text with local large language models to identify misdiagnosed limb fractures in radiology reports](#). *Artificial intelligence in medicine*, 159(103027):103027.
- Jinghui Liu and Anthony Nguyen. 2024. [Rephrasing electronic health records for pretraining clinical language models](#). In *Proceedings of the 22nd Annual Workshop of the Australasian Language Technology Association*, pages 164–172.
- Leibo Liu, Oscar Perez-Concha, Anthony Nguyen, Vicki Bennett, and Louisa Jorm. 2022b. [Hierarchical label-wise attention transformer model for explainable ICD coding](#). *Journal of biomedical informatics*, 133:104161.
- Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024b. [Lost in the middle: How language models use long contexts](#). *Transactions of the Association for Computational Linguistics*, 12:157–173.
- Mohammad Loni, Fatemeh Poursalim, Mehdi Asadi, and Arash Gharehbaghi. 2025. [A review on generative AI models for synthetic medical text, time series, and longitudinal data](#). *npj digital medicine*, 8(1):281.
- Pratyush Maini, Vineeth Dorna, Parth Doshi, Aldo Carranza, Fan Pan, Jack Urbanek, Paul Burstein, Alex Fang, Alvin Deng, Amro Abbas, Brett Larsen, Cody Blakeney, Charvi Bannur, Christina Baek, Darren Teh, David Schwab, Haakon Mongstad, Haoli Yin, Josh Wills, and 10 others. 2025. [BeyondWeb: Lessons from scaling synthetic data for trillion-scale pretraining](#). *arXiv [cs.LG]*.
- Pratyush Maini, Skyler Seto, Richard Bai, David Grangier, Yizhe Zhang, and Navdeep Jaitly. 2024. [Rephrasing the web: A recipe for compute and data-efficient language modeling](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14044–14072.
- Sewon Min, Kalpesh Krishna, Xinxu Lyu, Mike Lewis, Wen-Tau Yih, Pang Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. [FActScore](#):

- Fine-grained atomic evaluation of factual precision in long form text generation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12076–12100, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Avijit Mitra, Zhichao Yang, Emily Druhl, Raelene Goodwin, and Hong Yu. 2025. Synth-SBDH: A synthetic dataset of social and behavioral determinants of health for clinical text. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 27887–27923, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Monica Munnangi, Akshay Swaminathan, Jason Alan Fries, Jenelle A Jindal, Sanjana Narayanan, Ivan Lopez, Lucia Tu, Philip Chung, Jesutofunmi Omiye, Mehr Kashyap, and Nigam Shah. 2025. FactEHR: A dataset for evaluating factuality in clinical notes using LLMs. In *Machine Learning for Healthcare Conference*. PMLR.
- Aaron Nicolson, Jinghui Liu, Jason Dowling, Anthony Nguyen, and Bevan Koopman. 2024. e-health CSIRO at RRG24: Entropy-augmented self-critical sequence training for radiology report generation. In *Proceedings of the 23rd Workshop on Biomedical Natural Language Processing*, pages 99–104.
- Guilherme Penedo, Hynek Kydlíček, Loubna Ben Allal, Anton Lozhkov, Margaret Mitchell, Colin Raffel, Leandro Von Werra, and Thomas Wolf. 2024. The FineWeb datasets: Decanting the web for the finest text data at scale. *arXiv [cs.CL]*.
- Qwen, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, and 24 others. 2024. Qwen2.5 technical report. *arXiv [cs.CL]*.
- Kirk Roberts. 2024. Large language models for reducing clinicians’ documentation burden. *Nature medicine*.
- Ananya B Sai, Akash Kumar Mohankumar, and Mitesh M Khapra. 2022. A survey of evaluation metrics used for NLG systems. *ACM Comput. Surv.*, 55(2):1–39.
- Arne Schwieger, Katrin Angst, Mateo de Bardeci, Achim Burre, Flurin Cathomas, Stefano Ferrea, Franziska Grätz, Marius Knorr, Golo Kronenberg, Tobias Spiller, David Troi, Erich Seifritz, Samantha Weber, and Sebastian Olbrich. 2024. Large language models can support generation of standardized discharge summaries - a retrospective study utilizing ChatGPT-4 and electronic health records. *International journal of medical informatics*, 192(105654):105654.
- Yijia Shao, Yucheng Jiang, Theodore Kanell, Peter Xu, Omar Khattab, and Monica Lam. 2024. Assisting in writing wikipedia-like articles from scratch with large language models. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 6252–6278, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Iliia Shumailov, Zakhar Shumaylov, Yiren Zhao, Nicolas Papernot, Ross Anderson, and Yarin Gal. 2024. AI models collapse when trained on recursively generated data. *Nature*, 631(8022):755–759.
- Yuqi Si and Kirk Roberts. 2019. Deep patient representation of clinical notes via Multi-Task learning for mortality prediction. *AMIA Joint Summits on Translational Science proceedings. AMIA Joint Summits on Translational Science*, 2019:779–788.
- William R Small, Jonathan Austrian, Luke O’Donnell, Jesse Burk-Rafel, Katherine A Hochman, Adam Goodman, Jonah Zaretsky, Jacob Martin, Stephen Johnson, Vincent J Major, Simon Jones, Christian Henke, Benjamin Verplanke, Jwan Osso, Ian Larson, Archana Saxena, Aron Mednick, Choumika Simonis, Joseph Han, and 12 others. 2025. Evaluating hospital course summarization by an electronic health record-based large language model. *JAMA network open*, 8(8):e2526339.
- Luca Soldaini. 2016. QuickUMLS: a fast, unsupervised approach for medical concept extraction. *MedIR workshop, sigir*.
- M Solnyshkina, R Zamaletdinov, L Gorodetskaya, and A Gabitov. 2017. Evaluating text complexity and flesch-kincaid grade level. *Journal of social studies education research*, 8(3):238–248.
- Sarvesh Soni and Dina Demner-Fushman. 2024. Toward relieving clinician burden by automatically generating progress notes using interim hospital data. *AMIA Annual Symposium Proceedings*, 2024:1059–1068.
- Sarvesh Soni and Kirk Roberts. 2020. Paraphrasing to improve the performance of electronic health records question answering. *AMIA Summits on Translational Science proceedings AMIA Summit on Translational Science*, 2020:626–635.
- Ruixiang Tang, Xiaotian Han, Xiaoqian Jiang, and Xia Hu. 2023. Does synthetic data generation of LLMs help clinical text mining? *arXiv [cs.CL]*.
- Dave Van Veen, Cara Van Uden, Louis Blanke-meier, Jean-Benoit Delbrouck, Asad Aali, Christian Bluethgen, Anuj Pareek, Malgorzata Polacin, Eduardo Pontes Reis, Anna Seehofnerová, Nidhi Rohatgi, Poonam Hosamani, William Collins, Neera Ahuja, Curtis P Langlotz, Jason Hom, Sergios Gavidis, John Pauly, and Akshay S Chaudhari. 2024. Adapted large language models can outperform medical experts in clinical text summarization. *Nature medicine*, 30(4):1134–1142.
- Hanyin Wang, Chufan Gao, Christopher Dantona, Bryan Hull, and Jimeng Sun. 2024a. DRG-LLaMA : tuning

LLaMA model to predict diagnosis-related group for hospitalized patients. *NPJ digital medicine*, 7(1):16.

Junda Wang, Zonghai Yao, Zhichao Yang, Huixue Zhou, Rumeng Li, Xun Wang, Yucheng Xu, and Hong Yu. 2024b. [NoteChat: A dataset of synthetic patient-physician conversations conditioned on clinical notes](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 15183–15201, Stroudsburg, PA, USA. Association for Computational Linguistics.

Yuxia Wang, Minghan Wang, Muhammad Arslan Manzoor, Fei Liu, Georgi Nenkov Georgiev, Rocktim Jyoti Das, and Preslav Nakov. 2024c. [Factuality of large language models: A survey](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 19519–19529, Stroudsburg, PA, USA. Association for Computational Linguistics.

Yiqing Xie, Sheng Zhang, Hao Cheng, Pengfei Liu, Zelalem Gero, Cliff Wong, Tristan Naumann, Hoifung Poon, and Carolyn Rose. 2024. [DocLens: Multi-aspect fine-grained medical text evaluation](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 649–679, Stroudsburg, PA, USA. Association for Computational Linguistics.

Justin Xu, Zhihong Chen, Andrew Johnston, Louis Blankemeier, Maya Varma, Jason Hom, William J Collins, Ankit Modi, Robert Lloyd, Benjamin Hopkins, Curtis Langlotz, and Jean-Benoit Delbrouck. 2024. [Overview of the first shared task on clinical text generation: RRG24 and “discharge me!”](#). In *Proceedings of the 23rd Workshop on Biomedical Natural Language Processing*, pages 85–98.

Wen-Wai Yim, Yujuan Fu, Asma Ben Abacha, Neal Snider, Thomas Lin, and Meliha Yetisgen. 2023. [Acibench: a novel ambient clinical intelligence dataset for benchmarking automatic visit note generation](#). *Scientific data*, 10(1):586.

Feiyang Yu, Mark Endo, Rayan Krishnan, Ian Pan, Andy Tsai, Eduardo Pontes Reis, Eduardo Kaiser Ururahy Fonseca, Henrique Min Ho Lee, Zahra Shakeri Hossein Abad, Andrew Y Ng, Curtis P Langlotz, Vasantha Kumar Venugopal, and Pranav Rajpurkar. 2023. [Evaluating progress in automatic chest X-ray radiology report generation](#). *Patterns*, 0(0).

Tianyi Zhang*, Varsha Kishore*, Felix Wu*, Kilian Q Weinberger, and Yoav Artzi. 2020. [BERTScore: Evaluating text generation with BERT](#). In *International Conference on Learning Representations*.

Hendrik Švalov, Mihkel Lepson, Veronika Kukk, Maria Malk, Neeme Ilves, Hele-Andra Kuulmets, and Raivo Kolde. 2025. [Using synthetic health care data to leverage large language models for named entity recognition: Development and validation study](#). *Journal of medical internet research*, 27(1):e66279.

A Details of Rephrasing MIMIC Notes

For clinical text synthesis, each LLM was provided with a system prompt consistent with its chat template and was instructed to synthesize a new version using a paraphrasing prompt found effective in prior studies (Liu and Nguyen, 2024; Maini et al., 2024), denoted as the “high quality” prompt:

Rephrasing prompt

[System Prompt] You are a medical artificial intelligence assistant. The assistant gives truthful, detailed, and professional answers to the requests.

For the following paragraph give me a diverse paraphrase of the same in high-quality English language as in clinical notes written by medical professionals:

We consider two rephrasing setups, one *by note* and the other *by chunk*. For rephrasing *by note*, we stratify the clinical notes by lengths — (0, 500], (500, 1000], (1000, 2000], (2000, 4000], and > 4000 — and set the corresponding maximum number of generated tokens to 1000, 2000, 4000, 8000, and 10,000, respectively. For rephrasing *by chunk*, we set the maximum number of generated tokens to 512 given the short chunk length (~150 words). Temperature was set to 0.75 and top-p 0.9 for all LLMs.

B Implementation Details on Extrinsic Evaluation

We followed previous studies (Liu et al., 2022a) to construct the in-hospital prediction cohorts (mortality and DRG) from MIMIC-III. For ICD coding, we used the benchmark (Edin et al., 2023) that comprehensively examined various types of ICD codes from MIMIC and focused on ICD-10 based on MIMIC-IV. Only ICD codes with a frequency higher than 10 were retained in the benchmark. This threshold was also applied to the DRG cohort with MS-DRG (Medicare Severity-DRG). We constructed a readmission cohort from MIMIC-IV by considering only the first admission of each patient and discarding immediate readmission cases that occurred within 24 hours after discharge. The statistics of the datasets are presented in Table 5.

We chose PLM-ICD (Huang et al., 2022; Edin et al., 2023) as the classification model given its strong performance on ICD coding and adaptability to handle long texts with sliding window. More

	# Total words (billion)	# Avg words per note	# Avg words per chunk
Llama (by notes)	1.21	260	-
Mistral (by notes)	1.34	289	-
Qwen (by notes)	1.13	263	-
Llama (by chunks)	2.24	482	206
Mistral (by chunks)	1.74	376	160
Qwen (by chunks)	1.41	315	146
Total	9.1	332	-

Table 4: Statistics of MIMIC-EXT-SYNNOTES generated by LLMs. When rephrasing based on chunks, each note from MIMIC is segmented into 2.3 chunks on average. The synthetic notes contain 9.1B words in total.

specifically: 1) it leverages a pre-trained encoder (BERT), allowing for domain-specific initialization; 2) it processes text via a sliding-window mechanism, making it efficient for long clinical documents; 3) it has demonstrated strong performance, particularly for ICD coding, compared to other deep learning methods (Edin et al., 2023); and 4) it is adaptable across tasks by modifying only the classification layer.

We initialized the model with a biomedical RoBERTa checkpoint (Lewis et al., 2020). Training was conducted using a batch size of 16 for up to 10 epochs, with early stopping triggered after no improvement for 3 epochs. The learning rate was set to $5e-5$, with 2,000 warmup steps followed by linear decay. This training recipe was applied to all versions of input notes, either human-written or LLM-rephrased, to ensure consistent comparisons. The trained model was then evaluated on the test set of the corresponding task. We averaged the results over five runs with different random seeds to report the main results.

C Fact-checking Implementation

Following previous works (Munnangi et al., 2025; Xie et al., 2024), we performed fact-checking of the synthetic notes by decomposing the synthetic notes into atomic facts and measuring entailment. We use Gemini-2.5-Pro, a thinking model, to decompose a note into atomic facts, and Gemini-2.5-Flash to assess entailment. We used the prompts from Munnangi et al. (2025) for fact decomposition and entailment judgment, which are presented in the following.

We sampled 25 notes of varied length (10 radiology reports and 15 discharge summaries) for fact-checking due to the cost related to querying Gemini. We also consider *long notes* over 2K words, constituting five discharge summaries from the 25 notes. We show results (Table 2) on both these 25

notes and the five long notes. For synthetic notes rephrased *by chunk*, we also only present the results on the five long notes. Due to the prolonged content, these five notes take significantly longer to decompose and judge, sometimes requiring 500+ API calls for a single pair. Given the cost consideration, we relied on this small subset for this section. We also only focused on Llama and Mistral for this analysis, which contributed 5.2K and 6.1K facts, respectively.

D Omitted and Added Semantic Types

Table 6 presents the top-10 semantic types of the medical concepts that are either omitted or added by the LLMs after rephrasing. These are based on MedCON (Yim et al., 2023) but mapped CUI to TUI. The added semantic types can be considered as false positives (FPs) and omitted ones as false negatives (FNs) when comparing against the human-written notes.

E Alternative Prompts

We consider the following alternative prompts to explore their impact on generated text.

Rearranged Structure *Rewrite the following paragraph by rearranging sentences and phrases, while retaining the original clinical content and meaning:*

Minimal Changes *Rephrase the following paragraph with minimal changes, keeping the original structure and meaning intact. Make only slight adjustments for clarity and conciseness while preserving the clinical details:*

Formal Clinical Language *Paraphrase the following paragraph using strictly formal clinical documentation language, maintaining the same level of detail and medical accuracy:*

	# Patients	% Train/Val/Test	# Classes	Input notes	Metrics
In-hospital Mortality	33,515 (MIMIC-III)	80%/10%/10%	2 (11.1%)	Nursing notes, physician notes, etc.	AUROC AUPRC
30-Day Readmission	33,518 (MIMIC-IV)	80%/10%/10%	2 (13.0%)	Discharge summary	AUROC AUPRC
DRG Prediction	18,681 (MIMIC-III)	70%/20%/10%	369	Nursing notes, physician notes, etc.	Accuracy F1
ICD Coding	122,278 (MIMIC-IV)	73%/16%/11%	7,942	Discharge summary	Micro F1 Macro F1

Table 5: Description of cohorts and prediction tasks. For diagnosis-related classifications, DRG prediction is based on MS-DRG and ICD coding on ICD-10. All tasks have a single label per patient except for ICD coding, which is multi-label multi-class classification.

Substantial Revision *Produce a substantially revised version of the following paragraph in professional clinical English, changing sentence structure and vocabulary while preserving all critical medical details:*

Simpler English *Paraphrase the following paragraph into simpler clinical English, minimizing jargon and long sentences but keeping the key medical details:*

Entailment judgment prompt

You are an expert on natural language entailment.
Your task is to deduce whether premise statements entail hypotheses.
Return only '1' if the hypothesis can be fully entailed by the premise.
Return only '0' if the hypothesis contains information that cannot be entailed by the premise.
Generate the answer in JSON format with the following keys:
'entailment_prediction': 1 or 0, whether the claim can be entailed.
Only return the JSON-formatted answer and nothing else.

Premise: `{{premise}}`
Hypothesis: `{{hypothesis}}`

Here is the JSON-formatted answer:

Fact decomposition prompt

Please breakdown the following text into independent facts as a string delimited by "/" to separate the facts

Example 1:

Note: "There is a dense consolidation in the left lower lobe."

Atomic facts:

There is a consolidation. // The consolidation is dense. // The consolidation is on the left. // The consolidation is in a lobe. // The consolidation is in the lower portion of the left lobe.

Example 2:

Note: "The patient has been having intermittent shortness of breath for the last two years."

Atomic facts:

The patient has been having shortness of breath. // The shortness of breath is intermittent. // The shortness of breath has been present for the last two years.

Do not include any other text, or say "Here is the list..."

Note: `{{note}}`

F Alternative Models

Table 8 presents the numerical results presented in Figure 2. We provide additional results on rephrasing with Gemma-2 (9B) and Mixtral-v0.1 (8x7B) by chunks with one seed.

Synthetic text rephrased by Llama	
<i>Added Semantic Types (FP)</i>	<i>Omitted Semantic Types (FN)</i>
<i>Total counts: 9,321</i>	<i>Total counts: 8,979</i>
Natural Phenomenon or Process (19.1%)	Medical Device (13.0%)
Clinical Attribute (15.2%)	Body Space or Junction (12.4%)
Organ or Tissue Function (10.9%)	Mental Process (11.7%)
Pharmacologic Substance (9.9%)	Biomedical or Dental Material (11.3%)
Phenomenon or Process (8.1%)	Gene or Genome (9.9%)
Organism Attribute (7.9%)	Body System (9.1%)
Body Substance (7.6%)	Indicator, Reagent, or Diagnostic Aid (8.8%)
Organism Function (7.4%)	Cell Component (8.5%)
Laboratory or Test Result (7.3%)	Organ or Tissue Function (7.9%)
Body Part, Organ, or Organ Component (6.6%)	Physiologic Function (7.3%)

Synthetic text rephrased by Mistral	
<i>Added Semantic Types (FP)</i>	<i>Omitted Semantic Types (FN)</i>
<i>Total counts: 7,409</i>	<i>Total counts: 4,238</i>
Natural Phenomenon or Process (15.9%)	Gene or Genome (14.6%)
Clinical Attribute (11.9%)	Body Space or Junction (13.0%)
Phenomenon or Process (11.4%)	Mental Process (11.3%)
Organ or Tissue Function (11.2%)	Medical Device (10.3%)
Organism Attribute (9.9%)	Cell Component (9.4%)
Organism Function (8.5%)	Biomedical or Dental Material (8.9%)
Body Space or Junction (8.3%)	Biologic Function (8.8%)
Laboratory or Test Result (8.2%)	Body System (8.0%)
Pharmacologic Substance (7.8%)	Indicator, Reagent, or Diagnostic Aid (7.9%)
Body Substance (6.9%)	Organ or Tissue Function (7.8%)

Synthetic text rephrased by Qwen	
<i>Added Semantic Types (FP)</i>	<i>Omitted Semantic Types (FN)</i>
<i>Total counts: 5,641</i>	<i>Total counts: 5,288</i>
Natural Phenomenon or Process (13.1%)	Mental Process (12.4%)
Clinical Attribute (11.9%)	Body Space or Junction (11.4%)
Pharmacologic Substance (11.0%)	Clinical Attribute (11.2%)
Phenomenon or Process (10.4%)	Gene or Genome (10.7%)
Cell (9.8%)	Medical Device (10.3%)
Element, Ion, or Isotope (9.5%)	Body Location or Region (9.3%)
Laboratory or Test Result (8.8%)	Nucleic Acid, Nucleoside, or Nucleotide (9.3%)
Molecular Function (8.6%)	Cell Component (8.6%)
Organism Function (8.6%)	Indicator, Reagent, or Diagnostic Aid (8.4%)
Biomedical or Dental Material (8.3%)	Organism Attribute (8.4%)

Table 6: Changes of semantic types in the rephrased notes by the three LLMs. Top 10 semantic types are listed for false positive (FP) and false negative (FN) compared to the original, human-written notes. The ratio is calculated based on total counts of the presented semantic types for each LLM for FP and FN, respectively.

Claim	Error Category	Comment
“The patient presented with fever.”	No error	<i>This is a correct claim.</i>
“The patient continued to have significant oxygen requirements.”	No error	<i>Correct.</i>
“An MRI later confirmed the focus was a subacute infarction.”	Misinterpretation	<i>This is likely, but not confirmed.</i>
“The anemia was treated with iron supplementation.”	Fabricated claim	<i>Patient had iron-binding study. No supplementation intake was documented.</i>
“The surgery occurred on discharge day.”	Misinterpretation	<i>Wrong interpretation of “Surgery Discharge part”.</i>
“The patient was cleared by occupational therapy (OT) for discharge to rehabilitation.”	Misinterpretation	<i>Cleared by PT.</i>
“The patient was scheduled for a follow-up appointment with her gastroenterologist.”	Misinterpretation	<i>Appointment scheduled with primary care physician.</i>
“The patient’s name is Mr. John Doe.”	Fabricated claim	<i>Made-up name.</i>
“The patient is a non-smoker.”	Misinterpretation	<i>Incorrect; tobacco use is mentioned.</i>
“The patient consumes approximately 1–2 drinks per day.”	Measurement/Numeracy	<i>EtOH: [**4–20**] drinks/day.</i>
“The adjusted medications include simvastatin.”	Fabricated claim	<i>No adjustment.</i>
“The patient is asymptomatic with the low blood sugar levels.”	Misinterpretation	<i>“He states that he is asymptomatic,” but note shows symptomatic hypoglycemia.</i>
“In the emergency department, the patient’s vital signs were stable.”	Misinterpretation	<i>Incorrect; labs showed elevated WBC.</i>
“The patient’s most recent pulse was 95 bpm.”	Temporal/Recency	<i>Not most recent.</i>
“The patient’s most recent urine output was 2 L.”	Temporal/Recency	<i>Not most recent.</i>
“The patient is a former smoker.”	Temporal/Recency	<i>Current smoker.</i>
“All pathology samples of bone and cartilage showed hyaline cartilage with focal acute inflammation.”	Misinterpretation	<i>Incorrect summarization of pathology results: not all samples show hyaline cartilage with focal acute inflammation.</i>
“In the ER, the patient’s respiration rate was 142/min.”	Measurement/Numeracy	<i>142 is blood pressure.</i>
“In the ER, the patient’s blood pressure was 18/97 mmHg.”	Measurement/Numeracy	<i>18 is respiratory rate; 96 (instead of 97) is oxygen saturation.</i>
“The patient’s discharge blood pressure was 168/45 mmHg.”	Misinterpretation	<i>BP in ED, not discharge.</i>

Table 7: Categorization and analysis of example false positive claims in synthetic texts.

Input source	In-hospital Mortality		30-day Readmission		ICD Coding		DRG Classification	
	AUROC	AUPRC	AUROC	AUPRC	Micro F1	Macro F1	Accuracy	F1
Human-written	0.873 (0.858–0.888)	0.507 (0.467–0.547)	0.627 (0.566–0.688)	0.233 (0.168–0.297)	0.564 (0.562–0.566)	0.180 (0.172–0.189)	0.314 (0.308–0.321)	0.132 (0.115–0.149)
<i>Synthetic text by notes</i>								
Llama	0.874 (0.853–0.895)	0.521 (0.480–0.561)	0.602 (0.514–0.691)	0.219 (0.149–0.289)	0.458 (0.456–0.460)	0.138 (0.131–0.144)	0.306 (0.301–0.311)	0.125 (0.106–0.144)
Mistral	0.879 (0.871–0.888)	0.528 (0.502–0.555)	0.637 (0.609–0.664)	0.241 (0.212–0.269)	0.522 (0.521–0.524)	0.158 (0.154–0.161)	0.315 (0.311–0.318)	0.138 (0.125–0.150)
Qwen	0.861 (0.849–0.872)	0.477 (0.447–0.508)	0.564 (0.500–0.628)	0.182 (0.119–0.244)	0.452 (0.447–0.458)	0.122 (0.110–0.135)	0.299 (0.295–0.303)	0.124 (0.108–0.139)
<i>Synthetic text by chunks</i>								
Llama	0.874 (0.857–0.892)	0.521 (0.464–0.579)	0.654 (0.626–0.683)	0.265 (0.237–0.293)	0.539 (0.538–0.541)	0.172 (0.161–0.183)	0.305 (0.298–0.312)	0.127 (0.119–0.136)
Mistral	0.881 (0.872–0.890)	0.530 (0.504–0.555)	0.643 (0.596–0.690)	0.244 (0.187–0.301)	0.549 (0.547–0.551)	0.177 (0.167–0.187)	0.309 (0.305–0.314)	0.133 (0.118–0.149)
Qwen	0.868 (0.862–0.874)	0.492 (0.458–0.526)	0.618 (0.552–0.683)	0.229 (0.161–0.296)	0.519 (0.517–0.521)	0.162 (0.153–0.170)	0.293 (0.290–0.296)	0.121 (0.101–0.142)
<i>Additional synthesis by chunks</i>								
Gemma	0.870	0.512	0.654	0.285	0.537	0.171	0.293	0.109
Mixtral	0.882	0.548	0.662	0.275	0.548	0.169	0.314	0.123

Table 8: Downstream modeling results for the four tasks based on different sources of clinical notes. Values are shown as mean with 95% confidence intervals in parentheses. The results for the three main LLMs are also presented in Figure 2.