

Privacy-Preserving Generation of Clinical Narratives from Medical Terminologies

Yuping Wu¹, Viktor Schlegel^{1,2,3†}, Warren Del-Pinto¹, Srinivasan Nandakumar²,
Iqra Zahid², Yidan Sun², Hao Li³, Usama Farghaly Omar⁴, Amirah Jasmine³,
Arun-Kumar Kaliya-Perumal⁵, Chun Shen Tham¹, Gabriel Connors¹,
Anil Anthony Bharath^{2,3}, Goran Nenadic¹

¹University of Manchester, Department of Computer Science ²Imperial Global Singapore

³Imperial College London, Department of Bioengineering

⁴Khoo Teck Puat Hospital, Orthopaedic Surgery Department, Singapore

⁵Nanyang Technological University, Rehabilitation Research Institute of Singapore

Abstract

In high-stakes domains such as healthcare, privacy concerns severely limit the use of real-world training data. Differentially private (DP) synthetic data offers a promising alternative with formal privacy guarantees, but achieving strong utility remains challenging for clinical note generation due to domain specificity and long-form text complexity. We present **Term2Note***, a method for synthesising full-length clinical notes under DP constraints. By structurally separating content and form, Term2Note generates section-wise note content conditioned on medical terms, with terms and notes privatised under separate DP constraints, and applies a DP quality maximiser to improve outputs. Experiments demonstrate that Term2Note produces synthetic notes with statistical properties closely aligned with real clinical notes, and that downstream models trained on these notes achieve performance comparable to those trained on real clinical data. Compared to existing DP text generation baselines, Term2Note substantially improves both fidelity and utility, without relying on label distribution assumptions, highlighting its effectiveness as a practical privacy-preserving alternative to real clinical notes.

1 Introduction

The scaling law of neural language models (Kaplan et al., 2020) suggests that model performance improves with larger training corpora. As large language models (LLMs) continue to scale in size, with models such as Llama 3 (Meta, 2024), Gemma 3 (Google, 2025), and Qwen3 (Yang et al., 2025) ranging from 0.6B to over 405B parameters, the demand for large-scale, high-quality training data has risen accordingly. Synthetic data generation with LLMs has therefore emerged as a promising approach, with instruction-following synthetic

datasets (Schick and Schütze, 2021; Taori et al., 2023; Li et al., 2025b) proving effective for model pretraining and fine-tuning. This is especially relevant for high-stakes domains such as healthcare (Li et al., 2025a), where real data is sensitive, highly regulated, and difficult to share (Schlegel et al., 2025; Sun et al., 2025a). Although large amounts of clinical data exist within healthcare institutions, access to these datasets remains extremely limited. A practical and privacy-conscious solution is to share synthetic versions of sensitive clinical data instead of the raw data itself. However, because even synthetic texts can leak private information (Meeus et al., 2025; Sun et al., 2025b), to make such sharing viable, formal privacy guarantees are essential.

Differential privacy (DP) provides a principled framework for this purpose (Alzoubi and Mishra, 2025). By bounding the influence of any individual record on the synthesised dataset, DP offers quantifiable privacy guarantees. Prior work on DP-based text generation has focused mainly on short-form texts in low-risk domains such as reviews (Yue et al., 2023; Kurakin et al., 2023; Mattern et al., 2022; Flemings and Annaram, 2024). In biomedical settings, efforts have been restricted to synthesising (public) PubMed abstracts (Xie et al., 2024) and relatively short clinical passages (Al Aziz et al., 2021; Ramesh et al., 2024), with no prior work addressing the more complex task of generating full-length clinical notes under DP constraints.

Synthesising clinical notes with DP presents two key challenges. First, *generation complexity* (Kweon et al., 2024a; Weetman et al., 2021): clinical notes are long and exhibit diverse structures and free-form content, making it more difficult for generative models to maintain coherence and quality, particularly under privacy constraints. Second, *domain specificity* (Adnan et al., 2010): clinical notes typically contain extensive domain-specific terminologies, which require expert knowledge to

[†]Correspondence: v.schlegel@imperial.ac.uk

*Code: <https://github.com/yuping-wu/Term2Note>

understand and reproduce accurately.

In this paper, we tackle this underexplored and challenging task by proposing Term2Note, a novel methodology for DP synthetic clinical note generation. As illustrated in Figure 1, Term2Note addresses the challenge of long-form generation by leveraging domain-specific document structures to decompose the task into smaller, section-wise subtasks. To handle domain specificity, generation is conditioned not on generic metadata (e.g., class labels or diagnosis codes such as *Enterocolitis due to Clostridium difficile*), but on salient clinical terms (e.g., [diarrhea, Clostridium difficile colitis, Vancomycin]), which are subjected to an additional layer of DP protection as rare combinations of such terms might bear patient-identifying information. This term-based conditioning strategy allows the model to generate text that is both clinically meaningful and structurally coherent, while also enabling more fine-grained privacy control.

Our experimental results show that Term2Note: (1) produces synthetic notes with high structural and semantic fidelity to real clinical data; (2) provides high utility for downstream tasks, such as ICD code prediction; (3) satisfies formal DP guarantees, facilitating data sharing with quantifiable privacy protection. Moreover, Term2Note consistently demonstrates stronger performance than baselines across a wide range of evaluation metrics. In summary, Term2Note offers a promising solution for privacy-preserved clinical notes sharing.

2 Background & Related Work

De-identification is widely used to protect sensitive data by removing explicit personal identifiers. However, it does not provide formal privacy guarantees, and individuals can often be re-identified through quasi-identifiers or linkage with external datasets. For example, users in the anonymised Netflix dataset were successfully re-identified by linking it with IMDb ratings (Narayanan and Shmatikov, 2006). In contrast, differential privacy (DP) provides formal, quantifiable privacy guarantees against such inference attacks. We briefly review the theoretical foundations of DP below.

Definition 1. (Dwork et al., 2006) A randomised algorithm $\mathcal{M} : \mathbb{N}^{|\mathcal{X}|} \rightarrow R$ is said to be (ϵ, δ) -*differentially private (DP)*, if, for any two neighboring datasets D and D' differing in one single instance, and for all subsets S of the output space of $\mathcal{M}$, it has

$$\mathbb{P}[\mathcal{M}(D) \in S] \leq e^\epsilon \mathbb{P}[\mathcal{M}(D') \in S] + \delta.$$

The definition implies the probability distributions induced by $\mathcal{M}$ on neighboring datasets must be close, with their likelihood ratios bounded by a multiplicative factor of e^ϵ and an additive slack of δ . Smaller ϵ values indicate stronger privacy, while δ denotes the (typically negligible) probability of a privacy breach exceeding the e^ϵ bound.

Theorem 1. (Post-Processing) (Dwork and Roth, 2014) Let $\mathcal{M} : \mathbb{N}^{|\mathcal{X}|} \rightarrow R$ be a randomised algorithm that is (ϵ, δ) -DP. Let $f : R \rightarrow R'$ be an arbitrary randomised function. Then $f \circ \mathcal{M} : \mathbb{N}^{|\mathcal{X}|} \rightarrow R'$ is also (ϵ, δ) -DP.

Theorem 2. (Parallel Composition) (McSherry, 2009) Let dataset $D = D_1 \cup \dots \cup D_k$, and $D_i \cap D_j = \emptyset$ for $i \neq j$. Let $\mathcal{M}_i : \mathbb{N}^{|\mathcal{X}_i|} \rightarrow R_i$ be a randomised algorithm that is (ϵ_i, δ_i) -DP, for $i \in [k]$. Then $\mathcal{M}(D) = (\mathcal{M}_1(D_1), \dots, \mathcal{M}_k(D_k))$ is $(\max_i \epsilon_i, \max_i \delta_i)$ -DP.

The post-processing property states that applying any function to the output of an (ϵ, δ) -DP algorithm preserves the same privacy guarantee. Thus, a DP-trained generative model remains private during downstream use. The parallel composition property states that applying DP mechanisms to disjoint data subsets results in an overall privacy loss bounded by the maximum individual (ϵ, δ) values, enabling clinical terms and notes to be privatised independently in our method.

DP in Deep Learning is commonly achieved by adding random noise, typically Gaussian noise, during training. DP-SGD (Abadi et al., 2016) enforces DP by clipping per-sample gradients and injecting noise at each optimisation step, limiting the influence of individual training examples on model parameters. Subsequent work has improved the efficiency of DP-SGD (Lee and Kifer, 2020; Bu et al., 2023; Yousefpour et al., 2021) by reducing the time and memory overhead associated with per-example gradient computation. Unless otherwise specified, we adopt FastDP (Bu et al., 2023) for DP fine-tuning throughout this paper.

Synthetic text generation has progressed rapidly with the rise of LLMs, especially through instruction-following datasets that enhance downstream performance (Taori et al., 2023; Peng et al., 2023; Schick and Schütze, 2021). Generating synthetic clinical text, however, is substantially more challenging due to the domain complexity, expert knowledge requirements, and privacy concerns. Prior attempts have often produced datasets




Figure 1: Overview of **Term2Note**. An original clinical note is split into sections and associated medical terms, which are embedded and optionally privatised via DP perturbation. A term generation model produces synthetic terms from these embeddings, and a DP-trained note generation model synthesises section-wise notes conditioned on them. A DP quality maximiser then selects the final synthetic note. Appendix A illustrates the whole process with a concrete example.

with limited downstream utility (Li et al., 2023a; Schlegel et al., 2023). Recent work has sought to improve realism by prompting LLMs with control code (e.g., ICD codes) (Falisi et al., 2024) or transforming biomedical abstracts (e.g., PubMed content) into clinical-style text (Kweon et al., 2024b). While promising, these methods do not address privacy, a central concern in clinical settings. A straightforward strategy is de-identifying private information and prompting LLMs to fill in the gaps (Sarkar et al., 2025), but a more principled approach is to fine-tune or instruction-tune generative models with DP for formal protection (Al Aziz et al., 2021; Ramesh et al., 2024; Baumel et al., 2024). However, prior studies operate on relatively short clinical texts, either naturally brief or deliberately truncated, making generation considerably easier. In contrast, DP-constrained generation of full-length clinical notes, as addressed in this paper, is more challenging due to their length, unstructured format, and high variability.

3 Methodology

3.1 Problem Statement

Given a private dataset D^{src} consisting of clinical notes, our goal is to develop a mechanism $\mathcal{M}$ that satisfies (ϵ, δ) -DP and produces a synthetic dataset D^{syn} as its output. Let $X^{\text{src}} \in D^{\text{src}}$ denote an original clinical note and $X^{\text{syn}} \in D^{\text{syn}}$ its corresponding synthetic version. To support the development and

evaluation of the mechanism, we partition the private dataset into a training set $D_{\text{train}}^{\text{src}}$ and a test set $D_{\text{test}}^{\text{src}}$. The training set is used to develop $\mathcal{M}$, while the test set remains completely unseen by $\mathcal{M}$ to provide an unbiased evaluation. Additionally, we assume access to a public dataset of clinical terms, denoted as D_{public} , which can be automatically derived from publicly available medical resources and therefore used freely without privacy constraints.

We introduce **Term2Note**, a section-wise DP generation framework for clinical notes. An overview is illustrated in Figure 1, with the detailed procedure in Algorithm 1. In the following, we first elaborate on Term2Note alongside the algorithm, and then present the implementation details.

3.2 Format and Term Identification

Since our framework synthesises clinical notes via section-wise generation conditioned on medical terms, we first standardise the structure of the notes and identify salient clinical terms. Let SECSPLIT denote an automatic section segmentation module. Given an original clinical note X^{src} , SECSPLIT outputs a list of m segmented sections: $[\text{SEC}_1^{\text{src}}, \dots, \text{SEC}_m^{\text{src}}] = \text{SECSPLIT}(X^{\text{src}})$, where $\text{SEC}_i^{\text{src}} = ""$ if the i -th section is absent. Next, we apply an automatic term extraction module, denoted as TERMEXT, to each section to identify clinically salient terms, resulting in a list of section-specific medical terms: $[T_1^{\text{src}}, \dots, T_m^{\text{src}}] =$

$[\text{TERMEXT}(\text{SEC}_1^{\text{src}}), \dots, \text{TERMEXT}(\text{SEC}_m^{\text{src}})]$.

3.3 Clinical Terms Generation

Clinical terms from private notes may still contain sensitive information. For example, unique combinations of diagnoses and procedures could re-identify patients. A uniqueness analysis on a large subset of the MIMIC-IV dataset reveals that (1) although clinical terms are often shared across notes, more than half of the notes contain at least one clinical term that is unique to that note, and (2) certain sections, such as clinical course and medication, are highly distinctive, with over 96% of notes containing unique combination of terms in these sections. The presence of note-specific clinical terms increases the risk of patient re-identification and privacy exposure. To mitigate this risk, we introduce an optional DP step for term generation. We formulate it as a reconstruction task, fine-tuning a generative model on the public term dataset D_{public} to recover term lists from embeddings, and then applying it to private data under DP constraints.

Specifically, any given (section-wise) term list is first embedded using a clinical text embedder, denoted EMB. A projection layer PROJ, parameterised by θ_p , maps the embeddings to the hidden dimensionality required by the generative model TERMGEM, parameterised by θ_t . The model then reconstructs the original term list from the projected embeddings. To ensure privacy, we adapt the DPRP schema (Gondara and Wang, 2020), a model-agnostic DP mechanism originally proposed for tabular data, to term embeddings derived from private notes, denoted as DPRP*. The procedure perturbs private embeddings in four steps: (1) add dimension-wise random noise to input embeddings; (2) compute the covariance matrix of the input embeddings and add random noise; (3) perform singular value decomposition (SVD) on the noisy covariance matrix; (4) reconstruct the inputs from the noisy embeddings and the right singular vectors. The pseudocode of DPRP* is provided in Appendix B. To minimise the distribution difference between embeddings used in training and those perturbed during inference, we additionally add Gaussian noise to the embeddings during training. Formally, the process is defined as follows:

$$E' = \begin{cases} \text{EMB}(D_{\text{public}}) + \mathcal{N}(0, \sigma_{\text{emb}}), & \text{if training,} \\ \text{DPRP}^*(\text{EMB}(T^{\text{src}}), \frac{\epsilon_t}{m}, \frac{\sigma_t}{m}), & \text{otherwise.} \end{cases} \quad (1)$$

$$T^{\text{syn}} \sim \text{TERMGEM}_{\theta_t}(\text{PROJ}_{\theta_p}(E')) \quad (2)$$

Here, ϵ_t and σ_t denote the privacy parameters of DPRP*, and m is the number of sections in a single note. Since E' is computed at the section level, the overall privacy cost for an entire note accumulates across sections. To account for this, we distribute the privacy budget evenly by scaling the cost for each section to $\frac{1}{m}$ of the total budget.

3.4 Clinical Note Generation

We define section-wise clinical note generation as a conditional text generation task. Given the task instruction I , the content of the previous (generated) sections $[\text{SEC}_1, \dots, \text{SEC}_{i-1}]$, and a list of clinical terms T_i for the current section, a generative model NOTEGEM, parameterised by θ_n , is trained to produce the i -th section of the note under (ϵ_n, δ_n) -DP. During inference, each section is sampled sequentially as:

$$\text{SEC}_i^{\text{syn}} \sim \text{NOTEGEM}_{\theta_n}(I, [\text{SEC}_1^{\text{syn}}, \dots, \text{SEC}_{i-1}^{\text{syn}}], T_i) \quad (3)$$

Here, T_i can be the original extracted terms T_i^{src} or the synthetic privatised terms T_i^{syn} , depending on the privacy configuration. With the section group name provided, the instruction I is defined as shown in Figure 1. Finally, a synthetic full note is obtained by concatenating the generated sections: $X^{\text{syn}} = [\text{SEC}_1^{\text{syn}}, \dots, \text{SEC}_m^{\text{syn}}]$.

3.5 DP Quality Maximiser

To improve the quality of the synthetic data, we introduce a quality maximisation strategy during inference by leveraging the generative capabilities of LLMs. Specifically, instead of generating a single synthetic note, we perform preference sampling on k candidate notes for X^{src} , denoted $X^{\text{syn}}[1 : k]$. Notably, this sampling procedure preserves the DP guarantee due to the post-processing property of DP. To select the most fluent and coherent output among the candidates, we use perplexity as the preference model. Perplexity reflects the likelihood of a sequence under an LLM, computed as the exponentiated average negative log-likelihood of the tokens. Lower perplexity indicates higher linguistic plausibility. To avoid bias from the generator itself, we compute perplexity scores using a reference domain-specialised LLM, denoted LLM_{ppl} . This ensures a more objective assessment of sequence quality. Formally, the perplexity of a candidate note $X^{\text{syn}}[i]$ is given by $\text{PPL}(X^{\text{syn}}[i]) = \exp(-\frac{1}{t} \sum_{i=1}^t \log \text{LLM}_{\text{ppl}}(d_i | d_{<i}))$, where $d_{1,\dots,t}$ are tokens in $X^{\text{syn}}[i]$. The synthetic note

with the lowest perplexity score is selected.

3.6 Privacy Analysis

The overall privacy guarantee of Term2Note depends on the composition of its two DP components: TERMGEN and NOTEGEN. Specifically, TERMGEN is trained on D_{public} and can optionally be applied to privatise terms for $D_{\text{test}}^{\text{src}}$, while NOTEGEN is trained on $D_{\text{train}}^{\text{src}}$ under DP. Since TERMGEN and NOTEGEN operate on disjoint subsets of private data, the overall privacy loss can be computed using parallel composition. Formally, the total privacy guarantee (ϵ, δ) is defined as below, and the proof is provided in Appendix C.

$$(\epsilon, \delta) = (\max(\epsilon_n, \epsilon_t), \max(\delta_n, \delta_t)) \quad (4)$$

3.7 Implementation Details

SECSPLIT Clinical notes lack a universally standardised structure, and while formats such as SOAP are common, they often require manual annotation for accurate segmentation (Gao et al., 2022). To enable automatic sectioning, we analyse formatting patterns in the raw notes and design a rule-based segmentation approach using regular expressions to identify section titles. Section spans are determined greedily, from each detected title to the next. We automatically curate a list of frequent section titles and group them into six semantic categories: *Patient Information*, *Clinical Course & History*, *Examinations & Findings*, *Laboratory & Imaging Results*, *Hospital Stay & Treatment*, and *Medications & Discharge Plan*. This taxonomy underpins SECSPLIT, which splits each note into up to six standardised sections. The full list of section titles and groupings is provided in Appendix D.

TERMEXT Biomedical vocabularies vary by taxonomy. In this work, we focus on terms from SNOMED CT, a comprehensive clinical vocabulary widely used in electronic health records (EHRs) and included in the Unified Medical Language System (UMLS) metathesaurus (U.S. NLM, 2025). To extract medical terms from clinical text, we use QuickUMLS (Soldaini, 2016), an unsupervised approximate string-matching tool over UMLS concepts, and retain only SNOMED CT terms.

Backbone Models For clinical term embeddings, we use the encoder-only model MedEmbed-large (Balachandran, 2024) as the embedder EMB, which is fine-tuned for medical and clinical text. For the two generative modules in Term2Note, we adopt

lightweight yet effective language models: GPT2-Large (Radford et al., 2019) for the term generator TERMGEN, and Llama-3.2-1B (Meta, 2024) or Gemma-3-1B (Google, 2025) for the note generator NOTEGEN. Using both Llama and Gemma increases architectural diversity and reduces model-specific bias while enabling reproducible evaluation. To compute the perplexity of generated notes, we use Asclepius-Llama3-8B (Kweon et al., 2024b) as the reference model LLM_{ppl}. Pre-trained on clinical text and supporting up to 8192 tokens as input, it mitigates domain mismatch and accommodates full-length clinical notes.

4 Experimental Setting

Datasets The MIMIC dataset series is one of the most widely used resources for clinical NLP. We use discharge notes from MIMIC-III (Johnson et al., 2016) and MIMIC-IV (Johnson et al., 2023) for different purposes. MIMIC-III serves as the *public* dataset to train the term generation model TERMGEN, while MIMIC-IV is treated as the *private* dataset for training the note generation model NOTEGEN under DP constraints. For both datasets, we filter out notes without associated ICD codes. We further exclude notes annotated in the SNOMED CT Entity Linking Challenge (Hardman et al., 2025) from MIMIC-IV notes and reserve them as the test set. This results in three datasets (statistics in Appendix E) for our experiments: D_{public} , with 52.7k MIMIC-III notes (500 held out for development); $D_{\text{train}}^{\text{src}}$, with 122k MIMIC-IV notes after exclusion; $D_{\text{test}}^{\text{src}}$, consisting of 204 SNOMED-annotated notes.

Hyperparameters Following common practice in previous work (Yu et al., 2021; De et al., 2022; Baumel et al., 2024), we experiment with different privacy budgets by varying the overall $\epsilon \in [2, 5, 8]$. Accordingly, the privacy budget for term generation (ϵ_t) or note generation (ϵ_n) is set to one of these values, and the corresponding δ value is set as $\frac{1}{N \log N}$ where N denotes the size of the private dataset. All experiments are conducted on up to two Nvidia A100 80GB GPUs. More training details, including learning rate, number of epochs, batch size, etc., are provided in Appendix F.

Baselines We compare our method with existing DP approaches for synthetic text generation: (1) DP-SGD with control codes (Yue et al., 2023), where a language model is fine-tuned under DP

Method	Fidelity				Utility					
	Length	Unary/Binary Term		Semantic	F1		AUC		Precision@k	
	KL Div.↓	Jaccard↑	KL Div.↓	MAUVE↑	Micro	Macro	Micro	Macro	k = 3	k = 5
Original Data					57.03	30.80	82.01	58.88	68.93	62.14
$\epsilon = \infty$										
AUG-PE ($\epsilon_{\text{meta}} = \infty, \epsilon_n = \infty$)	11.96	0.14/0.02	7.59/16.34	0.01	45.82	14.84	79.52	54.35	68.48	61.77
FastDP ($\epsilon_{\text{meta}} = \infty, \epsilon_n = \infty$)	1.04	0.53/0.28	0.32/0.84	0.12	<u>53.02</u>	<u>25.51</u>	79.35	51.05	<u>69.77</u>	60.39
Term2Note ($\epsilon_{\text{meta}} = \infty, \epsilon_n = \infty$)	0.25	0.52/0.20	0.22/1.08	0.59	49.95	21.89	81.40	55.43	<u>69.77</u>	<u>61.96</u>
Term2Note ($\epsilon_t = \infty, \epsilon_n = \infty$)	0.68	0.43/0.14	0.45/1.95	0.46	49.24	21.90	80.24	51.54	67.81	61.08
Term2Note ($\epsilon_t = \infty, \epsilon_n = 8$)	0.26	0.39/0.13	0.50/1.19	0.35	48.16	20.63	78.72	50.01	65.35	58.72
Term2Note ($\epsilon_t = \infty, \epsilon_n = 5$)	0.28	0.38/0.13	0.53/1.23	0.27	51.00	22.73	78.80	50.37	67.15	59.91
Term2Note ($\epsilon_t = \infty, \epsilon_n = 2$)	0.43	0.37/0.12	0.60/1.35	0.36	48.57	20.31	79.56	51.75	68.64	<u>60.41</u>
$\epsilon = 8$										
AUG-PE ($\epsilon_{\text{meta}} = \infty, \epsilon_n = 8$)	11.71	0.19/0.03	5.03/12.18	0.01	40.73	13.28	78.13	53.49	63.24	59.03
FastDP ($\epsilon_{\text{meta}} = \infty, \epsilon_n = 8$)	4.51	0.31/0.10	2.88/5.88	0.02	48.58	16.40	<u>80.74</u>	51.57	69.79	<u>61.59</u>
Term2Note ($\epsilon_{\text{meta}} = \infty, \epsilon_n = 8$)	0.39	<u>0.40/0.13</u>	<u>0.47/1.14</u>	<u>0.53</u>	49.71	21.28	80.03	52.80	67.49	61.48
Term2Note ($\epsilon_t = 8, \epsilon_n = 8$)	0.16	0.38/0.13	0.62/1.15	0.37	52.31	<u>26.50</u>	78.19	50.17	67.81	57.36
Term2Note ($\epsilon_t = 5, \epsilon_n = 8$)	0.15	0.39/0.13	0.64/1.16	0.46	49.52	21.36	79.08	50.11	68.29	60.68
Term2Note ($\epsilon_t = 2, \epsilon_n = 8$)	0.19	0.38/0.12	0.61/1.18	0.38	<u>53.25</u>	24.66	78.85	49.41	68.31	59.91
$\epsilon = 5$										
AUG-PE ($\epsilon_{\text{meta}} = \infty, \epsilon_n = 5$)	11.70	0.11/0.01	7.75/13.58	0.01	48.10	17.44	77.78	53.33	63.01	56.94
FastDP ($\epsilon_{\text{meta}} = \infty, \epsilon_n = 5$)	3.40	0.29/0.09	2.97/5.67	0.04	49.30	16.23	<u>80.54</u>	<u>54.22</u>	<u>67.31</u>	61.98
Term2Note ($\epsilon_{\text{meta}} = \infty, \epsilon_n = 5$)	<u>0.20</u>	<u>0.41/0.14</u>	<u>0.42/1.17</u>	0.39	47.94	20.31	79.29	51.19	66.04	61.69
Term2Note ($\epsilon_t = 5, \epsilon_n = 5$)	<u>0.20</u>	0.38/0.13	0.63/1.19	<u>0.43</u>	54.83	28.96	78.20	50.32	64.56	57.18
Term2Note ($\epsilon_t = 2, \epsilon_n = 5$)	0.36	0.38/0.12	0.64/1.32	0.36	51.26	21.45	79.05	50.44	66.36	60.49
$\epsilon = 2$										
AUG-PE ($\epsilon_{\text{meta}} = \infty, \epsilon_n = 2$)	12.11	0.19/0.03	5.42/11.85	0.01	40.90	13.57	78.29	<u>53.38</u>	63.74	60.10
FastDP ($\epsilon_{\text{meta}} = \infty, \epsilon_n = 2$)	9.67	0.14/0.03	5.79/10.89	0.01	51.06	20.04	<u>79.98</u>	51.31	66.32	59.99
Term2Note ($\epsilon_{\text{meta}} = \infty, \epsilon_n = 2$)	<u>0.43</u>	<u>0.39/0.12</u>	<u>0.48/1.17</u>	<u>0.31</u>	51.78	<u>23.36</u>	79.00	50.60	67.00	59.52
Term2Note ($\epsilon_t = 2, \epsilon_n = 2$)	0.48	0.37/0.12	0.68/1.29	<u>0.31</u>	<u>51.87</u>	23.06	79.43	51.30	<u>69.45</u>	60.31

Table 1: Fidelity and utility evaluation of synthetic datasets generated by different methods using Llama-3.2-1B. Term2Note with $\epsilon_{\text{meta}} = \infty$ indicates that note generation is conditioned on ground-truth labels, i.e., without perturbing the input labels, whereas ϵ_t denotes privatised input labels. Appendix I summarises the privacy budget allocation for the experiments in this table. Utility metrics are averaged across 5-fold cross-validation. The **best overall result** is shown in bold, while the best result under the same privacy budget is underlined. Additional results on Gemma-3-1B and further evaluation details are provided in Appendix J.

constraints using FastDP and guided by prepended control codes; and (2) AUG-PE (Xie et al., 2024), a private evolution-based method that generates synthetic text without model training. For fair comparison, both baselines are adapted to clinical note generation by prepending ICD codes as control codes, consistent with the conditioning setup in TERM2NOTE. We use ϵ_{meta} to denote the privacy budget for ICD codes and ϵ_n for note generation.

Evaluation **Automatic evaluation** of synthetic data typically considers fidelity, privacy, and utility. Although privacy is formally guaranteed by DP, we additionally assess personal information leakage and conduct a preliminary empirical privacy analysis (Appendix G), showing encouraging results.

Fidelity measures how closely synthetic notes resemble real data in structural, syntactic, and semantic aspects. Structural similarity is evaluated using

Kullback–Leibler (KL) divergence over text length distributions. Syntactic similarity is measured by Jaccard similarity and KL divergence over unary and binary clinical term sets, where binary terms denote co-occurring term pairs. Semantic similarity is assessed using MAUVE (Pillutla et al., 2021) with domain-relevant semantic embeddings from BioMistral-7B (Labrak et al., 2024) to compare global distributional similarity.

Utility evaluates the usefulness of synthetic notes for downstream clinical tasks, such as **ICD coding prediction**; evaluation of additional tasks is presented in the Appendix J. ICD-9 and ICD-10 codes are normalised into 20 chapter-level groups (mapping provided in Appendix H). We compare Train-Real-Test-Real and Train-Synthetic-Test-Real settings using micro/macro F1, AUC, and Precision@k. Clinical-Longformer (Li et al., 2023b) is used as the classifier due to its suitability

for long clinical documents. Evaluation is conducted on the private test set $D_{\text{test}}^{\text{src}}$ (204 notes), which is split into training and testing subsets. The classifier is trained on $D_{\text{test-train}}^{\text{src/syn}}$ and evaluated on $D_{\text{test-test}}^{\text{src}}$. Owing to the limited data size, we use 5-fold cross-validation with an 80:20 train-test split.

Human Evaluation is additionally conducted with three licensed physicians to assess the clinical quality of the generated notes. The evaluation follows a pairwise comparison protocol, where each physician is presented with a randomly selected pair of notes and asked to indicate which one is clinically better. The pairs are sampled from outputs generated by three different models under $\epsilon = 8$, as well as from the original (real) notes. Each physician evaluates a minimum of 100 pairs, resulting in a total of 412 pairwise comparisons across all annotators. Based on these annotations, we estimate pairwise model preferences and infer a global ranking using the Bradley-Terry (BT) (Bradley and Terry, 1952) model. The BT model also allows us to estimate the probability that one model $\mathcal{M}_1$ is preferred over another model $\mathcal{M}_2$, based on the aggregated comparison outcomes.

5 Results

5.1 Main Results

Table 1 presents the main results.

Term2Note consistently achieves better structural, syntactic, and semantic similarity to the original data, despite operating under stronger privacy constraints and fewer assumptions. In terms of *structural similarity*, Term2Note achieves the lowest KL divergence in text length distribution (as low as 0.15), indicating faithful preservation of note structure. For *syntactic similarity*, it obtains the (almost) highest Jaccard scores for both unary and binary clinical terms, alongside the lowest KL divergence in term frequency distribution, suggesting close alignment with real clinical content. In *semantic space*, Term2Note substantially outperforms both baseline methods in MAUVE score, confirming that its generated text is significantly more aligned with the original notes in the semantic space. These advantages remain even under a strict privacy budget of $\epsilon = 2$ and when further enforcing DP on clinical term generation (i.e., $\epsilon_t = 2$), underscoring the robustness of the approach. Additionally, Figure 2 shows the distribution of n-gram frequencies, where Term2Note exhibits a distribu-




Figure 2: Distribution of n-gram frequencies in clinical notes generated by different DP methods under $\epsilon = 8$.




Figure 3: Human evaluation results summarised using the BT model. Term2Note (FP) denotes the *full privacy* setting, where both terms and notes are synthesised under DP constraints.

tion more closely aligned with the original notes compared to the baseline methods. More results are presented in Appendix J.

Regarding utility, Term2Note demonstrates strong overall performance and consistently preserves clinical utility across varying privacy budgets, even under stricter constraints on both term and note generation. Across all privacy levels, Term2Note achieves the highest F1 scores, outperforming both AUG-PE and FastDP, and showing the closest performance to the original data. While AUC and Precision@k scores are comparable, rather than uniformly superior, to those of AUG-PE and FastDP, this variation reflects the different aspects of model behaviour captured by each metric. Notably, even when the privacy budget for note generation is reduced to $\epsilon_n = 2$ and additional constraints are applied to privatise clinical terms (with $\epsilon_t = 2$), Term2Note maintains strong utility, with F1 and AUC scores comparable to or better than the baselines operating under looser privacy conditions.

Term2Note is consistently preferred by human experts over AUG-PE, with minimal quality loss under full privacy. As shown in Figure 3,

Method	Fidelity			
	Length	Unary/Binary Term		Semantic
	KL Div. $\downarrow$	Jaccard $\uparrow$	KL Div. $\downarrow$	MAUVE $\uparrow$
$\epsilon_{meta} = \infty, \epsilon_n = 8$				
Term2Note	0.39	0.40/0.13	0.47/ 1.14	0.53
w/o SECSPLIT	7.10	0.38/0.16	2.41/2.05	0.02
w/o DP Maximiser	0.39	0.41/0.13	0.43/ 1.14	0.42
w/ LLM _{ppi} (Meta-Llama-3-8B)	0.24	0.47/0.17	0.38/1.16	0.38
$\epsilon_{meta} = \infty, \epsilon_n = \infty$				
Term2Note	0.25	0.52/0.20	0.22/ 1.08	0.59
Term2Note (Llama-3.3-70b)	0.87	0.55/0.23	0.17/2.23	0.38
RAG	0.75	0.43/0.17	0.62/1.60	0.22

Table 2: Ablation and scaling analysis results.

Term2Note achieves a win rate of 52.2%–53.5% across DP settings, substantially outperforming AUG-PE (12.4%). The pairwise win probabilities (right panel) further demonstrate that Term2Note reliably outperforms AUG-PE across all conditions. Importantly, introducing full privacy constraints, where both clinical terms and notes are protected, has only a marginal effect on human preference. These results suggest that Term2Note maintains high perceived quality while offering stronger privacy guarantees.

Term privatisation maintains a good level of semantic coherence while effectively abstracting away from potentially privacy-leaking details. We evaluate the semantic alignment of the term generation model **TERMGEN** on D_{test}^{src} by computing the cosine similarity between the embeddings of original and generated term lists, using the clinical term encoder EMB. Without the DP mechanism DPRP*, the mean cosine similarity is high (0.82), indicating strong recovery of original terms. When DPRP* is applied, the mean similarity drops to 0.61, reflecting the expected privacy-induced noise. This drop indicates that the generated terms stay semantically coherent without closely matching the originals, reducing the risk of revealing sensitive information (see Appendix K for case studies).

5.2 Ablation and Scaling Analysis

To evaluate the contribution of each component in TERM2NOTE, we conduct an ablation study using Llama-3.2-1B with $\epsilon_{meta} = \infty$ and $\epsilon_n = 8$. Fidelity results are reported in Table 2. Removing SECSPLIT, i.e., generating the full note in a single pass, causes the largest quality drop, reflected by substantially lower MAUVE and higher KL divergence in clinical term distributions, indicating

Repetition	ϵ	Perplexity Rank	Leaked Canaries
100	∞	1/10000	60%
	8	3425/10000	0%

Table 3: Privacy leakage across privacy budgets.

greater deviation from real notes. Removing the DP maximiser or replacing LLM_{ppi} with a general-domain LLM also degrades performance, though less severely. Overall, the results highlight the importance of all three components in TERM2NOTE.

Table 2 also includes a preliminary comparison with Llama-3.3-70B[†] under a non-private setting ($\epsilon_{meta} = \infty, \epsilon_n = \infty$) to justify our choice of model scale. We evaluate two approaches: Retrieval-Augmented Generation (RAG) and LoRA fine-tuning. For RAG, the top-5 most similar sections from the training set D_{train}^{src} are retrieved to assist section-wise generation. For LoRA fine-tuning, the model is quantised to 4-bit precision and trained on 10k samples due to limited computational resources. Although neither approach provides privacy guarantees, they serve as reference points for large-scale models. The results show that LoRA fine-tuning substantially outperforms RAG, demonstrating the importance of parameter adaptation for clinical note generation. The LoRA-finetuned 70B model achieves performance comparable to our fully fine-tuned 1B model and could potentially perform better if trained on the full dataset. However, fine-tuning the 70B model is prohibitively expensive: training on only 10k samples requires more than 80 hours, making large-scale training impractical. In contrast, the 1B model offers a much more practical trade-off between performance and computational cost.

5.3 Leakage of Private Information

Following the canary injection experiment of Yue et al. (2023), we evaluate potential privacy leakage in synthetic notes generated by Term2Note. We inject canary sequences containing sensitive information (e.g., names, addresses, dates of birth, emails, and unit numbers) into the training set, with each repeated 100 times. As shown in Table 3, without DP ($\epsilon = \infty$), three out of five canaries are reproduced in the synthetic notes, yielding a leakage rate of 60%, with an extremely high exposure risk

[†]<https://huggingface.co/meta-llama/Llama-3.3-70B-Instruct>

		Fidelity			
ϵ_t	ϵ_n	Length	Unary/Binary Term		Semantic
		KL Div.↓	Jaccard↑	KL Div.↓	MAUVE↑
∞	∞	1.25	0.47/0.14	0.56/3.47	0.37
∞	8	0.75	0.45/0.18	0.63/2.23	0.44
8	8	0.30	0.44/0.16	0.62/1.79	0.35
5	8	0.36	0.44/0.15	0.61/1.69	0.40

		Utility					
ϵ_t	ϵ_n	F1		AUC		Precision@k	
		Micro	Macro	Micro	Macro	k = 3	k = 5
∞	∞	49.17	23.04	78.77	52.14	66.83	58.63
∞	8	52.77	25.44	79.62	53.33	68.47	59.24
8	8	48.42	21.46	78.55	49.58	63.08	58.23
5	8	53.33	26.28	78.55	50.98	67.67	58.24

Table 4: Fidelity and utility of Term2Note with TERMGEN trained on PMC-Patients.

indicated by a perplexity rank of 1/10000. In contrast, under DP ($\epsilon = 8$), no canaries are leaked, and the perplexity rank decreases substantially to 3425/10000. These results demonstrate the effective privacy protection of our method.

5.4 Robustness to the Choice of Public Dataset

To evaluate the impact of the public dataset used to train TERMGEN in TERM2NOTE, we replace MIMIC-III with PMC-Patients (Zhao et al., 2023), a public dataset containing over 250k patient summaries extracted from PMC articles. As PMC-Patients originates from a different source and exhibits a potentially different data distribution, it provides a useful test of the robustness of our framework. We follow the same term extraction procedure, except that each patient summary is treated as a single section, and reserve 500 summaries for development. The extracted medical terms are then used to train TERMGEN. Results are reported in Table 4. Compared with the corresponding settings in Table 1, where TERMGEN is trained on MIMIC-III, the performance differences are mixed. For example, PMC-Patients achieves a higher MAUVE score when $\epsilon_t = \infty$ and $\epsilon_n = 8$, whereas MIMIC-III performs slightly better under other privacy budgets. Overall, although training TERMGEN on MIMIC-III generally yields better performance, using PMC-Patients produces comparable fidelity and utility, with only small differences across all evaluation metrics.

5.5 Qualitative Analysis

To further assess the quality of synthetic notes, two physicians each reviewed 20 samples generated by Term2Note, covering both standard and full privacy settings with $\epsilon = 5$. As intended by the design of

the evaluation, where *physicians were explicitly asked to identify any potential clinical issues, the feedback focuses on shortcomings rather than general plausibility*. Importantly, not all notes contained identifiable issues. In one physician’s review, 45% were judged to have no clinical problems, indicating that a substantial proportion of generations were considered clinically plausible. **Both physicians agreed that many synthetic notes were plausible as discharge summaries and generally exhibited sound structural organisation.** However, recurring issues emerged around clinical accuracy and coherence. The first physician highlighted problems such as missing or misordered sections, internal inconsistencies (e.g., conflicting medications), and vague or overly generic phrasing; under full privacy constraints, repetition was more common. The second physician, who reviewed a different set of examples, reported more content-level issues, including medication misclassifications (e.g., labelling omeprazole as an antibiotic), illogical or irrelevant narrative insertions, and errors in clinical reasoning. These observations suggest that while Term2Note performs well in preserving structural fidelity, improvements are needed in clinical fact consistency and terminology use. Figure 5 illustrates examples of final note sections, demonstrating the model’s ability to maintain coherence in longer contexts. Although Term2Note occasionally omits sections, its outputs more closely align with the structure of the original note compared to the baseline model.

6 Conclusions

In this paper, we introduce Term2Note, a novel framework for DP clinical note generation that synthesises section-wise content conditioned on medical terms while providing formal privacy guarantees. Experiments show that Term2Note consistently outperforms existing baselines, achieving the highest fidelity and closely matching real notes in structure, semantics, and medical term distributions. Models trained on its synthetic data achieve utility comparable to real data on downstream tasks, while human evaluation shows that clinical experts consistently prefer its outputs over baseline methods. Overall, Term2Note provides a promising solution to healthcare NLP data scarcity through high-quality, privacy-preserving clinical note generation and facilitates privacy-conscious data sharing.

Limitations

Utility metrics exhibit some instability across privacy budgets; performance does not always degrade monotonically as the budget decreases. This may stem from label imbalance relative to training data size, which poses challenges for multi-label prediction when certain labels have limited training examples. Besides, the correlation between clinical factual correctness of the generated synthetic notes and components in Term2Note has not been clearly established due to the limited resources of human experts. Moreover, Term2Note is currently tailored to clinical notes, while other forms of clinical narratives remain unexplored. Nevertheless, it demonstrates the broader applicability of generating clinical content from medical terms, which could be extended to additional clinical text formats in future work. Qualitative feedback also reveals limitations in clinical reliability. In particular, ensuring factual correctness and avoiding clinical inconsistencies (e.g., incorrect medication classes or contradictory clinical states) remain challenging. Future work could integrate clinical fact-checking mechanisms, structured consistency checks (e.g., validating medications against conditions), or lightweight clinical reasoning modules, either during decoding or as post-generation filters. We view these extensions as natural directions to further improve the clinical safety and usefulness of Term2Note.

Ethical Considerations

It is important to note that Term2Note is a research artifact rather than a deployment-ready clinical system. The synthetic clinical notes it generates are intended solely for research purposes and should not be used directly for clinical decision-making without further validation. The formal differential privacy guarantees apply only to the private dataset (i.e., MIMIC-IV in this work) at the note level; they do not provide per-patient privacy guarantees, nor do they apply to the public datasets (i.e., MIMIC-III and PMC-Patients). In addition, TERM2NOTE may generate factual inaccuracies or clinically implausible content. Therefore, rigorous validation, subgroup evaluation, and appropriate human oversight are required before any downstream clinical use. Finally, because the training data are derived from a single healthcare system, the generated notes may reflect demographic or institutional biases. Consequently, evaluation across diverse populations is necessary before broader deployment.

Acknowledgments

We thank the anonymous reviewers and the ethics reviewer of the May 2026 cycle for their feedback. This research is part of the IN-CYPHER programme and is supported by the National Research Foundation, Prime Minister's Office, Singapore under its Campus for Research Excellence and Technological Enterprise (CREATE) programme. This research is also supported by the North West Cyber Security Connect for Commercialisation (NW CyberCom), funded by the Research England Connecting Capability Fund, and UKRI Digital Research Infrastructure Fund, through the FORTRESS project under the DARE UK next-gen catalyst programme. We are grateful for the support provided by Research IT in form of access to the Computational Shared Facility at The University of Manchester and the computational facilities at the Imperial College Research Computing Service[‡].

References

- Martín Abadi, Andy Chu, Ian J. Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. 2016. Deep learning with differential privacy. In *CCS*, pages 308–318. ACM.
- Mehnaz Adnan, Jim Warren, and Martin Orr. 2010. Assessing text characteristics of electronic discharge summaries and their implications for patient readability. In *Proceedings of the Fourth Australasian Workshop on Health Informatics and Knowledge Management*, volume 108, pages 77–84.
- 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 Trans. Comput. Healthcare*, 3(1).
- Yehia Ibrahim Alzoubi and Alok Mishra. 2025. Differential privacy and artificial intelligence: potentials, challenges, and future avenues. *EURASIP J. Inf. Secur.*, 2025(1):18.
- Abhinand Balachandran. 2024. Medembed: Medical-focused embedding models.
- Tal Baumel, Andre Manoel, Daniel Jones, Shize Su, Huseyin A. Inan, Aaron Bornstein, and Robert Sim. 2024. Controllable synthetic clinical note generation with privacy guarantees. *CoRR*, abs/2409.07809.
- Ralph Allan Bradley and Milton E Terry. 1952. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345.

[‡]DOI: <https://doi.org/10.14469/hpc/2232>

- Zhiqi Bu, Yu-Xiang Wang, Sheng Zha, and George Karypis. 2023. Differentially private optimization on large model at small cost. In *ICML*, volume 202 of *Proceedings of Machine Learning Research*, pages 3192–3218. PMLR.
- Soham De, Leonard Berrada, Jamie Hayes, Samuel L. Smith, and Borja Balle. 2022. [Unlocking high-accuracy differentially private image classification through scale](#). *Preprint*, arXiv:2204.13650.
- Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam D. Smith. 2006. Calibrating noise to sensitivity in private data analysis. In *TCC*, volume 3876 of *Lecture Notes in Computer Science*, pages 265–284. Springer.
- Cynthia Dwork and Aaron Roth. 2014. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.*, 9(3-4):211–407.
- Matús Falis, Aryo Pradipta Gema, Hang Dong, Luke Daines, Siddharth Basetti, Michael Holder, Rose S. Penfold, Alexandra Birch, and Beatrice Alex. 2024. Can GPT-3.5 generate and code discharge summaries? *J. Am. Medical Informatics Assoc.*, 31(10):2284–2293.
- James Flemings and Murali Annavaram. 2024. Differentially private knowledge distillation via synthetic text generation. In *ACL (Findings)*, pages 12957–12968. Association for Computational Linguistics.
- YanJun Gao, Dmitriy Dligach, Timothy Miller, Dongfang Xu, Matthew M. Churpek, and Majid Afshar. 2022. Summarizing patients’ problems from hospital progress notes using pre-trained sequence-to-sequence models. In *COLING*, pages 2979–2991. International Committee on Computational Linguistics.
- Lovedeep Gondara and Ke Wang. 2020. Differentially private small dataset release using random projections. In *UAI*, volume 124 of *Proceedings of Machine Learning Research*, pages 639–648. AUAI Press.
- Google. 2025. [Gemma 3](#).
- Mandy Guo, Joshua Ainslie, David Uthus, Santiago Ontanon, Jianmo Ni, Yun-Hsuan Sung, and Yinfei Yang. 2021. Longt5: Efficient text-to-text transformer for long sequences. *arXiv preprint arXiv:2112.07916*.
- William Hardman, Max Banks, Richard Davidson, Daniel Truran, Nadya Wulansari Ayuningtyas, Huy Ngo, Alistair Johnson, and Tom Pollard. 2025. SNOMED CT Entity Linking Challenge (version 1.1.0). <https://doi.org/10.13026/qn8t-6e19>. PhysioNet. RRID:SCR_007345.
- Alistair Johnson, Tom Pollard, Steven Horng, Leo Anthony Celi, and Roger Mark. 2023. MIMIC-IV-Note: Deidentified free-text clinical notes (version 2.2). <https://doi.org/10.13026/1n74-ne17>. PhysioNet. RRID:SCR_007345.
- Alistair EW 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(1):1–9.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. Scaling laws for neural language models. *CoRR*, abs/2001.08361.
- Alexey Kurakin, Natalia Ponomareva, Umar Syed, Liam MacDermed, and Andreas Terzis. 2023. Harnessing large-language models to generate private synthetic text. *CoRR*, abs/2306.01684.
- Sunjun Kweon, Jiyoung Kim, Heeyoung Kwak, Dongchul Cha, Hangyul Yoon, Kwang Kim, Jeewon Yang, Seunghyun Won, and Edward Choi. 2024a. Ehrnoteqa: An LLM benchmark for real-world clinical practice using discharge summaries. In *NeurIPS*.
- Sunjun Kweon, Junu Kim, Jiyoung Kim, Sujeong Im, Eunbyeo 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. 2024b. Publicly shareable clinical large language model built on synthetic clinical notes. In *ACL (Findings)*, pages 5148–5168. Association for Computational Linguistics.
- Yanis Labrak, Adrien Bazoge, Emmanuel Morin, Pierre-Antoine Gourraud, Mickael Rouvier, and Richard Dufour. 2024. [BioMistral: A collection of open-source pretrained large language models for medical domains](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 5848–5864, Bangkok, Thailand. Association for Computational Linguistics.
- Jaewoo Lee and Daniel Kifer. 2020. [Scaling up differentially private deep learning with fast per-example gradient clipping](#). *CoRR*, abs/2009.03106.
- Hao Li, Bowen Deng, Chang Xu, Zhiyuan Feng, Viktor Schlegel, Yu-Hao Huang, Yizheng Sun, Jingyuan Sun, Kailai Yang, Yiyao Yu, and Jiang Bian. 2025a. MIRA: medical time series foundation model for real-world health data. *CoRR*, abs/2506.07584.
- Hao Li, Yu-Hao Huang, Chang Xu, Viktor Schlegel, Ren-He Jiang, Riza Batista-Navarro, Goran Nenadic, and Jiang Bian. 2025b. BRIDGE: bootstrapping text to control time-series generation via multi-agent iterative optimization and diffusion modelling. *CoRR*, abs/2503.02445.
- Hao Li, Yuping Wu, Viktor Schlegel, Riza Batista-Navarro, Thanh-Tung Nguyen, Abhinav Ramesh Kashyap, Xiao-Jun Zeng, Daniel Beck, Stefan Winkler, and Goran Nenadic. 2023a. Team: PULSAR at probsum 2023: PULSAR: pre-training with extracted healthcare terms for summarising patients’ problems and data augmentation with black-box large language models. In *BioNLP@ACL*, pages 503–509. Association for Computational Linguistics.

- Yikuan Li, Ramsey M Wehbe, Faraz S Ahmad, Hanyin Wang, and Yuan Luo. 2023b. A comparative study of pretrained language models for long clinical text. *Journal of the American Medical Informatics Association*, 30(2):340–347.
- Justus Mattern, Zhijing Jin, Benjamin Weggenmann, Bernhard Schölkopf, and Mrinmaya Sachan. 2022. Differentially private language models for secure data sharing. In *EMNLP*, pages 4860–4873. Association for Computational Linguistics.
- Frank McSherry. 2009. Privacy integrated queries: an extensible platform for privacy-preserving data analysis. In *SIGMOD Conference*, pages 19–30. ACM.
- Matthieu Meeus, Lukas Wutschitz, Santiago Zanella-Béguelin, Shruti Tople, and Reza Shokri. 2025. [The canary’s echo: Auditing privacy risks of LLM-generated synthetic text](#). In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 43557–43580, Vancouver, Canada. PMLR.
- Meta. 2024. [Llama 3](#).
- Arvind Narayanan and Vitaly Shmatikov. 2006. How to break anonymity of the netflix prize dataset. *arXiv preprint cs/0610105*.
- Baolin Peng, Chunyuan Li, Pengcheng He, Michel Galley, and Jianfeng Gao. 2023. Instruction tuning with GPT-4. *CoRR*, abs/2304.03277.
- Krishna Pillutla, Swabha Swayamdipta, Rowan Zellers, John Thickstun, Sean Welleck, Yejin Choi, and Zaïd Harchaoui. 2021. MAUVE: measuring the gap between neural text and human text using divergence frontiers. In *NeurIPS*, pages 4816–4828.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, and 1 others. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Krithika Ramesh, Nupoor Gandhi, Pulkit Madaan, Lisa Bauer, Charith Peris, and Anjalie Field. 2024. Evaluating differentially private synthetic data generation in high-stakes domains. In *EMNLP (Findings)*, pages 15254–15269. Association for Computational Linguistics.
- Atiquer Rahman Sarkar, Yao-Shun Chuang, Xiaoqian Jiang, and Noman Mohammed. 2025. Not fully synthetic: Llm-based hybrid approaches towards privacy-preserving clinical note sharing. *AMIA Summits on Translational Science Proceedings*, 2025:441.
- Timo Schick and Hinrich Schütze. 2021. Generating datasets with pretrained language models. In *EMNLP (1)*, pages 6943–6951. Association for Computational Linguistics.
- Viktor Schlegel, Anil A. Bharath, Zilong Zhao, and Kevin Yee. 2025. Generating synthetic data with formal privacy guarantees: State of the art and the road ahead. *CoRR*, abs/2503.20846.
- Viktor Schlegel, Hao Li, Yuping Wu, Anand Subramanian, Thanh-Tung Nguyen, Abhinav Ramesh Kashyap, Daniel Beck, Xiao-Jun Zeng, Riza Theresa Batista-Navarro, Stefan Winkler, and Goran Nenadic. 2023. PULSAR at mediqa-sum 2023: Large language models augmented by synthetic dialogue convert patient dialogues to medical records. In *CLEF (Working Notes)*, volume 3497 of *CEUR Workshop Proceedings*, pages 1668–1679. CEUR-WS.org.
- Luca Soldaini. 2016. [Quickumls: a fast, unsupervised approach for medical concept extraction](#).
- Yidan Sun, Viktor Schlegel, Srinivasan Kolumam Nandakumar, Iqra Zahid, Yuping Wu, Warren Del-Pinto, Goran Nenadic, Lam Siew Kei, Jie Zhang, and Anil Bharath. 2025a. Evaluating differentially private generation of domain-specific text. In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management*, pages 5273–5278.
- Yidan Sun, Viktor Schlegel, Srinivasan Nandakumar, Iqra Zahid, Yuping Wu, Yulong Wu, Hao Li, Jie Zhang, Warren Del-Pinto, Goran Nenadic, and 1 others. 2025b. Synbench: A benchmark for differentially private text generation. *arXiv preprint arXiv:2509.14594*.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. Stanford alpaca: An instruction-following llama model. https://github.com/tatsu-lab/stanford_alpaca.
- U.S. NLM. 2025. UMLS Knowledge Sources [dataset on the internet]. <http://www.nlm.nih.gov/research/umls/licensedcontent/umlsknowledgesources.html>.
- Katharine Weetman, Rachel Spencer, Jeremy Dale, Emma Scott, and Stephanie Schnurr. 2021. What makes a “successful” or “unsuccessful” discharge letter? hospital clinician and general practitioner assessments of the quality of discharge letters. *BMC health services research*, 21(1):349.
- Chulin Xie, Zinan Lin, Arturs Backurs, Sivakanth Gopi, Da Yu, Huseyin A. Inan, Harsha Nori, Haotian Jiang, Huishuai Zhang, Yin Tat Lee, Bo Li, and Sergey Yekhanin. 2024. Differentially private synthetic data via foundation model apis 2: Text. In *ICML*. OpenReview.net.
- Justin Xu. 2024. [Discharge Me: BioNLP ACL’24 Shared Task on Streamlining Discharge Documentation](#). *PhysioNet*. Version 1.3.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao

Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025. Qwen3 technical report. *CoRR*, abs/2505.09388.

Ashkan Yousefpour, Igor Shilov, Alexandre Sablayrolles, Davide Testuggine, Karthik Prasad, Mani Malek, John Nguyen, Sayan Ghosh, Akash Bharadwaj, Jessica Zhao, Graham Cormode, and Ilya Mironov. 2021. *Opacus: User-friendly differential privacy library in pytorch*. *CoRR*, abs/2109.12298.

Da Yu, Huishuai Zhang, Wei Chen, and Tie-Yan Liu. 2021. *Do not let privacy overbill utility: Gradient embedding perturbation for private learning*. In *International Conference on Learning Representations*.

Xiang Yue, Huseyin A. Inan, Xuechen Li, Girish Kumar, Julia McAnallen, Hoda Shajari, Huan Sun, David Levitan, and Robert Sim. 2023. Synthetic text generation with differential privacy: A simple and practical recipe. In *ACL (1)*, pages 1321–1342. Association for Computational Linguistics.

Zhengyun Zhao, Qiao Jin, Fangyuan Chen, Tuorui Peng, and Sheng Yu. 2023. A large-scale dataset of patient summaries for retrieval-based clinical decision support systems. *Scientific data*, 10(1):909.

A Example Demonstration

Figure 4 presents the detailed process in Term2Note with example for each step.

Figure 5 presents examples of generation from different models.

B Algorithm Pseudocode

Algorithm 1 presents the pseudocode for Term2Note.

Algorithm 2 presents the pseudocode for DPRP*.

C Privacy Proof

Proof. We prove that for the full privatisation setting ($T_i = T_i^{\text{syn}}$), Term2Note achieves $(\max(\epsilon_n, \epsilon_t), \max(\delta_n, \delta_t))$ -DP by applying the parallel composition theorem.

Step 1: Parallel Composition Lemma

First, we establish the parallel composition property.

Lemma 3 (Parallel Composition). Let dataset $D = D_1 \cup D_2$, and $D_1 \cap D_2 = \emptyset$. Let $\mathcal{M}_1 : \mathbb{N}^{|\mathcal{X}_1|} \rightarrow R_1$ be (ϵ_1, δ_1) -DP and $\mathcal{M}_2 : \mathbb{N}^{|\mathcal{X}_2|} \rightarrow R_2$ be (ϵ_2, δ_2) -DP. Then $\mathcal{M}(D) = (\mathcal{M}_1(D_1), \mathcal{M}_2(D_2))$ is $(\max(\epsilon_1, \epsilon_2), \max(\delta_1, \delta_2))$ -DP.

Proof of Lemma. Let D and D' be neighboring datasets differing by one record. Since $D_1 \cap D_2 =$

$\emptyset$, the differing record is in either D_1 or D_2 , but not both.

Case 1: The differing record is in D_1 , so $D_1 \neq D'_1$ but $D_2 = D'_2$.

For any measurable sets $B_1 \subseteq R_1, B_2 \subseteq R_2$:

$$\begin{aligned} \Pr[\mathcal{M}(D) \in B_1 \times B_2] &= \Pr[\mathcal{M}_1(D_1) \in B_1] \cdot \Pr[\mathcal{M}_2(D_2) \in B_2] \\ &\leq (e^{\epsilon_1} \Pr[\mathcal{M}_1(D'_1) \in B_1] + \delta_1) \cdot \Pr[\mathcal{M}_2(D_2) \in B_2] \\ &= (e^{\epsilon_1} \Pr[\mathcal{M}_1(D'_1) \in B_1] + \delta_1) \cdot \Pr[\mathcal{M}_2(D'_2) \in B_2] \\ &= e^{\epsilon_1} \Pr[\mathcal{M}(D') \in B_1 \times B_2] + \delta_1 \Pr[\mathcal{M}_2(D'_2) \in B_2] \\ &\leq e^{\epsilon_1} \Pr[\mathcal{M}(D') \in B_1 \times B_2] + \delta_1 \end{aligned}$$

Case 2: The differing record is in D_2 , so $D_1 = D'_1$ but $D_2 \neq D'_2$. Similarly:

$$\begin{aligned} \Pr[\mathcal{M}(D) \in B_1 \times B_2] &\leq e^{\epsilon_2} \Pr[\mathcal{M}(D') \in B_1 \times B_2] + \delta_2 \end{aligned}$$

Combining cases: For arbitrary neighbouring datasets, we have:

$$\begin{aligned} \Pr[\mathcal{M}(D) \in B_1 \times B_2] &\leq e^{\max(\epsilon_1, \epsilon_2)} \Pr[\mathcal{M}(D') \in B_1 \times B_2] + \max(\delta_1, \delta_2) \end{aligned}$$

Therefore, $\mathcal{M}$ is $(\max(\epsilon_1, \epsilon_2), \max(\delta_1, \delta_2))$ -DP. $\square$

Step 2: Application to Term2Note

Now we apply the parallel composition lemma to Term2Note.

We have:

- $\mathcal{M}_1 = \text{NOTEGEN}$ training on D_{train} , which is (ϵ_n, δ_n) -DP
- $\mathcal{M}_2 = \text{TERMGEN}$ processing on D_{test} , which is (ϵ_t, δ_t) -DP
- $D_{\text{train}} \cap D_{\text{test}} = \emptyset$

Term2Note can be written as:

$$\text{Term2Note}(D) = f(\mathcal{M}_1(D_{\text{train}}), \mathcal{M}_2(D_{\text{test}}))$$

where f is a deterministic function that applies the trained NOTEGEN model to the synthetic terms from TERMGEN.




Figure 4: Example to demonstrate Term2Note.

Term2Note	Original	FastDP
Medications on Admission: The Preadmission Medication List is accurate and complete. 1. ibuprofen 600 mg PO Q6H 2. acetaminophen 650 mg PO Q8H 3. furosemide 40 mg PO DAILY 4. capsaipine extended-release 50 (...) 5. albuterol inhaler 1 inhaler q4h 6. metoprolol succinate XL (...) Discharge Medications: 1. Albuterol Inhaler 1 inhaler q4h 2. Capsaicine Extended-Release Tablets 50 mg PO BID 3. Metoprolol Succinate XL Tablet Oral Q24 4. Acetaminophen 325 (...) 4. Amoxicillin CLINICALLY CONFIRMED (PRN)(...) RX furosemide sodium hydrochloride ___ mg PO DAILY Disp #60 (...) Discharge Disposition: Home With Service: Service: Primary Care/Family Practice Disposition Code: ERX Facilities: ___	Medications on Admission: Aspirin X mg daily; Omeprazole X mg daily Saline Nasal spray daily; (...) Discharge Medications: Patient expired Discharge Disposition: Expired Discharge Diagnosis: Hypoxic respiratory failure; Anuric renal failure Atrial fibrillation with RVR; Delirium Discharge Condition: Patient expired after having been transitioned to CMO after progressive multiorgan system failure prompted discussion with his family and HCP. Discharge Instructions: Patient expired. Followup Instructions: ___	Medications That Need To Be Taken After Hospitalization: [list specific medication names and dosages] Instructions for Care: [insert instructions on care after hospitalization] Follow-up Schedule: [Insert follow-up schedule] Special Instructions: [Insert special instructions] Other Notes: [Any other important information you want included in this document.]

Figure 5: Examples of the last section in clinical notes generated by different models. For illustration purposes, some content is redacted with “(...)”, and numeric values in the original note are de-identified. No other modifications were made.

Algorithm 1 Term2Note

Input: Private note X^{src} , public term lists D_{public} , privacy parameters for term generation (ϵ_t, δ_t) and note generation (ϵ_n, δ_n) , #sections m , instruction I , noise variance σ_{emb} , #candidates k

Output: Synthetic note X^{syn} with (ϵ, δ) -DP guarantee

```
1:  $\text{SEC}_{1:m}^{\text{src}} \leftarrow \text{SECSPLIT}(X^{\text{src}})$  // 3.2 Format and Term Identification
2:  $T_{1:m}^{\text{src}} \leftarrow \text{TERMEXT}(\text{SEC}_{1:m}^{\text{src}})$ 
3:  $E_{1:m}^{\text{src}}, E_{\text{public}} \leftarrow \text{EMB}(T_{1:m}^{\text{src}}), \text{EMB}(D_{\text{public}})$  // 3.3 Clinical Terms Generation
4: if train then
5:    $\theta_t, \theta_p \leftarrow \mathcal{L}(\text{TERMGEN}_{\theta_t}(\text{PROJ}_{\theta_p}(E_{\text{public}} + \mathcal{N}(0, \sigma_{\text{emb}}))), D_{\text{public}})$  // Train on  $D_{\text{public}}$ 
6: else
7:    $T_{1:m}^{\text{syn}} \sim \text{TERMGEN}_{\theta_t}(\text{PROJ}_{\theta_p}(\text{DPRP}^*(E_{1:m}^{\text{src}}, \frac{\epsilon_t}{m}, \frac{\delta_t}{m})))$  // Infer for  $X^{\text{src}}$ 
8: end if
9: if train then
10:   $\theta_n \leftarrow \text{FastDP}(\mathcal{L}(\text{NOTEGEN}_{\theta_n}(I, \text{SEC}_{<i}^{\text{src}}, T_i^{\text{src}}), X^{\text{src}}), \epsilon_n, \delta_n)$  // 3.4 Clinical Note Generation
11: end if
12:  $\text{SEC}_j^{\text{syn}}[i] \sim \text{NOTEGEN}_{\theta_n}(I, \text{SEC}_{<j}^{\text{syn}}[i], T_j^{\text{src} / \text{syn}})$  for  $i = 1, \dots, k; j = 1, \dots, m$ 
13:  $X^{\text{syn}} \leftarrow \arg \min_{i \in [k]} \text{PPL}(\text{SEC}_{1:m}^{\text{syn}}[i])$  // 3.5 DP Quality Maximiser
14: return  $X^{\text{syn}}$ 
```

Algorithm 2 DPRP*

Input: Embeddings E , privacy parameters (ϵ, δ) , privacy allocation $b = 0.85$

Output: Privatised Embeddings E_{DP} with (ϵ, δ) -DP

```
1:  $(\epsilon_1, \delta_1), (\epsilon_2, \delta_2) \leftarrow 0.85 * (\epsilon, \delta), 0.15 * (\epsilon, \delta)$ 
2: Derive  $\sigma_i$  from  $(\epsilon_{1i}, \delta_{1i}); i \in [1, 2]$ 
3:  $E' = E + \mathcal{N}(0, \sigma_1^2)$ 
4:  $E'_C = E'^T E + \mathcal{N}(0, \sigma_2^2)$ 
5:  $V' \Sigma' V'^T = \text{SVD}(E'_C)$ 
6:  $V'_k = V'[1, \dots, k]; k = 0.6 * E_{\text{hdim}}$ 
7:  $E_{\text{DP}} = E' V_k'^T + V_k'^T$  // + refers to the Moore-Penrose pseudoinverse
8: return  $E_{\text{DP}}$ 
```

Since f is a post-processing function applied to the outputs of the parallel composition, and post-processing preserves differential privacy, we have:

Term2Note(D) is $(\max(\epsilon_n, \epsilon_t), \max(\delta_n, \delta_t))$ -DP

□

D Section Grouping

Table 5 presents the section grouping taxonomy for our SECSPLIT.

E Dataset Statistics

Table 6 summarises key statistics of the three datasets used in this study, including the number of clinical notes, and average note length, among other relevant attributes.

F Hyperparameters

TERMGEN We fine-tune GPT-2-large on section-wise clinical terms extracted from D_{public} for up to 5 epochs. The final model is selected based on the checkpoint with the highest F1 score, evaluated on a held-out set of 500 notes. During training, we set the embedding perturbation scale $\sigma_{\text{emb}} = 0.05$, with a batch size of 8 and a learning rate of $2e-5$. At inference, we use a batch size of 16 and a maximum generation length of 512 tokens. To ensure reproducibility, decoding is performed with a temperature of 0.1 and top- p set to 1.0.

NOTEGEN Training: We fine-tune Llama-3.2-1B-Instruct or Gemma-3-1B-IT on $D_{\text{train}}^{\text{src}}$ for up to 2 epochs using 2 GPUs. The batch size per device is 2, with a gradient accumulation step of 64 and a learning rate of $5e-5$. We enable DeepSpeed ZeRO Stage 3 to optimise memory usage. **Inference:** We

Group Name	Sections
Patient Information	“Name”, “Unit No”, “Admission Date”, “Discharge Date”, “Date of Birth”, “Sex”, “Service”, “Allergies”, “Attending”
Clinical Course & History	“Chief Complaint”, “Major Surgical or Invasive Procedure”, “History of Present Illness”, “Review of Systems”, “Past Medical History”, “Social History”, “Family History”
Examinations & Findings	“Physical Exam”
Laboratory & Imaging Results	“Pertinent Results”
Hospital Stay & Treatment	“Brief Hospital Course”
Medications & Discharge Plan	“Medications on Admission”, “Discharge Medications”, “Discharge Disposition”, “Discharge Diagnosis”, “Discharge Condition”, “Discharge Instructions”, “Followup Instructions”

Table 5: The grouped section titles.

adapt vLLM for faster generation, with decoding parameters set as temperature = 0.1, top- p = 1.0, repetition penalty = 1.2 and max tokens per section = 2048 across all experiments. Llama-3.2 tends to generate overly long outputs during section-wise generation, so we apply a logit bias on the EOS token to encourage early stopping. This bias is set between 0.5 and 6.0: for DP-enabled models, the value is 0.5 or 1.0; for the non-private setting ($\epsilon = \infty$), it is set to 6.0. Additionally, DP-enabled models use a frequency penalty of 0.4 to further discourage repetition.

For the FastDP baseline, which produces relatively short outputs, we apply only a repetition penalty during inference. Before applying the DP quality maximiser, we generate multiple candidates per input: 4 for Term2Note and FastDP, and 7 for AUG-PE, using the same decoding settings described above.

Corpus	MIMIC-III	MIMIC-IV	
Dataset	D_{public}	$D_{\text{train}}^{\text{src}}$	$D_{\text{test}}^{\text{src}}$
# notes	52,722	122,202	204
avg. # tokens	3327.93	3360.60	2818.63
avg. # sections	4.59	5.77	5.79
avg. # terms	176.26	203.20	173.33
avg. # ICD codes	-	-	6.73

Table 6: Dataset statistics. avg. refers to the average of. # tokens is calculated by taking the average of tokens in each note, tokenised by Llama-3.2-1B-Instruct.

G Distance-based Privacy Evaluation

A preliminary privacy evaluation is conducted to assess the privacy-preserving properties of the synthetic clinical notes using a membership inference attack (MIA) framework. We compare the distribution of minimum cosine distances between the synthetic notes generated by Term2Note and the original training data to a baseline of real, non-member test notes. As illustrated in Figure 6, the synthetic notes are, on average, located significantly further from the training data than the test notes. Notably, the setting with $\epsilon_n = \infty$, which is generated without DP, exhibited the most overlap with the test set’s distribution. This expected outcome highlights the privacy benefits of the other DP settings, and provides a clear baseline for comparison.

In future work, we intend to expand our privacy evaluation using canary-based membership inference attacks. This approach involves injecting specially crafted canaries into the training data to establish a worst-case lower bound on privacy risks.

H ICD Codes Grouping

Table 7 shows the mapping between our combined ICD categories and the corresponding ICD-9 and ICD-10 chapter headings. For fine-tuning Clinical-Longformer on this classification task, we train for 30 epochs per setting (i.e., model and data fold), with a batch size of 8 and a learning rate of 2e-5.

I Privacy Accounting

Table 8 presents the summary of how privacy budgets are allocated across components.




Figure 6: Distribution of minimum cosine distances for all evaluated synthetic strategies compared to the baseline of original test notes. Here, ϵ_t and ϵ_n are the privacy budgets for TERMGEN and NOTEGEN, respectively.

J Supplementary Experimental Results

J.1 Additional Downstream Tasks

To further evaluate the utility of the synthetic notes, we conduct experiments on two additional downstream tasks.

Continued pre-training. We perform a small-scale continued pre-training experiment using dmis-lab/biobert-base-cased-v1.2 on the synthetic notes generated by Term2Note. The resulting model (“BioBERT-tuned”) is evaluated on three biomedical benchmarks: named entity recognition (BC5CDR), relation extraction (GAD), and question answering (BioASQ 7b), and compared with the original BioBERT. Tables 9 and 10 show that continued pre-training consistently improves F1 on both NER and relation extraction, demonstrating that the synthetic notes provide useful domain-specific supervision.

Summarisation. We further evaluate the synthetic notes on the BioNLP 2024 Shared Task 2 (Discharge Me!) (Xu, 2024). Following the shared-task setup, the “Brief Hospital Course” (BHC) section is treated as the reference summary and removed from each clinical note. Term2Note is trained on the remaining content to generate synthetic notes. Owing to the large training set (over 36k notes) and limited computational resources, the

DP quality maximiser is omitted in this experiment. We fine-tune google/long-t5-tglobal-large (Guo et al., 2021) on either the original or synthetic training set and evaluate on the original test set. As shown in Table 11, models trained on synthetic notes perform worse than those trained on the original data across all metrics. However, the performance gaps remain relatively small on BLEU, ROUGE-2, and BERTScore, suggesting that the synthetic notes preserve much of the information required for summarisation.

J.2 Fidelity

Text Length Comparison under $\epsilon = 8$ Figure 7 shows the distribution of sequence lengths for clinical notes generated by different methods under a fixed privacy budget of $\epsilon = 8$. While all synthetic methods shift the length distribution away from the original data to some extent, Term2Note exhibits the closest alignment. Its distribution more faithfully captures the broad length range of the original notes than the baselines. In contrast, AUG-PE produces much shorter and more narrowly distributed sequences, indicating a loss of structural richness. FastDP also generates relatively short sequences, with a sharp peak around 500 tokens. These deviations suggest that Term2Note is better able to preserve the structural properties of real clinical notes, which is crucial for downstream utility and

Combined ICD Category	ICD-9	ICD-10
Certain Infectious And Parasitic Diseases	Infectious And Parasitic Diseases	Certain Infectious And Parasitic Diseases
Neoplasms	Neoplasms	Neoplasms
Endocrine, Nutritional And Metabolic Diseases, And Immunity Disorders	Endocrine, Nutritional And Metabolic Diseases, And Immunity Disorders	Endocrine, Nutritional And Metabolic Diseases
Diseases Of The Blood And Blood-Forming Organs And Certain Disorders Involving The Immune Mechanism	Diseases Of The Blood And Blood-Forming Organs	Diseases Of The Blood And Blood-Forming Organs And Certain Disorders Involving The Immune Mechanism
Mental And Behavioural Disorders	Mental Disorders	Mental And Behavioural Disorders
Diseases Of The Nervous System And Sense Organs	Diseases Of The Nervous System And Sense Organs	Diseases Of The Nervous System
Diseases Of The Circulatory System	Diseases Of The Circulatory System	Diseases Of The Circulatory System
Diseases Of The Respiratory System	Diseases Of The Respiratory System	Diseases Of The Respiratory System
Diseases Of The Digestive System	Diseases Of The Digestive System	Diseases Of The Digestive System
Diseases Of The Genitourinary System	Diseases Of The Genitourinary System	Diseases Of The Genitourinary System
Complications Of Pregnancy, Childbirth, And The Puerperium	Complications Of Pregnancy, Childbirth, And The Puerperium	Complications Of Pregnancy, Childbirth, And The Puerperium
Diseases Of The Skin And Subcutaneous Tissue	Diseases Of The Skin And Subcutaneous Tissue	Diseases Of The Skin And Subcutaneous Tissue
Diseases Of The Musculoskeletal System And Connective Tissue	Diseases Of The Musculoskeletal System And Connective Tissue	Diseases Of The Musculoskeletal System And Connective Tissue
Congenital Malformations, Deformations And Chromosomal Abnormalities	Congenital Anomalies	Diseases Of The Musculoskeletal System And Connective Tissue
Congenital Malformations, Deformations And Chromosomal Abnormalities	-	Congenital Malformations, Deformations And Chromosomal Abnormalities
Certain Conditions Originating In The Perinatal Period	Certain Conditions Originating In The Perinatal Period	Certain Conditions Originating In The Perinatal Period
Symptoms, Signs And Abnormal Clinical And Laboratory Findings, Not Elsewhere Classified	Symptoms, Signs, And Ill-Defined Conditions	Symptoms, Signs And Abnormal Clinical And Laboratory Findings, Not Elsewhere Classified
Injury, Poisoning And Certain Other Consequences Of External Causes	Injury And Poisoning	Injury, Poisoning And Certain Other Consequences Of External Causes
External Causes Of Morbidity And Mortality, Injury and Poisoning	External Causes Of Injury And Poisoning	External Causes Of Morbidity And Mortality, Injury and Poisoning
Factors Influencing Health Status And Contact With Health Services	Factors Influencing Health Status And Contact With Health Services	Factors Influencing Health Status And Contact With Health Services
Diseases Of The Eye And Adnexa	-	Diseases Of The Eye And Adnexa
Diseases Of The Ear And Mastoid Process	-	Diseases Of The Ear And Mastoid Process
Codes For Special Purposes	-	Codes For Special Purposes

Table 7: The grouped ICD codes.

Term Generator	Note Generator	ϵ_n	ϵ_t	ϵ (total)
×	✓	-	∞	∞
✓	✓	$\infty/8/5/2$	∞	∞
✓	✓	8	$8/5/2$	8
✓	✓	5	$5/2$	5
✓	✓	2	2	2

Table 8: Summary of privacy accounting.

realism in synthetic data.

DP Quality Maximiser Beyond perplexity, we investigate a range of reference-free (RF) metrics to guide the selection of high-quality generations, including maximum and mean sentence length (in words and characters), self-BLEU, and distinct- n variants. To evaluate these metrics, we manually annotate a small set of synthetic sections as “good”




Figure 7: Text length distribution.

Task-Dataset	Model	Precision	Recall	F1
NER-BC5CDR	BioBERT	85.23	88.26	86.67
	BioBERT-tuned	85.78	88.24	86.98
RL-GAD	BioBERT	80.00	79.23	79.34
	BioBERT-tuned	82.13	80.81	80.95

Table 9: NER and relation extraction results.

Task-Dataset	Model	Strict Accuracy	Lenient Accuracy	Mean Reciprocal Rank
QA-BioASQ 7b	BioBERT	65.55	97.21	78.99
	BioBERT-tuned	62.94	97.21	77.21

Table 10: Question answering results.

or “bad” based on readability, with approximately 12% labelled as “bad”. Metrics are evaluated on their ability to identify these poor-quality sections via scalar thresholds. We then assess how well each metric identifies poor-quality sections using scalar thresholds. Our results indicate that metrics based on sentence length—particularly maximum sentence character count—align most closely with human annotations. A rejection threshold of 2181 characters yields the strongest correspondence. Table 12 reports the KL divergence and MAUVE values of synthetic notes after integrating this metric into the inference process. Specifically, if a generated section exceeds the threshold, it is discarded and regenerated until acceptance. Incorporating this simple criterion yields measurable improvements, suggesting that lightweight, reference-free filters can enhance the realism of DP synthetic text. Future work may extend this approach by combining multiple RF metrics for greater robustness.

J.3 Utility

Table 13 presents the detailed precision and recall scores for the downstream task evaluation.

J.4 Gemma

Experimental results for Gemma are reported in Table 14. Both fidelity and utility metrics are comparable to those of Llama in Table 1, although the MAUVE score for Gemma without DP (i.e., $\epsilon = \infty$) is higher than that of Llama. Overall, the same trend holds across both models: stricter privacy guarantees lead to reduced fidelity, while full privatisation still preserves strong fidelity and utility.

K Case Studies for Term Generation

Table 15 presents examples of synthetic clinical terms generated with and without the application

Method	BLEU	ROUGE-1/2/L	BERTScore	Meteor
Original Data	6.51	27.28/7.32/15.48	50.85	18.83
Term2Note ($\epsilon_n = \infty$)	6.39	13.06/5.35/9.19	46.09	10.56
Term2Note ($\epsilon_n = 8$)	4.93	15.72/5.15/10.74	48.03	11.36

Table 11: Utility evaluation on BioNLP 2024 Shared Task 2 Discharge Me!.

Dataset	KL Divergence↓	MAUVE↑
w/o RF metric	1.99±0.37	0.24±0.04
w/ RF metric	1.03±0.10	0.36±0.07

Table 12: Fidelity evaluation of synthetic notes generated with and without integrating the RF metric (maximum sentence character count) into the inference process.

of DPRP*. The original list contains five salient terms extracted from a real clinical note. When no DP is applied, the generated list recovers only two of these terms (“air” and “discharge”), suggesting limited coverage despite the absence of privacy constraints. In contrast, the DP-enabled output does not directly replicate any of the original terms beyond “discharge”, but instead generates a substantially longer and more diverse list of medically plausible terms.

This illustrates a key trade-off: the DP mechanism introduces sufficient variability to obscure direct term recovery, thus enhancing privacy protection. At the same time, the generated list remains semantically coherent and clinically relevant, containing realistic phrases such as “hemodynamically stable,” “chronic low back pain,” and “pulmonary vein,” which contribute to the naturalness and utility of the resulting synthetic note. These findings align with our earlier quantitative analysis, confirming that DPRP* balances semantic fidelity with privacy-preserving diversity.

Method	F1		Precision		Recall		AUC		Precision@k	
	Micro	Macro	Micro	Macro	Micro	Macro	Micro	Macro	k = 3	k = 5
Original Data	57.03±3.59	30.80±2.20	60.14±4.51	34.66±4.56	54.40±4.26	30.88±2.66	82.01±1.36	58.88±2.73	68.93±4.53	62.14±3.20
$\epsilon = \infty$										
AUG-PE ($\epsilon_{\text{meta}} = \infty, \epsilon_n = \infty$)	45.82±2.33	14.84±2.97	67.30±5.10	17.91±5.33	34.94±3.19	16.27±2.67	79.52±1.23	54.35±1.75	68.48±6.11	61.77±3.03
Fast-DP ($\epsilon_{\text{meta}} = \infty, \epsilon_n = \infty$)	53.02±3.03	25.51±1.92	56.19±5.62	26.83±2.31	50.99±6.67	27.27±4.23	79.35±1.55	51.05±2.48	69.77±4.99	60.39±4.44
Term2Note ($\epsilon_{\text{meta}} = \infty, \epsilon_n = \infty$)	49.95±4.77	21.89±3.90	65.37±6.31	28.69±1.26	41.02±6.64	21.05±3.86	81.40±1.78	55.43±2.40	69.77±3.52	61.96±5.70
Term2Note ($\epsilon_t = \infty, \epsilon_n = \infty$)	49.24±2.63	21.9±2.01	61.07±4.02	26.75±2.9	41.61±4.73	21.64±3.29	80.24±1.56	51.54±1.81	67.81±3.66	61.08±3.71
Term2Note ($\epsilon_t = \infty, \epsilon_n = 8$)	48.16±4.39	20.63±3.18	59.07±4.04	26.64±5.14	40.84±5.3	20.68±3.41	78.72±0.85	50.01±4.19	65.35±5.21	58.72±3.9
Term2Note ($\epsilon_t = \infty, \epsilon_n = 5$)	51.0±1.41	22.73±2.22	56.7±4.95	24.89±4.24	46.79±3.93	24.62±3.03	78.8±1.78	50.37±1.78	67.15±5.88	59.91±3.12
Term2Note ($\epsilon_t = \infty, \epsilon_n = 2$)	48.57±0.89	20.31±1.25	59.92±5.19	24.63±2.69	41.23±3.46	20.75±2.05	79.56±1.08	51.75±0.85	68.64±3.55	60.41±3.43
$\epsilon = 8$										
AUG-PE ($\epsilon_{\text{meta}} = \infty, \epsilon_n = 8$)	40.73±9.22	13.28±4.71	60.68±4.70	15.14±4.77	31.43±9.80	16.01±5.81	78.13±1.85	53.49±3.08	63.24±4.65	59.03±4.59
FastDP ($\epsilon_{\text{meta}} = \infty, \epsilon_n = 8$)	48.58±5.93	16.40±4.01	64.79±1.33	16.70±3.67	39.49±8.78	18.89±5.16	80.74±1.41	51.57±3.46	69.79±2.25	61.59±4.42
Term2Note ($\epsilon_{\text{meta}} = \infty, \epsilon_n = 8$)	49.71±1.90	21.28±1.28	61.37±4.28	25.17±1.55	41.96±2.98	21.31±1.63	80.03±1.44	52.80±1.15	67.49±4.42	61.48±4.45
Term2Note ($\epsilon_t = 8, \epsilon_n = 8$)	52.31±4.45	26.5±3.95	53.04±3.34	26.9±2.62	51.87±6.45	28.66±4.65	78.19±0.91	50.17±2.33	67.81±5.88	57.36±3.3
Term2Note ($\epsilon_t = 5, \epsilon_n = 8$)	49.52±4.44	21.36±3.11	58.74±6.32	24.59±5.36	43.23±5.68	22.04±3.52	79.08±1.34	50.11±4.08	68.29±5.36	60.68±4.65
Term2Note ($\epsilon_t = 2, \epsilon_n = 8$)	53.25±1.1	24.66±1.82	56.93±2.49	28.23±4.52	50.16±2.52	26.94±2.17	78.85±0.98	49.41±2.86	68.31±5.53	59.91±3.34
$\epsilon = 5$										
AUG-PE ($\epsilon_{\text{meta}} = \infty, \epsilon_n = 5$)	48.10±2.08	17.44±2.70	60.59±8.29	18.28±3.69	40.73±5.59	21.02±4.88	77.78±3.11	53.33±2.60	63.01±12.42	56.94±5.20
Fast-DP ($\epsilon_{\text{meta}} = \infty, \epsilon_n = 5$)	49.30±4.04	16.23±3.04	64.49±6.50	15.57±2.73	40.84±7.33	19.70±4.87	80.54±1.36	54.22±1.97	67.31±7.22	61.98±3.34
Term2Note ($\epsilon_{\text{meta}} = \infty, \epsilon_n = 5$)	47.94±4.47	20.31±3.40	60.88±4.95	25.76±4.44	40.15±6.22	20.20±3.71	79.29±1.49	51.19±3.40	66.04±5.73	61.69±4.86
Term2Note ($\epsilon_t = 5, \epsilon_n = 5$)	54.83±2.24	28.96±1.81	51.64±4.44	29.07±4.58	58.92±4.33	34.06±2.99	78.2±0.8	50.32±3.94	64.56±4.1	57.18±3.09
Term2Note ($\epsilon_t = 2, \epsilon_n = 5$)	51.26±1.97	21.45±1.88	58.92±4.45	25.08±3.48	45.63±3.55	23.28±1.99	79.05±1.37	50.44±3.5	66.36±5.89	60.49±4.89
$\epsilon = 2$										
AUG-PE ($\epsilon_{\text{meta}} = \infty, \epsilon_n = 2$)	40.9±7.43	13.57±3.88	64.75±7.31	15.2±4.82	30.7±9.22	14.83±5.29	78.29±0.8	53.38±1.79	63.74±5.0	60.1±4.37
Fast-DP ($\epsilon_{\text{meta}} = \infty, \epsilon_n = 2$)	51.06±5.7	20.04±4.34	58.78±5.13	19.77±4.27	45.94±8.57	23.72±5.52	79.98±1.95	51.31±3.32	66.32±6.66	59.99±5.35
Term2Note ($\epsilon_{\text{meta}} = \infty, \epsilon_n = 2$)	51.78±3.99	3.36±3.87	57.45±5.97	25.59±4.84	47.21±3.16	25.26±2.64	79.00±1.67	50.60±2.75	67.00±3.39	59.52±5.59
Term2Note ($\epsilon_t = 2, \epsilon_n = 2$)	51.87±2.73	23.06±2.51	57.77±6.18	26.58±5.04	47.37±2.88	24.77±2.01	79.43±1.41	51.3±2.96	69.45±4.73	60.31±5.39

Table 13: Supplementary results for utility evaluation: F1, Precision, Recall, AUC, and Precision@k, with **mean±standard deviation** values reported.

Method	Fidelity				Utility					
	Length	Unary/Binary Term		Semantic	F1		AUC		Precision@k	
	KL Div.↓	Jaccard↑	KL Div.↓	MAUVE↑	Micro	Macro	Micro	Macro	k = 3	k = 5
Original Data					57.03	30.80	82.01	58.88	68.93	62.14
$\epsilon = \infty$										
Term2Note ($\epsilon_{\text{meta}} = \infty, \epsilon_n = \infty$)	0.18	0.61/0.34	0.23/1.55	0.80	52.56	25.32	80.81	55.09	67.30	60.78
Term2Note ($\epsilon_t = \infty, \epsilon_n = \infty$)	0.36	0.40/0.18	0.77/1.52	0.66	53.77	25.72	79.94	51.86	66.98	60.71
Term2Note ($\epsilon_t = \infty, \epsilon_n = 8$)	0.40	0.37/0.13	0.64/1.82	0.38	47.63	20.19	78.87	50.5	66.83	59.53
Term2Note ($\epsilon_t = \infty, \epsilon_n = 5$)	0.25	0.37/0.13	0.65/1.76	0.32	49.98	21.50	79.77	52.57	68.14	59.91
Term2Note ($\epsilon_t = \infty, \epsilon_n = 2$)	0.31	0.36/0.13	0.69/1.59	0.35	48.11	20.87	79.49	52.83	65.51	57.17
$\epsilon = 8$										
Term2Note ($\epsilon_{\text{meta}} = \infty, \epsilon_n = 8$)	0.40	0.39/0.13	0.53/1.66	0.48	48.41	20.89	78.55	48.89	67.50	59.31
Term2Note ($\epsilon_t = 8, \epsilon_n = 8$)	0.40	0.37/0.13	0.75/2.05	0.32	55.46	26.73	79.68	49.92	69.30	60.60
Term2Note ($\epsilon_t = 5, \epsilon_n = 8$)	0.41	0.38/0.13	0.70/1.71	0.38	51.82	21.67	80.44	51.28	70.92	60.88
Term2Note ($\epsilon_t = 2, \epsilon_n = 8$)	0.51	0.36/0.12	0.76/1.89	0.30	52.34	22.50	80.02	51.71	69.62	60.01
$\epsilon = 5$										
Term2Note ($\epsilon_{\text{meta}} = \infty, \epsilon_n = 5$)	0.36	0.38/0.13	0.56/1.52	0.31	51.20	21.18	80.03	53.10	67.65	61.38
Term2Note ($\epsilon_t = 5, \epsilon_n = 5$)	0.39	0.36/0.12	0.75/1.79	0.32	52.79	23.74	79.54	49.64	67.99	58.92
Term2Note ($\epsilon_t = 2, \epsilon_n = 5$)	0.41	0.36/0.13	0.77/1.83	0.21	55.06	27.86	79.42	49.19	69.30	59.82
$\epsilon = 2$										
Term2Note ($\epsilon_{\text{meta}} = \infty, \epsilon_n = 2$)	0.49	0.36/0.12	0.60/1.65	0.27	49.26	21.17	79.21	50.33	67.16	59.24
Term2Note ($\epsilon_t = 2, \epsilon_n = 2$)	0.46	0.35/0.12	0.80/1.78	0.31	53.66	24.43	79.91	50.26	69.28	61.08

Table 14: Fidelity and utility evaluation of synthetic datasets generated by Term2Note with Gemma-3-1B as the base model for NOTEGEN.

Method	Terms
Original	["Physical", "Discharge", "Laparoscopic", "incisions", "air"]
No DP	["air", "discharge"]
DP	["brief", "discharge", "negative", "medications", "placement", "drainage", "hemodynamically stable", "therapy", "chronic low back pain", "symptoms", "right chest", "referred to cardiac surgery", "chest discomfort", "pulmonary vein", "hyperlipidemia: he", "hypertension-", "difficulty", "surgical service", "anticoagulation", "discontinued", "increased", "afebrile", "asymptomatic", "admission", "intervention", "hospitalization", "cardiac enzymes x3"]

Table 15: Example of synthetic terms.