

Publicly Shareable Clinical Large Language Model Built on Synthetic Clinical Notes

Sunjun Kweon^{1*}, Junu Kim^{1*}, Jiyou Kim¹, Sujeong Im¹, Eunbyeol Cho¹
Seongsu Bae¹, Jungwoo Oh¹, Gyubok Lee¹, Jong Hak Moon¹, Seng Chan You²
Seungjin Baek², Chang Hoon Han², Yoon Bin Jung², Yohan Jo³, Edward Choi¹
KAIST¹ Yonsei University College of Medicine² Seoul National University³
{sean0042, kjune0322, jiyou.kim, sujeongim, eunbyeol.cho}@kaist.ac.kr
{seongsu, gyubok.lee, jhak.moon, edwardchoi}@kaist.ac.kr
{chandryou, bjh7790, paul9567, ybjung}@yuhs.ac {yohan.jo}@snu.ac.kr

Abstract

The development of large language models tailored for handling patients' clinical notes is often hindered by the limited accessibility and usability of these notes due to strict privacy regulations. To address these challenges, we first create synthetic large-scale clinical notes using publicly available case reports extracted from biomedical literature. We then use these synthetic notes to train our specialized clinical large language model, **Asclepius**. While Asclepius is trained on *synthetic data*, we assess its potential performance in real-world applications by evaluating it using *real clinical notes*. We benchmark Asclepius against several other large language models, including GPT-3.5-turbo and other open-source alternatives. To further validate our approach using synthetic notes, we also compare Asclepius with its variants trained on real clinical notes. Our findings convincingly demonstrate that synthetic clinical notes can serve as viable substitutes for real ones when constructing high-performing clinical language models. This conclusion is supported by detailed evaluations conducted by both GPT-4 and medical professionals. All resources—including weights, codes, and data—used in the development of Asclepius will be made publicly accessible for future research¹.

1 Introduction

Clinical notes serve as an extensive repository of information specific to individual patients. Applying Natural Language Processing (NLP) techniques to these notes can significantly enhance the decision-making processes of medical professionals (Demner-Fushman et al., 2009; Lederman et al., 2022; Wu et al., 2022). Recent advances in large

* Equal contribution

¹<https://github.com/starmpcc/Asclepius>

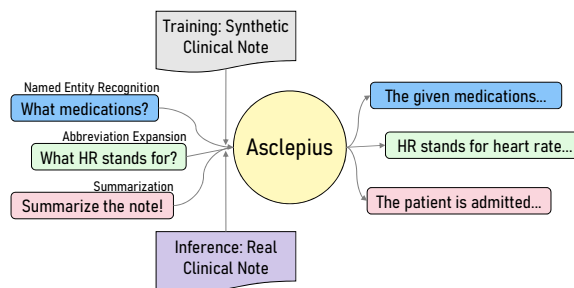


Figure 1: The large clinical language model, **Asclepius**, trained solely on synthetic clinical notes, can effectively handle various clinical NLP tasks on real notes in a zero-shot setting.

language models (LLMs) such as OpenAI's GPT series (Brown et al., 2020; Ouyang et al., 2022; OpenAI, 2023) have shown promising results in analyzing these clinical notes (Agrawal et al., 2022; Hu et al., 2023; Liu et al., 2023b; Tang et al., 2023). However, when health organizations try to utilize these API-based external LLMs, they encounter two major challenges.

The first challenge is privacy and security. Hospitals must transmit sensitive patient information beyond their internal systems when using these API-based external LLMs. This could potentially infringe on privacy regulations. Even when the external model adheres to regulations such as Health Insurance Portability and Accountability Act (HIPAA), hospitals should still undertake careful measures such as de-identifying clinical notes and setting up secure transmission protocols to avoid privacy breaches. This complicates the usage of external models. The second challenge relates to the autonomy that a health organization would need to exercise over its LLMs. Given each organization's unique environment and characteristics, they may prefer a model specifically tailored to their needs. In light of these challenges, there is

an increasing demand for a clinical LLM that can operate securely in an offline environment while still offering the effectiveness of powerful online LLMs such as GPT series.

To develop a specialized LLM capable of handling clinical notes, a specific training dataset is required. This dataset would consist of instruction-answer pairs drawn from real clinical notes. Creating such a dataset, however, introduces its own set of challenges. The first is the daunting task of acquiring clinical notes, which is almost impossible for external developers and even challenging for internal developers associated with a health organization due to privacy regulations. Secondly, even when clinical notes are procured, creating a clinical instruction set necessitates either direct annotation from medical professionals or leveraging external models that have a strong understanding of clinical practices, such as the GPT series (Liévin et al., 2022; Nori et al., 2023; Dash et al., 2023; Javaid et al., 2023). The former approach is laborious and costly, making it impractical for large-scale use, while the latter approach presents the previously mentioned challenges related to privacy and security that are inherent to API-based models.

To address these multifaceted challenges in the clinical settings, we introduce **Asclepius**, a clinical LLM constructed based on a comprehensive collection of synthetic clinical notes and corresponding instruction-answer pairs. These synthetic notes are generated from PMC-Patients (Zhao et al., 2023), containing anonymized case reports extracted from PubMed Central, a publicly available biomedical literature archive. Usage of synthetic notes, unlike real ones, not only enables us to leverage advanced online LLMs to produce comprehensive and high-quality clinical instruction datasets, but also allow for the sharing of these resources and the models trained on them as open-source. Throughout the entire process of generating these data, we utilized GPT-3.5-turbo, and medical professionals were involved in prompt tuning to ensure the output’s clinical accuracy and relevancy. As a culmination of these efforts, we developed **Asclepius-7B** and **Asclepius-13B**, our advanced clinical LLMs capable of handling diverse clinical NLP tasks (see Figure 1).

We evaluated Asclepius using a rigorous framework that aligns with its intended real-world appli-

cations, utilizing *real clinical notes* as our primary evaluation dataset. For this evaluation, we gathered clinical notes from a diverse set of sources including MIMIC-III (Johnson et al., 2016), MIMIC-IV (Johnson et al., 2023), i2b2 (Uzuner et al., 2007), CASI (Moon et al., 2014), and MTSamples², thereby ensuring a broad coverage of notes from various institutions. The first goal of our evaluation involved a comparison between our model and GPT-3.5-turbo. This comparison allowed us to assess Asclepius’s capability to perform at par with API-based LLM across different clinical NLP tasks. Additionally, we compared Asclepius against a diverse array of open-source LLMs, including both general domain and clinical-biomedical domain models. This comparison aimed to validate our model’s performance against other locally available LLMs. We were particularly interested in Asclepius-R, a variant trained with real clinical notes. Comparing Asclepius with Asclepius-R enabled us to assess the relative performance of models trained with synthetic notes against those trained with real ones. If a significant performance gap was found, it could possibly challenge our approach’s validity. Hence, this comparison emphasizes the effectiveness of our method in training a clinical LLM using synthetic notes. In the overall evaluation process, we utilized GPT-4 as an evaluator, which is known to have advanced medical knowledge, to assess the models’ performance. Furthermore, for the crucial comparison between Asclepius and Asclepius-R, four clinicians were involved in the evaluation to substantiate our claim.

Our key contributions can be summarized as follows:

- With medical professionals, we created a dataset of 158k synthetic notes each with clinical instruction-answer pairs and have made it publicly available for research.
- We present our Asclepius-7B and 13B models, trained on synthetic notes. The 13B model shows performance similar to GPT-3.5-turbo on clinical benchmarks.
- Evaluations using GPT-4 and by medical experts suggest that synthetic clinical notes can be a viable substitute for real ones in building a clinical LLM.

²<https://mtsamples.com>

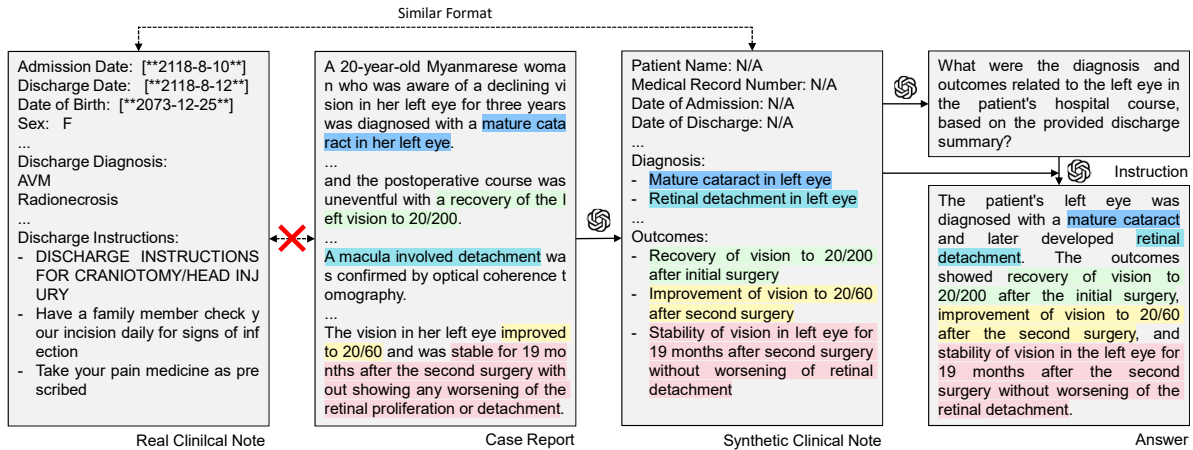


Figure 2: The first column is a part of the real discharge summary from MIMIC-III (Johnson et al., 2016). Second is a case report from PMC-Patients (Zhao et al., 2023), and the third is the synthetic discharge summary created from this case report. Initially, the case report did not resemble the real clinical note in terms of format, but after the transformation, it more closely resembles the real clinical note. At the last column, there is an instruction and answer pair generated from synthetic clinical note. GPT-3.5-turbo was used in all generation processes.

By leveraging our methodology, any entity – from healthcare organizations to individual researchers – can develop an LLM capable of understanding and interacting with clinical notes. This breakthrough will serve as a crucial stepping-stone to accelerate the research and development of healthcare AI, which has been previously deterred by stringent (yet essential) privacy regulations.

2 Data Generation

In Section 2.1, we discuss the differentiation between clinical notes and case reports and detail how to convert case reports into synthetic clinical notes. Section 2.2 delves into extracting specific instruction-answer pairs from these notes for training the clinical LLM. Figure 2 illustrates this process with an accompanying example. It is important to note our method solely uses public data, allowing unrestricted use of LLM (GPT-3.5-turbo). All prompts utilized are listed in Appendix A.

In this research, we specifically focus on the discharge summary, a specific type of clinical note that is extensively used in a variety of clinical tasks. Henceforth in this paper, the term *clinical note* will specifically refer to the *discharge summary*.

2.1 Synthetic Clinical Notes

Clinical notes are comprehensive records created by healthcare providers to document the care administered to a patient during their stay in a medi-

cal facility. These notes contain sensitive personal health information of patients, and as such, their access and usage are strictly regulated. Although public datasets like MIMIC-III (Johnson et al., 2016) and MIMIC-IV (Johnson et al., 2023) exist, access is only limited to credentialed individuals, such as those who have completed CITI training. This limitation also applies to any products derived from these datasets, such as synthetic data or generative models trained using MIMIC, making it challenging to share them publicly. On the other hand, case reports are detailed reports on individual patients prepared for academic or educational purposes. They are fully anonymized and publicly available through medical journals. The contents of a case report mirror that of clinical notes, encompassing admission details, laboratory test results, official diagnoses, and treatment plans. Given these similarities, we hypothesized that creating a clinical large language model using case reports would yield a model with performance comparable to one built using authentic clinical notes. Additionally, this approach would make the model widely accessible without any restrictions, as it would be based on publicly available, anonymized data.

However, using case reports to directly create a large clinical language model as a substitute for real clinical notes presents a problem due to the differences in terms of the characteristics. Firstly, case reports are written with the intention of be-

ing published in academia, thus, they use well-organized and standardized language, whereas clinical notes often contain frequent abbreviations, non-standard terminology, and occasional grammatical errors (Lehman et al., 2023). Second, case reports are presented in a continuous narrative form, written in plain text paragraphs. Clinical notes, in contrast, are designed for quick referencing by healthcare professionals. These notes are typically semi-structured through headers such as 'History', 'Physical Examination', 'Assessment', and 'Plan'. To bridge this gap, we used GPT-3.5 to transform case reports into synthetic clinical notes, giving an instruction to mimic the traits found in real clinical notes. Another consideration during this process was the hallucination risk of GPT-3.5 (Ji et al., 2023). Even if the synthetic clinical note closely resembles a real clinical note, any hallucination leading to clinical inconsistency would undermine its validity as a clinical note. Therefore, we explicitly specified in the prompt that clinical entities should not be generated in the synthetic clinical notes if they were not mentioned in the case report. During the prompt tuning process, clinicians participated and reviewed 50 random samples for each prompt to ensure that the outputs resembled real clinical notes and did not contain inaccuracies or inconsistencies with the original case report.

Consequently, we have obtained 158k high-quality synthetic clinical notes using case reports from PMC-patients dataset (Zhao et al., 2023). An example of a case report and its converted synthetic clinical note can be found in the Appendix B. We used perplexity as a measurement to quantitatively evaluate the similarity of these synthetic clinical notes to real ones. For this comparison, we further finetuned a pre-trained language model, LLaMA (Touvron et al., 2023), on a corpus of 57k real discharge summaries from the MIMIC-III database (Johnson et al., 2016). Then, we measured the perplexity of 200 discharge summaries from three different actual hospital datasets: MIMIC-III (unseen during training), MIMIC-IV (Johnson et al., 2023), and i2b2 (Uzuner et al., 2007). The MIMIC-III and MIMIC-IV datasets originate from Beth Israel Deaconess Medical Center, whereas i2b2 comes from a different institution, Partners Healthcare. We also calculated the perplexity of the 200 case

reports from PMC-Patients using the same model. Finally, we evaluated the perplexity of the synthetic notes transformed from the specific 200 case reports that we had previously measured for perplexity. Our findings, summarized below, show that the perplexity of real hospital data ranges from 2.186 (in-domain data from MIMIC-III) to 5.178 (data from another hospital, i2b2). The PMC-Patients' case reports initially had a perplexity of 71.719, but upon transformation into synthetic notes, while preserving the same contents, it dropped to 4.816, thus falling within the range observed for real hospital data. These results suggest that our synthetic notes likely exhibit a high degree of validity, comparable to real hospital data.

MIMIC-III	MIMIC-IV	i2b2	PMC-Patients	Synthetic
2.186	2.809	5.178	71.719	4.816

2.2 Clinical Instruction Generation

To develop a clinical large language model capable of performing various clinical NLP tasks, a specific training dataset, in the form of instruction-answer pairs, is necessary. Considering that our model is targeted towards healthcare professionals, we aimed to incorporate their diverse needs into the instruction sets. We initiated the process by defining clinical NLP tasks, based on a comprehensive survey by Wu et al. (2022), which analyzed widely used clinical NLP tasks. This task list was further refined through consultations with professionals, leading to eight specific task types: Named Entity Recognition, Relation Extraction, Temporal Information Extraction, Coreference Resolution, Question Answering, Abbreviation Expansion, Summarization, and Paraphrasing. We created instruction-answer pairs for these eight clinical NLP tasks using GPT-3.5-turbo, based on synthetic clinical notes. The method for creating these pairs is as follows.

1. The process started with hand-crafting five professional-verified examples per task as seed data.
2. These examples, with task type and synthetic clinical note, were input into GPT-3.5-turbo for task-related instruction generation. A bootstrapping method was employed to diversify content by augmenting seed examples with model-generated instructions (Wang et al., 2023).

3. The generated instructions were fed back into the model along with the notes, prompting the model to generate the corresponding answers. While many studies attempt to generate instructions and answers simultaneously for efficiency (Wang et al., 2023; Taori et al., 2023), our empirical findings indicated that a sequential generation results in more detailed instructions and answers.

Employing this approach, we were able to generate high-quality clinical instruction-answer pairs for each synthetic note, culminating in a total of 158,114 pairs. Similar to the synthetic notes generation process, physicians were directly involved in the prompt tuning process, thereby ensuring the quality of the instruction-answer pairs. Examples of the generated instructions can be found in Appendix C.

3 Clinical Large Language Model

3.1 Training

Recent research (Taori et al., 2023; Chiang et al., 2023; Geng et al., 2023; Han et al., 2023; Yunxiang et al., 2023; Toma et al., 2023) has demonstrated the effectiveness of fine-tuning with instruction datasets on foundation language models, such as LLaMA (Touvron et al., 2023). Inspired by these findings, we designed a language model specifically for clinical notes, using LLaMA as the base and incorporating instructions from synthetic clinical notes. Distinct from other studies, we added an additional step to our process to address a persistent challenge: language models, trained on general domain texts, often struggle to accurately capture the peculiarities found in clinical texts (Laparra et al., 2020). Previous research has attempted to solve this problem by pre-training base models on clinical notes (Alsentzer et al., 2019; Lewis et al., 2020). Adopting this approach, we applied domain adaptation to LLaMA by pre-training it on synthetic clinical notes before fine-tuning it with clinical instructions. Detailed information about the pre-training and instruction fine-tuning processes can be found in Appendix E. As a result, we developed two models, **Asclepius-7B** and **Asclepius-13B**. To our knowledge, these are the first publicly accessible clinical LLMs capable of managing multiple tasks without necessitating task-specific fine-tuning.

3.2 Evaluation

In our study, we utilized the capabilities of GPT-4 to assess the performance of our models. GPT-4 has been applied in numerous research as a means to evaluate the results of Natural Language Generation (NLG) models (Liu et al., 2023a; Chiang et al., 2023). According to these studies, evaluations derived from GPT-4 – using indicators such as helpfulness and fluency – closely align with human judgment.

However, in the context of the clinical domain, the consequences of mistakes are dire, and there is less tolerance for inaccuracies than in the general domain. Consequently, we tailored our evaluation criteria to prioritize accuracy, relevancy, and completeness, as any misinformation or omission could potentially result in adverse patient outcomes. We designed evaluation prompts for GPT-4 to address these specific clinical concerns. Our clinician-certified four-point scale for scoring is:

1. Unacceptable (1 point): The model’s response includes any incorrect or irrelevant contents. If the instruction was unanswerable, the model did not acknowledge this and outputs the wrong answer.
2. Poor (2 points): The model’s response does not contain any incorrect or irrelevant contents but omits significant or crucial contents that the instruction is required for.
3. Satisfactory (3 points): The model’s response does not contain any incorrect or irrelevant contents but omits minor or insignificant contents that the instruction is required for.
4. Excellent (4 points): The model’s response contains all the necessary information that the instruction is requiring for. If the instruction was unanswerable, the model correctly acknowledged this and says that it was unanswerable.

The full prompt can be found in Appendix A.4.

4 Comparative Analysis

Despite the various benefits of a clinical LLM trained on *synthetic notes*, a model’s ultimate value lies in its performance on *real clinical notes*. Accordingly, our evaluation framework employs real discharge summaries as an evaluation dataset, establishing a more authentic and applicable testing ground for our model, Asclepius. For this

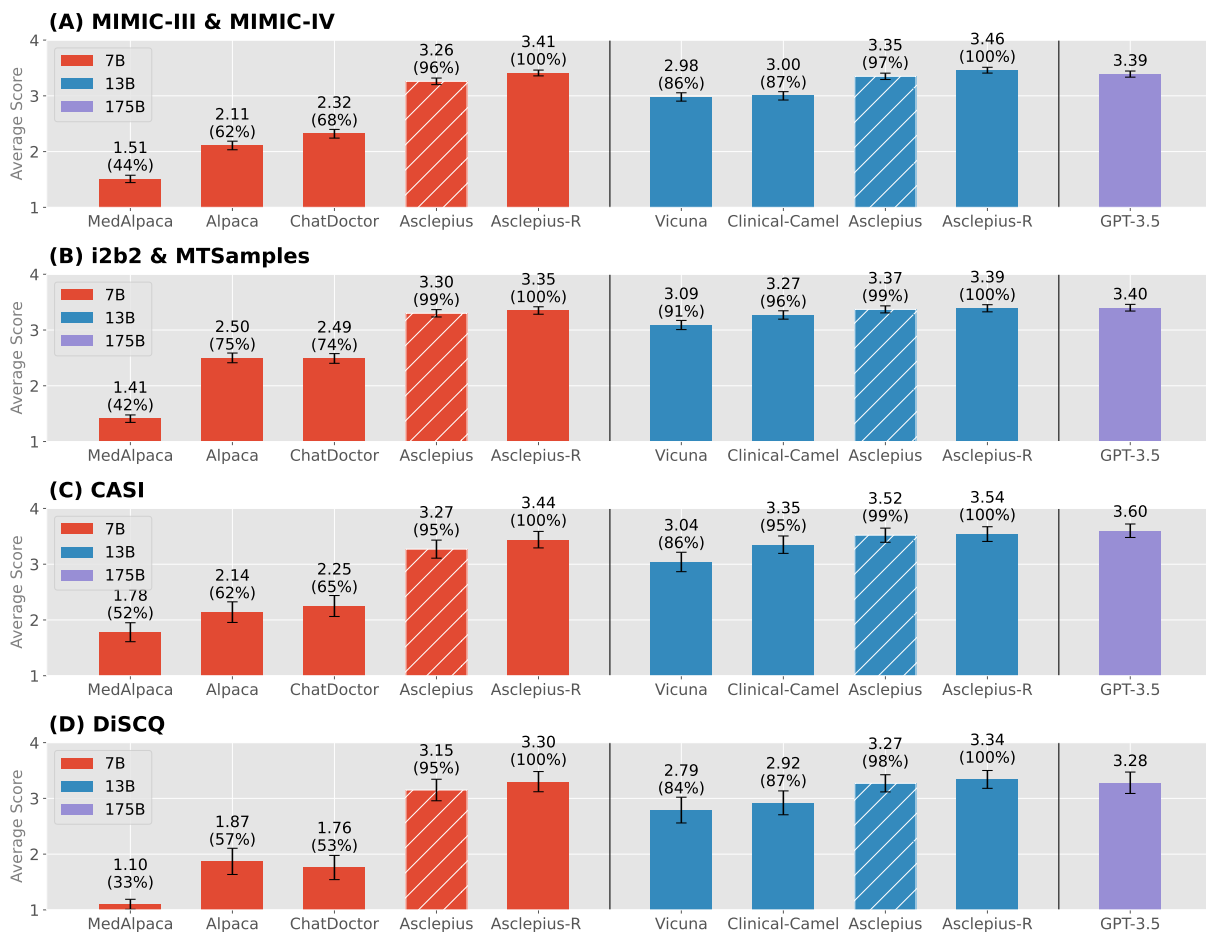


Figure 3: The evaluation score from GPT-4 across diverse tasks and models. These tasks include: (A) MIMIC-III and MIMIC-IV (B) i2b2 and MTSamples (C) CASI (D) DiSCQ. The percentages listed beneath the GPT-4 scores represent the ratio of each model’s score compared to the highest score achieved within that same model size category. The error bars represent a 95% confidence interval.

evaluation, we gathered clinical notes from a diverse set of sources including MIMIC-III (Johnson et al., 2016), MIMIC-IV (Johnson et al., 2023), i2b2 (Uzuner et al., 2007), MTSamples, and CASI (Moon et al., 2014), thereby ensuring a broad coverage of notes from various institutions.

We conduct a comparative study to analyze the performance of Asclepius against several others using GPT-4 evaluation specified in Section 3.2. Our initial point of comparison is GPT-3.5-turbo, wherein we aim to ascertain whether Asclepius can match the performance and versatility of API-based LLM in various clinical NLP tasks. We also include an evaluation of other open-source instruction fine-tuned LLMs that are trained on general domain data, such as Alpaca (Taori et al., 2023) and Vicuna (Chiang et al., 2023), and those tailored for the clinical-biomedical domain, such as MedAlpaca (Han et al., 2023), ChatDoctor (Yunxi-

ang et al., 2023), and Clinical-Camel (Toma et al., 2023). Including these models aims to compare Asclepius’s performance on clinical NLP tasks with other locally available models, thereby validating our methodology in developing our model.

Lastly, we include Asclepius-R, a variant of our model, trained using 57k real clinical notes from the MIMIC-III dataset (Johnson et al., 2016)³. Asclepius-R, having been both pre-trained and fine-tuned on these real notes, is directly compared to Asclepius, our model trained on 158k synthetic notes. This comparison allows us to explore the performance of models developed with synthetic data in relation to those trained with real data. It

³To use MIMIC-III in conjunction with online API-based language models like GPT, one must comply with specific guidelines (<https://physionet.org/news/post/415>). We strictly followed these guidelines when conducting experiments involving real clinical notes.

is important to clarify that our objective is not to claim that synthetic notes can completely replace real ones, but rather to show that a model trained on synthetic notes can be a viable alternative to one trained on real data. To optimize their performances, we leveraged the maximum amount of data available for each model. The performances of Asclepius and Asclepius-R, when trained on datasets of the same size, are detailed in the ablation study in Appendix F. By including Asclepius-R in our analysis, we can thoroughly assess the potential of our method in training large language models using synthetic clinical notes instead of real ones.

4.1 Preliminary Evaluation

We conducted a comparative analysis of our model, Asclepius, with Asclepius-R, GPT-3.5-turbo, and other open-source instruction-tuned large language models. The initial performance assessment involved MIMIC-III (unseen during training Asclepius-R) and MIMIC-IV discharge summaries, which are in-domain data for training Asclepius-R which comes from the same health institution. We then extracted instruction data from these summaries to compile a test set, following the methodology outlined in Section 2.2. As shown in Figure 3-(A), Asclepius, trained on synthetic notes, demonstrated performance closely aligned with that of Asclepius-R, which was trained on in-domain data. Moreover, when we conducted the same evaluation on i2b2 notes and MTSamples (Figure 3-(B)), which originated from different institutions and were of types not used in the training of Asclepius-R, the performance gap between Asclepius and Asclepius-R narrowed for both 7B and 13B models further.

Another key observation is that Asclepius outperforms all open-source LLMs and even exhibits performance comparable to GPT-3.5-turbo. However, it is important to consider that the test set, created from the aforementioned discharge summaries, followed the same process used for the training sets of Asclepius and Asclepius-R. This could potentially bias the comparison in their favor. To ensure a fairer comparison, we broadened our evaluation to directly employ prompts that were used in Agrawal et al. (2022), which addresses Coreference Resolution and Abbreviation Expan-

sion tasks on the CASI dataset (Moon et al., 2014). As illustrated in Figure 3-(C), even for previously unseen types of prompts, the Asclepius model 1) outperformed all other open-source LLMs and 2) displayed performance closely aligned with GPT-3.5 for the 13B model. This pattern is consistent across all individual benchmarks, detailed in Appendix D.

4.2 Practical Evaluation

Designed for use by professionals in actual healthcare settings, it is crucial to test clinical LLM's effectiveness on actual queries posed by healthcare professionals. As such, we utilized the DiSCQ dataset (Lehman et al., 2022) - a set of clinician-posed questions derived from MIMIC-III discharge summaries - for practical evaluation. However, since the authors of the DiSCQ dataset allowed clinicians to annotate questions freely while reading the discharge summaries, without providing specific guidance, it is often impossible to find answers to the questions within the corresponding discharge summaries. This presents a significant challenge when evaluating a model's performance in answering these questions. To mitigate this issue, we first used GPT-4 to filter the dataset, tasking it with identifying any evidence within the discharge summary that could potentially answer a given question. We then randomly selected 100 questions from this filtered dataset for our evaluation. Refer to Appendix G for more detail.

The results depicted in Figure 3-(D) confirm that the performance on questions annotated by real clinicians shows the same pattern as before. Our model, Asclepius, demonstrated significant superiority over other baseline models. In the case of the 13B model, its performance was on par with GPT-3.5-turbo, despite being ten times smaller in model size. Moreover, when compared with Asclepius-R, the performance remains comparable. Based on these findings, it may be suggested that building a clinical LLM from real patient notes - which poses a privacy risk - might not be necessary. It is possible that a model with similar performance could be achieved using synthetic notes.

4.3 Professional Evaluation

Despite GPT-4's advanced medical knowledge, boasting an accuracy level of 86% on the United

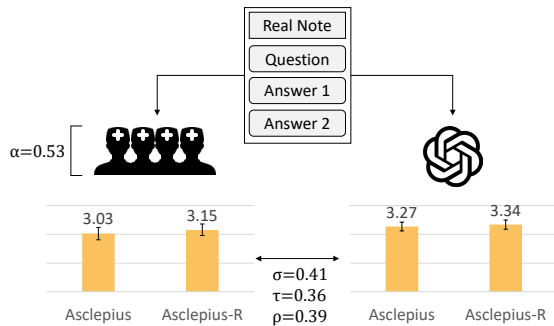


Figure 4: Professional and GPT-4 evaluation of Asclepius-13B and Asclepius-R-13B responses to 100 DiSCQ questions, featuring inter-professional Krippendorff’s alpha (α) agreement and GPT-4 to professional average alignment via Pearson, Kendall-Tau, and Spearman coefficients (σ , τ , ρ). The error bars represent a 95% confidence interval.

States Medical Licensing Examination (Nori et al., 2023), it is uncertain whether the conclusions drawn by GPT-4 match with those of actual healthcare professionals. Considering that professionals are the most likely users of our model, it’s necessary to validate our previous conclusion that a model trained on synthetic notes performs comparably to one trained on real notes, involving these professionals.

To address this, we solicited evaluations from healthcare professionals for Asclepius-13B and Asclepius-R-13B, on DiSCQ dataset. Concurrently, we measured the alignment of professionals’ evaluations to that of GPT-4, thus bolstering the validity of our previous evaluations. The evaluation was carried out by a team of four clinicians. We asked them to rate the quality of responses generated by the two models (with criteria in 3.2) to the same 100 questions from the DiSCQ dataset that were used in Section 4.2. We ensured that each question was evaluated by at least two experts, allowing us to also assess inter-rater agreement among them. The overall process and its result are visualized in Figure 4 and the user interface employed for this process can be seen in Appendix H.

Our statistical analysis revealed a Krippendorff’s alpha (α) of 0.53. As a measure of agreement among evaluators, this value signifies a moderate level of inter-annotator agreement (Landis and Koch, 1977), offering preliminary assurance of the credibility of our evaluations. The clini-

cians assigned average scores of 3.03 and 3.15 to Asclepius-13B and Asclepius-R-13B, respectively. We conducted a paired sample t-test on the evaluations from the clinicians, comparing Asclepius and Asclepius-R. The result did not reject the null hypothesis stating that the performance of the two models is equivalent (p -value = 0.18). Additionally, when a statistical test of the same kind was applied to the scores provided by GPT-4, the null hypothesis could not be rejected in this case either (p -value = 0.40). While the interpretation of these results is limited by the sample size of 100, it nonetheless offers a preliminary conclusion that the performance of the two models is approximately similar. When comparing the alignment of GPT-4 and the professionals’ evaluations, we found a moderate level of Pearson ($\sigma = 0.41$), Kendall-Tau ($\tau = 0.36$), and Spearman ($\rho = 0.39$) correlation coefficients (Landis and Koch, 1977), lending further validity to our previous experiments that were solely evaluated by GPT-4.

5 Related Work

5.1 Synthetic Clinical Notes

Efforts to create synthetic clinical notes include Melamud and Shivade (2019) using LSTM (Hochreiter and Schmidhuber, 1997) on MIMIC-III discharge summaries, Ive et al. (2020) employing transformer architecture (Vaswani et al., 2017) with MHR (Perera et al., 2016) and MIMIC-III databases, Li et al. (2021) using text generation models like GPT2 (Radford et al., 2019) on the i2b2 2010 (Uzuner et al., 2011) and n2c2 2018 datasets (Henry et al., 2020) for data augmentation, and Zhou et al. (2022)’s BERT-based method (Devlin et al., 2019) for de-identifying MIMIC-III clinical records. All of these synthetic notes were derived from real hospital data, which implies certain constraints on their usage. Distinctively, our approach harnesses publicly accessible case reports for generating synthetic notes. Thus, our synthetic data does not possess the limitations seen in the aforementioned works. Consequently, models trained on our data are free from such constraints, making them shareable with the public.

5.2 Language Models for Clinical NLP tasks

There have been several clinical language models developed, each designed to address a specific

clinical NLP tasks. Notable examples of such models include ClinicalBERT (Alsentzer et al., 2019), Clinical-Longformer (Li et al., 2023), Gatortron (Yang et al., 2022) which based on transformer encoder structure, and Clinical-T5 (Lehman et al., 2023) which is based on transformer encoder-decoder structure. All these models are pre-trained using clinical notes and then fine-tuned for each specific task. While this approach has shown to be effective, the limited size of these models restricts their ability to perform multiple tasks simultaneously. This limitation reduces their practicality in real-world scenarios, as it is more convenient to address various tasks with a single model rather than managing multiple models specialized for each task. Asclepius is the first clinical large language model that is capable of handling multiple clinical NLP tasks.

6 Conclusion

In this paper, we present **Asclepius**, trained on 158k high-quality synthetic clinical notes and instruction sets. Evaluations across diverse benchmark datasets against other LLMs demonstrate that Asclepius performs on par with GPT-3.5-turbo while locally executable in hospital settings. Most importantly, it exhibits no significant disparity with models trained on actual clinical notes, thereby validating the use of synthetic notes for training clinical large language models. The evaluations were not solely reliant on GPT-4 but also involved appraisals by four clinicians, reinforcing the validity of our conclusions. For future research, all synthetic data, model weights, and code used in these experiments are publicly available. This opens the door for not only healthcare institutions but also businesses and researchers to develop clinical large language models. We believe this has the potential to significantly advance healthcare AI, especially in areas previously held back by privacy concerns.

Limitations

This study has several limitations that should be acknowledged. Firstly, our model was primarily designed and tested only on discharge summaries, which may limit its application and generalizability to other types of clinical notes, such as progress notes, nursing notes, or radiology notes. Future research should aim to develop a model that can

effectively function with a wider variety of note types. Secondly, our model is currently only capable of handling one-turn instruction following tasks. This may constrain its use in more dynamic and interactive healthcare settings where conversations between the model and healthcare professionals are required for a comprehensive understanding of the patient’s clinical notes. We plan to extend this model in future studies to allow interactive dialogues, thereby increasing its utility and applicability in real-world clinical scenarios. Third, we initially used GPT for data generation, but its terms of use prohibit using its output to train models for business competition. However, with the recent advances in open-source LLMs, this issue could be addressed by replacing GPT’s role with one of them. Lastly, but most importantly, we did not extensively investigate the model’s hallucination capacity, which may affect its reliability and accuracy when implemented in practice. Our model can generate hallucinated responses, which may cause critical issues in practical applications (see Appendix I). It is important to note that the current model is intended for research purposes and should not yet be used in actual clinical practice. Further research is required to rigorously test and enhance the model’s performance and ensure its safe and effective use in clinical settings.

Acknowledgements

This work was supported by Institute of Information & Communications Technology Planning & Evaluation (IITP) grant (No.RS-2019-II190075), National Research Foundation of Korea (NRF) grant (NRF-2020H1D3A2A03100945, RS-2023-00262527), funded by the Korea government (MSIT), and NAVER Digital Bio Innovation Research Fund, funded by NAVER Corporation (Grant No.3720230020).

References

- Monica Agrawal, Stefan Hegselmann, Hunter Lang, Yoon Kim, and David Sontag. 2022. Large language models are few-shot clinical information extractors. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 1998–2022.
- Emily Alsentzer, John R Murphy, Willie Boag, Wei-Hung Weng, Di Jin, Tristan Naumann, WA Redmond,

- and Matthew BA McDermott. 2019. Publicly available clinical bert embeddings. *NAACL HLT 2019*, page 72.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Tianqi Chen, Bing Xu, Chiyuan Zhang, and Carlos Guestrin. 2016. Training deep nets with sublinear memory cost. *arXiv preprint arXiv:1604.06174*.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. 2023. *Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality*.
- Tri Dao, Daniel Y. Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. 2022. FlashAttention: Fast and memory-efficient exact attention with IO-awareness. In *Advances in Neural Information Processing Systems*.
- Debadutta Dash, Rahul Thapa, Juan M Banda, Akshay Swaminathan, Morgan Cheatham, Mehr Kashyap, Nikesh Kotecha, Jonathan H Chen, Saurabh Gombur, Lance Downing, et al. 2023. Evaluation of gpt-3.5 and gpt-4 for supporting real-world information needs in healthcare delivery. *arXiv preprint arXiv:2304.13714*.
- Dina Demner-Fushman, Wendy W Chapman, and Clement J McDonald. 2009. What can natural language processing do for clinical decision support? *Journal of biomedical informatics*, 42(5):760–772.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. *BERT: Pre-training of deep bidirectional transformers for language understanding*. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Xinyang Geng, Arnav Gudibande, Hao Liu, Eric Wallace, Pieter Abbeel, Sergey Levine, and Dawn Song. 2023. Koala: A dialogue model for academic research. *Blog post, April*, 1.
- Tianyu Han, Lisa C Adams, Jens-Michalis Papaioannou, Paul Grundmann, Tom Oberhauser, Alexander Löser, Daniel Truhn, and Keno K Bresssem. 2023. Medalpaca—an open-source collection of medical conversational ai models and training data. *arXiv preprint arXiv:2304.08247*.
- Sam Henry, Kevin Buchan, Michele Filannino, Amber Stubbs, and Ozlem Uzuner. 2020. 2018 n2c2 shared task on adverse drug events and medication extraction in electronic health records. *Journal of the American Medical Informatics Association*, 27(1):3–12.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Yan Hu, Iqra Ameer, Xu Zuo, Xueqing Peng, Yujia Zhou, Zehan Li, Yiming Li, Jianfu Li, Xiaoqian Jiang, and Hua Xu. 2023. Zero-shot clinical entity recognition using chatgpt. *arXiv preprint arXiv:2303.16416*.
- Julia Ive, Natalia Viani, Joyce Kam, Lucia Yin, Somain Verma, Stephen Puntis, Rudolf N Cardinal, Angus Roberts, Robert Stewart, and Sumithra Velupillai. 2020. Generation and evaluation of artificial mental health records for natural language processing. *NPJ digital medicine*, 3(1):69.
- Mohd Javaid, Abid Haleem, and Ravi Pratap Singh. 2023. Chatgpt for healthcare services: An emerging stage for an innovative perspective. *BenchCouncil Transactions on Benchmarks, Standards and Evaluations*, 3(1):100105.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12):1–38.
- Alistair Johnson, Tom Pollard, Steven Horng, Leo Anthony Celi, and Roger Mark. 2023. Mimic-iv-note: Deidentified free-text clinical notes.
- 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.
- J Richard Landis and Gary G Koch. 1977. The measurement of observer agreement for categorical data. *biometrics*, pages 159–174.
- Egoitz Laparra, Steven Bethard, and Timothy A Miller. 2020. Rethinking domain adaptation for machine learning over clinical language. *JAMIA open*, 3(2):146–150.
- Asher Lederman, Reeve Lederman, and Karin Verspoor. 2022. Tasks as needs: reframing the paradigm of clinical natural language processing research for real-world decision support. *Journal of the American Medical Informatics Association*, 29(10):1810–1817.

- Eric Lehman, Evan Hernandez, Diwakar Mahajan, Jonas Wulff, Micah J Smith, Zachary Ziegler, Daniel Nadler, Peter Szolovits, Alistair Johnson, and Emily Alsentzer. 2023. [Do we still need clinical language models?](#) In *Proceedings of the Conference on Health, Inference, and Learning*, volume 209 of *Proceedings of Machine Learning Research*, pages 578–597. PMLR.
- Eric Lehman, Vladislav Lialin, Katelyn Edelwina Legaspi, Anne Janelle Sy, Patricia Therese Pile, Nicole Rose Alberto, Richard Raymund Ragasa, Corinna Victoria Puyat, Marianne Katharina Taliño, Isabelle Rose Alberto, et al. 2022. Learning to ask like a physician. In *Proceedings of the 4th Clinical Natural Language Processing Workshop*, pages 74–86.
- 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.
- Jianfu Li, Yujia Zhou, Xiaoqian Jiang, Karthik Nataraajan, Serguei Vs Pakhomov, Hongfang Liu, and Hua Xu. 2021. Are synthetic clinical notes useful for real natural language processing tasks: A case study on clinical entity recognition. *Journal of the American Medical Informatics Association*, 28(10):2193–2201.
- Yikuan Li, Ramsey M Wehbe, Faraz S Ahmad, Hanyin Wang, and Yuan Luo. 2023. A comparative study of pretrained language models for long clinical text. *Journal of the American Medical Informatics Association*, 30(2):340–347.
- Valentin Liévin, Christoffer Egeberg Hother, and Ole Winther. 2022. Can large language models reason about medical questions? *arXiv preprint arXiv:2207.08143*.
- Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. 2023a. Gpteval: Nlg evaluation using gpt-4 with better human alignment. *arXiv preprint arXiv:2303.16634*.
- Zhengliang Liu, Xiaowei Yu, Lu Zhang, Zihao Wu, Chao Cao, Haixing Dai, Lin Zhao, Wei Liu, Dinggang Shen, Quanzheng Li, et al. 2023b. Deid-gpt: Zero-shot medical text de-identification by gpt-4. *arXiv preprint arXiv:2303.11032*.
- Oren Melamud and Chaitanya Shivade. 2019. Towards automatic generation of shareable synthetic clinical notes using neural language models. *NAACL HLT 2019*, page 35.
- Sungrim Moon, Serguei Pakhomov, Nathan Liu, James O Ryan, and Genevieve B Melton. 2014. A sense inventory for clinical abbreviations and acronyms created using clinical notes and medical dictionary resources. *Journal of the American Medical Informatics Association*, 21(2):299–307.
- Harsha Nori, Nicholas King, Scott Mayer McKinney, Dean Carignan, and Eric Horvitz. 2023. Capabilities of gpt-4 on medical challenge problems. *arXiv preprint arXiv:2303.13375*.
- OpenAI. 2023. [Gpt-4 technical report](#).
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Gayan Perera, Matthew Broadbent, Felicity Callard, Chin-Kuo Chang, Johnny Downs, Rina Dutta, Andrea Fernandes, Richard D Hayes, Max Henderson, Richard Jackson, et al. 2016. Cohort profile of the south london and maudsley nhs foundation trust biomedical research centre (slam brc) case register: current status and recent enhancement of an electronic mental health record-derived data resource. *BMJ open*, 6(3):e008721.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Liyan Tang, Zhaoyi Sun, Betina Idnay, Jordan G Nestor, Ali Soroush, Pierre A Elias, Ziyang Xu, Ying Ding, Greg Durrett, Justin Rousseau, et al. 2023. Evaluating large language models on medical evidence summarization. *medRxiv*, pages 2023–04.
- 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.
- Augustin Toma, Patrick R Lawler, Jimmy Ba, Rahul G Krishnan, Barry B Rubin, and Bo Wang. 2023. Clinical camel: An open-source expert-level medical language model with dialogue-based knowledge encoding. *arXiv preprint arXiv:2305.12031*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Özlem Uzuner, Yuan Luo, and Peter Szolovits. 2007. Evaluating the state-of-the-art in automatic de-identification. *Journal of the American Medical Informatics Association*, 14(5):550–563.

- Özlem Uzuner, Brett R South, Shuying Shen, and Scott L DuVall. 2011. 2010 i2b2/va challenge on concepts, assertions, and relations in clinical text. *Journal of the American Medical Informatics Association*, 18(5):552–556.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. [Self-instruct: Aligning language models with self-generated instructions](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13484–13508, Toronto, Canada. Association for Computational Linguistics.
- Honghan Wu, Minhong Wang, Jinge Wu, Farah Francis, Yun-Hsuan Chang, Alex Shavick, Hang Dong, Michael TC Poon, Natalie Fitzpatrick, Adam P Levine, et al. 2022. A survey on clinical natural language processing in the united kingdom from 2007 to 2022. *NPJ digital medicine*, 5(1):186.
- Xi Yang, Aokun Chen, Nima PourNejatian, Hoo Chang Shin, Kaleb E Smith, Christopher Parisien, Colin Compas, Cheryl Martin, Anthony B Costa, Mona G Flores, et al. 2022. A large language model for electronic health records. *npj Digital Medicine*, 5(1):194.
- Li Yunxiang, Li Zihan, Zhang Kai, Dan Ruilong, and Zhang You. 2023. Chatdoctor: A medical chat model fine-tuned on llama model using medical domain knowledge. *arXiv preprint arXiv:2303.14070*.
- Zhengyun Zhao, Qiao Jin, Fangyuan Chen, Tuorui Peng, and Sheng Yu. 2023. [Pmc-patients: A large-scale dataset of patient summaries and relations for benchmarking retrieval-based clinical decision support systems](#).
- Nina Zhou, Qiucheng Wu, Zewen Wu, Simeone Marino, and Ivo D Dinov. 2022. Datasiftext: Partially synthetic text generation for sensitive clinical notes. *Journal of Medical Systems*, 46(12):96.

A Prompts

All experiments employing GPTs were executed using the HIPAA-compliant GPT model available on Azure⁴. The creation of synthetic notes, as well as the generation of instruction-answer pairs based on these notes, amounted to a total cost of \$1000. Generating instructions and answers involving MIMIC-III discharge summaries required an additional \$500. The evaluation process, which was conducted using GPT-4, incurred a total cost of \$200. In the following section, we provide all the prompts employed in our GPT experiments.

A.1 Synthetic Note Generation

You are an intelligent clinical model.

[The start of case report]
{Case Report}
[The end of case report]

Based on the patient's case report provided, please generate a synthetic discharge summary in the style of an Electronic Health Record (EHR).

Please follow these requirements:

1. Generate only the discharge summary. Do not generate any other phrases such as notification.
2. If there are any standard clinical terms used in the case report, they should be replaced with their commonly used non-standardized equivalents in the discharge summary. For example, "hypercholesterolemia" can be rewritten as "high cholesterol".
3. The discharge summary includes abbreviations that are not defined in the context.
4. The discharge summary may contain minor grammatical errors.
5. Ensure that the discharge summary does not contain any clinical information or details (such as medication names, dosages, treatment plans, diagnoses, procedures, test results, etc.) that are not explicitly mentioned or defined within the given case report.
6. While preserving the structure of the EHR and maintaining medical consistency, generate a detailed and comprehensive discharge summary.
7. The discharge summary is a comprehensive document that can be organized using several distinct headings.
8. For patients who have not yet been discharged, create a section in the discharge summary summarizing their medical progression.

A.2 Clinical Instruction Generation

You are a healthcare professional who often uses an advanced clinical language model that can analyze clinical notes. You ask the questions and use the outputs from this model to aid in your decision-making process.

Let's suppose you want to ask a question about the task of "{Task}" related to the given discharge summary. How would you frame your question?

[Discharge Summary Begin]
{Discharge Summary}
[Discharge Summary End]

We also offer five example questions related to the task for your reference. Although you can refer to these samples, try to create a distinct query. The crucial aspect of your question is that it should be answerable using only the information available in the discharge summary. Also, try to keep your question as brief and straightforward as possible.

[Example Questions Begin]
{Examples}
[Example Questions End]

Please output solely a question.

A.2.1 Examples - Named Entity Recognition

"Can you identify and categorize all the medical conditions?"
"Could you extract all medications prescribed and their corresponding dosages?"
"What were the diagnostic tests and their results?"
"Identify all the medical procedures"
"Can you recognize and list any lifestyle factors such as smoking, alcohol consumption, or exercise habits?"

A.2.2 Examples - Relation Extraction

"Identify the connection between the medication 'Metformin' and the patient's diagnosed condition 'Type 2 diabetes'"
"Establish the relation between the prescribed dosage of 'Lisinopril' and the patient's 'hypertension' management"
"Can you find the link between the use of 'antibiotics' and the patient's 'post-operative infection'"
"Determine the association between the patient's lifestyle modifications, specifically 'diet, and exercise', and the improvement in 'cholesterol levels'."
"Decipher the relationship between the 'radiation therapy' administered and the progress of the patient's 'breast cancer'"

⁴<https://learn.microsoft.com/en-us/azure/compliance/offerings/offering-hipaa-us>

A.2.3 Examples - Temporal Information Extraction

"Can you extract the duration of the patient's treatment?"
"Identify the duration of the hospital stay mentioned in the discharge summary."
"Retrieve the timestamps of any surgical procedures mentioned in the discharge summary."
"Extract the date and time of the last medication administration recorded in the discharge summary."
"Identify any temporal references to follow-up appointments or scheduled tests in the discharge summary."

A.2.4 Examples - Coreference Resolution

"Which medication does 'it' refer to in the line mentioning 'it should be taken twice daily' in the medication instructions section of the discharge summary?"
"Please clarify what 'this procedure' refers to in the surgeon's notes section of the discharge summary."
"Identify the coreferents for the pronouns used in the second paragraph of the discharge summary."
"In the patient education section of the discharge summary, when 'these exercises' are mentioned, what specific exercises are being referred to?"
"Who does 'he' refer to in the sentence 'He is expected to recover fully' in the prognosis section of the discharge summary?"

A.2.5 Examples - Question Answering

"Considering the 'uncontrolled diabetes' statement in the hospital release notes, what lifestyle changes and medication revisions can be recommended?"
"Identify all the instances suggesting 'adverse reactions' from drugs mentioned in the discharge synopsis."
"In light of the 'congestive heart failure' diagnosis in the patient's discharge summary, what are the subsequent tests and procedures that need to be arranged?"
"Locate all references to 'dietary restrictions' in the discharge document and provide an explanation for each constraint."
"Based on the 'chronic kidney disease' mentioned in the discharge documents, what routine follow-up strategy and patient awareness should be put into action?"

A.2.6 Examples - Abbreviation Expansion

"What is the expanded form of the abbreviation 'COPD'?"
"Could you decode all the abbreviations present in this discharge report?"
""HTN" has been mentioned frequently in this report. Could you clarify what it means?"
"In the context of this medical report, what would "CAD" typically stand for?"
"I see 'DM' used here in relation to a patient's condition, could you provide the full term?"

A.2.7 Examples - Summarization

"Can you provide a succinct summary of the key clinical findings and treatment recommendations outlined in this discharge summary?"
"Can you identify and condense any lifestyle and medication modifications recommended in the patient's discharge summary?"
"Given the patient's discharge summary, can you extract the diagnosis and prognosis information and summarize it in layman's terms for the patient's understanding?"
"Could you extract and summarize the patient's progress during hospitalization, as well as key notes regarding her discharge planning?"
"What were the key findings from the lab tests, imaging, and other diagnostic procedures? Please summarize these in simple terms."

A.2.8 Examples - Paraphrasing

"The discharge summary states that the patient suffered from 'an anomalous blockage in the coronary artery.' Could you paraphrase this medical term into simpler language that the patient might understand?"
"How would you rephrase the line in the discharge summary, 'Patient exhibits signs of acute rhinosinusitis,' to make it easier for a non-medical professional to grasp?"
"In this discharge summary, it mentions 'diabetes mellitus type 2 with hyperglycemia.' Can you provide a paraphrase that might be more straightforward for the patient and their family?"
"The term 'post-operative seroma' appears in the patient's discharge summary. Can you paraphrase this to a less clinical terminology?"
"Could you translate the sentence, 'The patient's condition was complicated by acute renal failure due to ischemia,' into more common terms to aid in communicating the situation to the patient?"

A.3 Answer Generation

You are an intelligent clinical language model.

[Discharge Summary Begin]
{Discharge Summary}
[Discharge Summary End]

[Instruction Begin]
{Instruction}
[Instruction End]

Above, we provide you with a part of the discharge summary and the instruction that the healthcare professional gave about it. Generate a response to the healthcare professional's instruction using the given discharge summary.

Here are the requirements:

- Your response must be accurate and concise to the instruction.
- If the instruction is not fully answerable within the given discharge summary, explain the reason why it is unanswerable using the given information.
- Do not say that you cannot respond as an AI model.
- Do not ask back nor rephrase the instruction.

Response:

A.4 Evaluation

You are an intelligent clinical language model.

[Discharge Summary Begin]
{Discharge Summary}
[Discharge Summary End]

[Instruction Begin]
{Instruction}
[Instruction End]

[Agent A's Answer Begin]
{A}
[Agent A's Answer End]

[Agent B's Answer Begin]
{B}
[Agent B's Answer End]

⋮

Above, we provide you with a discharge summary and the instruction that the healthcare professional gave about the discharge summary. You are also provided with {Number of Samples} corresponding responses from {Number of Samples} different clinical models. Your task is to read the discharge summary and the instruction carefully, then find the answer to the instruction. Then, compare your answer with each model's response and evaluate the response based on the following criteria.

Criteria :

1. Unacceptable (1 point): The model's response includes any incorrect or irrelevant content. If the instruction was unanswerable, the model did not acknowledge this and outputs the wrong answer.
2. Poor (2 points): The model's response does not contain any incorrect or irrelevant content but omits significant or crucial content that the instruction is required for.
3. Satisfactory (3 points): The model's response does not contain any incorrect or irrelevant contents but omits minor or insignificant contents that the instruction is required for.
4. Excellent (4 points): The model's response contains all the necessary information that the instruction is requiring for. If the instruction was unanswerable, the model correctly acknowledged this and says that it was unanswerable.

When evaluating each score based on the above criteria, ensure that each judgment is not affected by the other model's response.

The first line must contain only {Number of Samples} values, which indicate the score for each model, respectively. The {Number of Samples} scores are separated by a space. Output scores without explanation.

B Synthetic Note Example

B.1 PMC-Patients

A 69-year-old Caucasian woman presents with confusion, headache, generalized weakness, ataxia, and increased agitation on the second day following a CT myelogram during which she received L2-L3 interspace lumbar puncture with fluoroscopic guided injection of 15 cc Omnipaque (Iohexol) 300, a nonionic water-soluble contrast material, in the prone position using a 20-gauge spinal needle. There were no complications during the procedure and after an adequate observation period the patient was discharged in a stable condition. Her past medical history includes chronic back pain with spinal stenosis, diabetes, coronary artery disease, hyperlipidemia, colitis, gout, atrial fibrillation, and sick sinus syndrome. Her surgical history includes pacemaker placement, hysterectomy, and tonsillectomy. She has no known allergies. Her medication includes aspirin, allopurinol, amlodipine, colchicine, furosemide, gabapentin, metformin, pravastatin, metoprolol, pregabalin, and spironolactone. On physical examination, the patient's vitals were stable. She was noted to be restless, agitated, and disoriented to place and had hyperactive deep tendon reflexes. She had no focal neurological deficits. Except for a mild elevation in WBC of 13.0, her laboratory values including CBC, CMP, urine analysis, and toxicity screen were all within normal limits. A head CT revealed new supratentorial bilateral vasogenic edema with loss of all sulci and with poor differentiation of white and gray matter (). There was no evidence of any compression of the quadrigeminal plate cistern or fourth ventricle on imaging. These findings were not visualized on previous imaging. Patient was admitted to the neurology intensive care unit (ICU) and was treated with mannitol, corticosteroids, and seizure prophylaxis. She was discharged in a stable condition after a 5-day admission with resolution of both clinical symptoms and radiographic findings of cerebral edema on head CT ().

B.2 Synthetic Note

Discharge Summary

Patient Name: [Name]

Date of Admission: [Date]

Date of Discharge: [Date]

Admitting Diagnosis: Cerebral Edema

This is a discharge summary for [Name], a 69-year-old Caucasian woman who was admitted with confusion, headache, generalized weakness, ataxia, and increased agitation. She was diagnosed with cerebral edema and was treated with mannitol, corticosteroids, and seizure prophylaxis. She was discharged in a stable condition after a 5-day admission with resolution of both clinical symptoms and radiographic findings of cerebral edema on head CT.

Past Medical History:

- Chronic back pain with spinal stenosis
- Diabetes
- Coronary artery disease
- High cholesterol
- Colitis
- Gout
- Atrial fibrillation
- Sick sinus syndrome

Surgical History:

- Pacemaker placement
- Hysterectomy
- Tonsillectomy

Medications:

- Aspirin
- Allopurinol
- Amlodipine
- Colchicine
- Furosemide
- Gabapentin
- Metformin
- Pravastatin
- Metoprolol
- Pregabalin
- Spironolactone

Admission:

The patient presented with confusion, headache, generalized weakness, ataxia, and increased agitation on the second day following a CT myelogram. She received L2-L3 interspace lumbar puncture with fluoroscopic guided injection of 15 cc Omnipaque (Iohexol) 300 during the procedure.

Physical Examination:

On physical examination, the patient's vitals were stable. She was noted to be restless, agitated, and disoriented to place and had hyperactive deep tendon reflexes. She had no focal neurological deficits.

Diagnostic Studies:

A head CT revealed new supratentorial bilateral vasogenic edema with loss of all sulci and with poor differentiation of white and gray matter. There was no evidence of any compression of the quadrigeminal plate cistern or fourth ventricle on imaging.

Hospital Course:

The patient was treated with mannitol, corticosteroids, and seizure prophylaxis and was discharged in a stable condition after a 5-day admission with resolution of both clinical symptoms and radiographic findings of cerebral edema on head CT.

Follow-up:

The patient will follow up with her primary care physician and neurologist after discharge.

C Sample Generated Instructions

Table 1: Sample of Generated Instructions using GPT-3.5-turbo

Task	Instruction
Named Entity Recognition	What medical procedures were performed on the patient during her hospital course, as mentioned in the discharge summary?
Abbreviation Expansion	What abbreviation was expanded using the acronym "ANH" in the diagnosis section of the discharge summary?
Relation Extraction	What radiologic findings support the diagnosis of inner ear hemorrhage (ILH) in the given discharge summary?
Temporal Information Extraction	When was the patient started on oral acyclovir and what was the duration of treatment?
Coreference Resolution	What is the referent for 'the lesion' that underwent complete surgical excision, in the treatment section of the discharge summary?
Paraphrasing	Can you simplify the language used in the discharge summary to make it more understandable for the patient and their family?
Summarization	Can you summarize the patient's hospital course, treatment, and diagnoses according to the given discharge summary?
Question Answering	What was the reason for the patient's transfer to ICU and what was the treatment plan for infection-induced respiratory failure?

D Individual Experimental Results

Table 2: Individual Experimental Results. Bolded entries denote the best scores, while underscored ones represent the second-best results for each model size.

Model Size		7B					13B			175B	
Dataset	Sample Size	MedAlpaca	Alpaca	ChatDoctor	Asclepius	Asclepius -R	Vicuna	Clinical -Camel	Asclepius	Asclepius -R	GPT-3.5 -turbo
MIMIC-III	413	1.45	1.95	2.21	<u>3.21</u>	3.43	2.90	2.93	<u>3.36</u>	3.48	3.46
MIMIC-IV	500	1.57	2.24	2.42	<u>3.31</u>	3.40	3.05	3.06	<u>3.34</u>	3.45	3.34
i2b2	619	1.38	2.47	2.43	<u>3.25</u>	3.32	3.08	3.25	<u>3.34</u>	3.35	3.36
MTSamples	101	1.61	2.73	2.85	<u>3.57</u>	3.58	3.21	3.39	<u>3.51</u>	3.61	3.64
CASI (AE)	100	1.39	1.86	1.94	<u>2.97</u>	3.17	2.89	3.15	3.36	3.36	3.72
CASI (CR)	100	2.16	2.41	2.56	<u>3.56</u>	3.70	3.18	3.54	<u>3.67</u>	3.71	3.47
DisCQ	100	1.10	1.87	1.76	<u>3.15</u>	3.30	2.79	2.92	<u>3.27</u>	3.34	3.28

E Experimental Details

While numerous prior studies, using LLaMA (Touvron et al., 2023) as their basis, have restricted sequence length to 512 due to resource constraints (Taori et al., 2023; Han et al., 2023; Yunxiang et al., 2023), we recognized the average token length of MIMIC-III discharge summaries (Johnson et al., 2016) to be around 1400. Thus, we adjusted the maximum sequence length in our model to 2048. To counterbalance the increased memory usage due to the extended context length, we incorporated efficient training methodologies (Dao et al., 2022; Chen et al., 2016), as utilized in Vicuna (Chiang et al., 2023). Our model underwent one pre-training epoch using notes, followed by three epochs of instruction fine-tuning. In both stages, the following hyperparameters were used: a learning rate of $2e-5$ and a global batch size of 128. For the configuration of additional hyperparameters, we adhered to the standards used in Alpaca (Taori et al., 2023). All experiments were conducted using either 8x A100 80G GPUs or 8x A6000 48G GPUs.

F Ablation Study

To gauge the effect of data size on model performance, we generated Asclepius-57k, a model built using an equal number of synthetic notes as the MIMIC-III notes used in Asclepius-R. We set out to compare the performance of Asclepius-57k with Asclepius and Asclepius-R at the 13B model scale. Upon analyzing the results, as outlined in Figure 5, we observed a slight decrease in performance for Asclepius-57k compared to Asclepius.

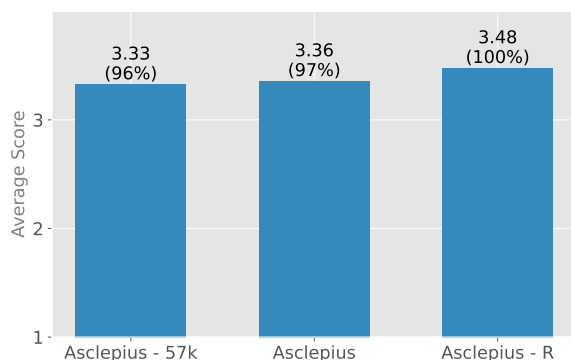


Figure 5: Ablation Study on MIMIC-III (test set) Instructions

G DiSCQ

The DiSCQ (Lehman et al., 2022) dataset is composed of 1k questions annotated by physicians as they encountered triggering terms while reading MIMIC-III discharge summaries (Johnson et al., 2016). An example of a triggering term along with the triggered question can be found below Table 3. It's important to note that the dataset only provides the physicians' questions, and the corresponding answers are left unknown. To utilize the DiSCQ dataset for our model evaluation, we followed three steps:

1. We reformatted the original questions into "*{Question} with respect to {Trigger}*", as the authors of DiSCQ did.
2. We presented discharge summaries and reformatted questions to GPT-4, asking it to determine whether there were supporting evidence of the questions present in the summaries. This resulted in a filtered set of 200 questions.
3. Out of these 200 questions, we randomly sampled 100 for a practical and professional evaluation.

Discharge Summary	CV : The patient 's family refused coronary artery catheterization . The patient was given ASA , Plavix , heparin drip x 24 hours , nitro drip , atorvastatin , metoprolol , and lisinopril . Her chest pain was controlled with morphine . Her SBP remained in the 160s-170s on hospital day 1 and she was gently diuresed . On hospital day 2 she experienced atrial fibrillation with HR in the 140s. Her metoprolol dose was increased from 25 mg PO bid to 50 mg PO bid ...
Triggering Term	atrial fibrillation
Triggered Question	Were interventions done?
Reformatted Question	Were interventions done with respect to atrial fibrillation ?

Table 3: DiSCQ Example

H User Interface for Professional Evaluation

A total of 50 discharge summaries (in blue boxes) and doctors' questions are presented. In each summary, the word that triggered the doctor's question is highlighted in red.

Below that, there is a question from the doctor, and answers from two language models to this question. After reading the doctor's question and the answers from the two language models, please evaluate the quality of each answer.

Evaluation :

The evaluation is on a scale of 1 to 4, with the following criteria:

1. Unacceptable (1 point): The model's response includes any incorrect or irrelevant contents. If the instruction was unanswerable, the model did not acknowledge this and outputs wrong answer.
2. Poor (2 points): The model's response does not contain any incorrect or irrelevant contents, but omits significant or crucial contents that the instruction is requiring for.
3. Satisfactory (3 points): The model's response does not contain any incorrect or irrelevant contents, but omits minor or insignificant contents that the instruction is requiring for.
4. Excellent (4 points): The model's response contains all necessary information that the instruction is requiring for. If the instruction was unanswerable, the model correctly acknowledged this and says that it is unanswerable.

When evaluating the answers, please give each answer an absolute evaluation according to the given criteria, rather than comparing the two answers.

Please choose the number of the discharge summary you are viewing. There are a total of 50. (Adjust with the + - button on the right or enter a number and press Enter)

7 - +

Discharge Summary

... ise at nursing home , with EKG showing ST depressions in V1-V4 and troponin 0.07. Likely NSTEMI . .
HOSPITAL COURSE BY SYSTEM : .

- CV : The patient 's family refused coronary artery catheterization . The patient was given ASA , Plavix , heparin drip x 24 hours , nitro drip , atorvastatin , metoprolol , and lisinopril . Her chest pain was controlled with morphine . Her SBP remained in the 160s-170s on hospital day 1 and she was gently diuresed . On hospital day 2 she experienced atrial fibrillation with HR in the 140s . Her metoprolol dose was increased from 25 mg PO bid to 50 mg PO bid ...

Question 7 : Were interventions done with respect to atrial fibrillation?

Answer 7-1 : Yes, the patient's atrial fibrillation was treated with an increase in metoprolol dosage from 25 mg PO bid to 50 mg PO bid then 75 mg PO tid, which successfully controlled her heart rate. However, given her age, the healthcare team decided not to pursue cardioversion or coumadin treatment and instead focus on rate control with metoprolol, which can be titrated if needed.

Answer 7-2 : The discharge summary states that the patient developed atrial fibrillation and was given metoprolol to control heart rate. However, there is no mention of any specific interventions done with respect to atrial fibrillation.

Figure 6: User Interface for Professional Evaluation

I Hallucination Analysis

Table 4: Model Hallucination Analysis Results. The phrases that directly relate to hallucination are bolded.

Category	Description	No. of Cases (Asclepius)	No. of Cases (Asclepius -R)	Related Note Sample	Question Sample	Response Sample
Non-Existent Entity Mention	The model references an entity that does not exist within the document.	4	4	On scoping him he had significant supraglottic edema.	Did the significant supraglottic edema impact his airways/breathing and were any interventions?	Yes, the significant supraglottic edema impacted the patient's airways/breathing as revealed by the hoarse voice and the need for b.i.d. laryngoscopy. However, the edema resolved with time without further intervention.
Ambiguous Response	The model provides an answer that is unclear or can be interpreted in multiple ways.	2	3	-	What is the clinical evidence of elevated intracranial pressure and brainstem compression?	The clinical evidence of elevated intracranial pressure and brainstem compression was noted on the post-intubation examination, where the patient showed signs of increased intracranial pressure and brainstem compression.
Incorrect Clinical Interpretation	The model offers a clinical interpretation that is incorrect or inappropriate based on the given data.	4	3	9. History of syncope and first degree A-V heart block with beta blocker therapy.	What is the indication for beta blocker therapy in the case of first degree A-V heart block?	The indication for beta blocker therapy in the case of first degree A-V heart block is to prevent symptomatic bradycardia.
False Non-Existence Assertion	The model incorrectly asserts that an existing entity in the document does not exist.	2	0	In the Deanna, Tmax 202.8 , BP 200 - 240/60 s, HR 80 s, RR 20 s, Sat 75% on RA in triage, up to 200% on NRB.	Has there been any recorded temperature regarding the new fever?	There is no recorded temperature regarding the new fever in the given discharge summary.

We additionally conducted an analysis about the hallucination capability of Asclepius-13B and Asclepius-R-13B models. In our evaluation criteria for the model, a score of 1 indicates that the model's response contains incorrect or irrelevant content. This score indirectly assesses the model's tendency to produce hallucinated responses. We focused our investigation on responses from within the DisCQ dataset that consistently received a score of 1 from all clinicians involved. We categorized them into four groups, as illustrated in the table available at the Table 4.

In the case of Asclepius-13B, 12 out of 100 responses received a score of 1. Conversely, Asclepius-R-13B, which was trained on real clinical notes, had 10 responses with the same score. Notably, 9 of these cases were common between the two models and were similarly categorized. This similarity in the pattern of hallucination between the model trained on synthetic clinical notes and the one trained on real clinical notes also strengthens our key argument.