

ClinicalMamba: A Generative Clinical Language Model on Longitudinal Clinical Notes

Zhichao Yang¹, Avijit Mitra¹, Sunjae Kwon¹, Hong Yu^{1,2}

¹ College of Information and Computer Sciences, University of Massachusetts Amherst

² Department of Computer Science, University of Massachusetts Lowell

{zhichaoyang, avijitmitra, sunjaekwon}@umass.edu hong_yu@uml.edu

Abstract

The advancement of natural language processing (NLP) systems in healthcare hinges on language models' ability to interpret the intricate information contained within clinical notes. This process often requires integrating information from various time points in a patient's medical history. However, most earlier clinical language models were pretrained with a context length limited to roughly one clinical document. In this study, We introduce ClinicalMamba, a specialized version of the Mamba language model, pretrained on a vast corpus of longitudinal clinical notes to address the unique linguistic characteristics and information processing needs of the medical domain. ClinicalMamba models, with 130 million and 2.8 billion parameters, demonstrate superior performance in modeling clinical language across extended text lengths compared to Mamba and other clinical models based on longformer and Llama. With few-shot learning, ClinicalMamba achieves notable benchmarks in speed and performance, outperforming existing clinical language models and large language models like GPT-4 in longitudinal clinical tasks.

1 Introduction

Clinical narratives, such as patient histories, consultation notes, and discharge summaries, contain detailed and complex information that extends over long text sequences (Wu et al., 2019). To fully understand a patient's condition, treatments, and outcomes, NLP systems need to integrate information from various parts of these narratives, which often requires understanding the context provided in those long form text (Blumenthal, 2010).

Understanding the sequence of health events is crucial for diagnoses, treatment plans, and patient monitoring (Wang et al., 2024; Yang et al., 2023; Kraljevic et al., 2023; Eva, 2005). This often involves putting together information from different

time points within a patient's health history (Gao et al., 2024). Long context enables NLP systems to perform temporal reasoning by tracking events over time longitudinally, which is essential for tasks like predicting disease progression or extracting medical relation (Chen et al., 2023; Jia et al., 2019; Wiegrefe et al., 2019).

It becomes imperative to design models for the need for processing longer texts (Parmar et al., 2023; Tay et al., 2020). Prior studies have introduced Mamba (Gu and Dao, 2023), a selective state space model, that selects and compresses all necessary information into latent space from context, and achieves linear-time efficiency with context length. While these advancements have been primarily directed towards processing general domain text, the unique linguistic features of clinical narratives differ significantly from general domain (Lehman et al., 2023), motivating us to develop specialized Mamba models in the clinical domain.

In this work, we build and publicly release ClinicalMamba - a Mamba model pretrained on longitudinal clinical notes. Furthermore, we demonstrate that ClinicalMamba outperforms multiple language models on longitudinal clinical NLP tasks. In particular, our contributions are as follows:

- We publicly release ClinicalMamba with 130m and 2.8b parameters trained on MIMIC-III (Johnson et al., 2016).¹
- Through distributed training, ClinicalMamba-2.8b model was pretrained in under 60 hours on 4 A100 GPUs and it is the first clinical autoregressive language model with a 16k maximum token length.
- Through few-shot prompt-based finetuning, we demonstrate both ClinicalMamba outperforms original Mamba, GPT4, and other existing clinical long context language models

¹<https://github.com/whaleloops/ClinicalMamba>

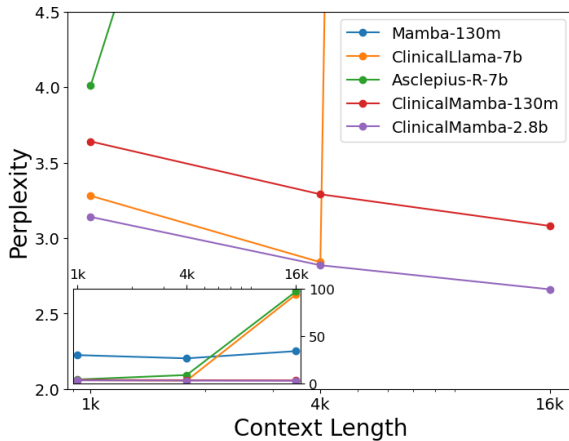


Figure 1: Perplexity of different generative language models on MIMIC-III when evaluated at various preceding context lengths (1k, 4k, and 16k tokens). The X-axis is in the log scale. The subfigure is a zoom-out plot with perplexity ranges 0-100. Experiment settings and detailed results are in section 5.

on well-established long context clinical information extractions tasks: cohort selection for clinical trial and international classification of diseases (ICD) coding.

2 Related Work

2.1 Pretraining clinical narratives

The rapid expansion of the utilization of electronic health records (EHRs) into the healthcare landscape underscores an urgent need for a clinical language model (Kang et al., 2019). Previous work, such as Alsentzer et al. (2019) on Clinical BERT embeddings and Huang et al. (2019) with ClinicalBERT and Lewis et al. (2020) with ClinicalRoberta, adapted general-purpose language models to the clinical domain to enhance performance on clinical tasks. These models have been pivotal in demonstrating the effectiveness of adapting general-purpose NLP tools to the intricacies of clinical text. Similarly, the creation of GatorTron (Yang et al., 2022a) scales up clinical language models to billions of parameters, while NYUTron (Jiang et al., 2023) harness billions of unstructured data found in EHRs. Both underscores the potential of domain-adapted language models to advance clinical NLP by improving performance across various tasks such as concept extraction and outcome prediction (Yang et al., 2022a; Jiang et al., 2023).

To handle complex and nuanced tasks, recent studies investigated training generative models with prompt (Kweon et al., 2023; Peng et al., 2023;

Wang et al., 2023; Lu et al., 2022; Wang and Sun, 2022). These models not only excel in classification but also in generating clinically relevant text that can be indistinguishable from human-written notes. Most previous methods focus on pretraining transformer models with a context window less than 2k tokens. However, we pretrained a selective state space model with a context window of 16k tokens, which includes more than 98% of the visits in MIMIC-III.

2.2 Clinical information extraction on long document

Handling long texts in clinical NLP has always been challenging. Traditional methods of information extraction tackle this by marking specific locations within the sentence, but such labeling is not always available, and hiring annotators can be costly (Fu et al., 2020; Mitra et al., 2023). Recent advancements in document information extraction involve pairing labels with documents (Kwon et al., 2022; Deshpande et al., 2024). However, BERT and Roberta struggle with processing these lengthy documents directly.

To address this, prior research introduced *Hierarchical-ClinicalRoberta*, which involves breaking down long documents into shorter segments of 512 tokens, applying ClinicalRoberta to each segment to obtain embeddings, and then using additional layers to leverage these embeddings for label classification (Huang et al., 2022; Zhang and Jankowski, 2022). However, this method combines information from each segment only at the final layer, which can hinder performance when training data is limited.

To mitigate this issue, *ClinicalLongformer* is designed to efficiently process longer context length by employing a self-attention mechanism across all layers, which is key to its proficiency in managing dense information exchanges within a specified contextual range (Li et al., 2022; Ji et al., 2023). This mechanism, while powerful, is limited by its focus on a predetermined window of text, restricting its scope to what falls within this window.

To overcome these limitations, the Mamba model emerges as a revolutionary approach. Mamba employs a selective state space model strategy to meticulously choose critical data for incorporation into its state (Gu and Dao, 2023), thereby, enhancing its capability to manage information beyond the conventional self-attention window. In

general domain language modeling, Mamba surpasses Transformers of equivalent size in task performance and speed.

3 Methods

3.1 Pretraining

We gather 82,178 hospital visits along with their deidentified free-text clinical notes (2,083,180) from 46,520 patients in MIMIC-III (Johnson et al., 2016). Rather than breaking down the notes into chunks of 512 tokens to act as individual data instances, we aggregate all notes related to a visit longitudinally. The distribution of token counts per data instance is detailed in Table A.1. For information on our text pre-processing methods, please refer to section A.1.

Following previous works (Li et al., 2019; Lee et al., 2019), we continue to pretrain Mamba using MIMIC-III clinical notes with the causal language modeling objective. This pretraining process utilizes 4 Nvidia A100-80GB GPUs. It’s important to note that some of our downstream evaluation tasks utilize a small subset (6,049) of hospital visits from MIMIC-III, so we exclude them from the pretraining data. A comprehensive training recipe is available in section A.2.

3.2 Prompt based fine-tuning

We leverage the inherent capabilities of pre-trained language models by introducing a novel fine-tuning strategy that aligns with the specific demands of few-shot learning in clinical NLP. Recognizing the limitations of traditional fine-tuning methods when applied to clinical NLP tasks with limited labeled data, we adapted a prompt-based fine-tuning mechanism following previous works (Gao et al., 2021; Yang et al., 2022b; Taylor et al., 2023). Specifically, we first identify a set of representative prompts that encapsulate key aspects of the clinical tasks, such as the patient’s alcohol consumption. These prompts are then appended after each input clinical note and incorporated into the fine-tuning phase, where the language model learns to associate them with label tokens (Yes/No) based on a limited dataset. The generated label tokens are then mapped to label space.

As shown in Figure 2, we transfer the downstream information extraction task into label token generation, which is similar to next token prediction during pretraining. We prompt based fine-tuned mamba and other baseline models unless

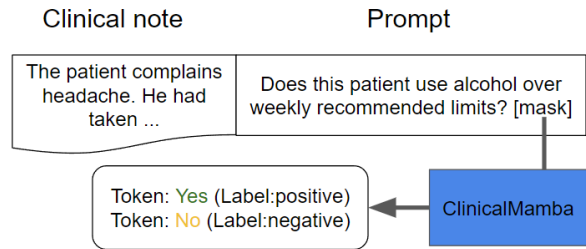


Figure 2: Illustration of Prompt-based fine-tuning.

otherwise specified on the following tasks.

3.3 Fine-tuning tasks

Cohort selection for clinical trial addresses the challenge of interpreting unstructured clinical narratives to streamline the patients selection process (Wornow et al., 2024; Jin et al., 2023; Wong et al., 2023). It aims to classify patients based on whether they meet 13 specific eligibility criteria, such as the usage of aspirin to prevent myocardial infarction, excessive alcohol consumption, and HbA1c values between 6.5 and 9.5%, among others. The input contains multiple clinical notes with a total length of 4924 tokens on average. This dataset was released as part of n2c2 challenge (track 1) in 2018 (Stubbs et al., 2019).

ICD coding interprets complex clinical narratives, translating them into standardized codes that facilitate accurate billing, statistical analysis, and healthcare management. It aims to extract patient’s disease and procedure codes from clinical text. We followed general instructions from Mullenbach et al. (2018) in building this task, but instead of using a single discharge summary as input, we used all previous discharge summaries and assigned ICD code descriptions from previous visits. We further filtered 50 infrequent codes as Code-rare and 50 frequent codes as Code-common following Yang et al. (2022b). The average length is 4,223 and 7,062 tokens respectively. Detailed dataset statistics are shown in Table A.1.

4 Experiments

For each fine-tuning task, we measured the micro precision, micro recall, micro F1 scores, and micro receiver operating characteristic/area under the curve (ROCAUC). We compared our model with the following baselines:

GPT-4 is a large language model designed to understand and generate human-like text based on the input it receives. We applied zero-shot prompting

to each downstream task, using ACAN and original prompt introduced in [Wornow et al. \(2024\)](#). GPT-4 (version 2023-12-01-preview) was accessed securely through the Azure OpenAI API. We opt out of human review of the data by signing the Additional Use Case Form². We set the sampling temperature for decoding to 0.1.

Asclepius-R ([Kweon et al., 2023](#)) is a clinical generative language model trained on MIMIC-III discharge summaries and corresponding instruction-answer pairs. It has 7 billion parameters with a maximum input of 4096 tokens.

ClinicalLlama2 (CLlama2) is similar to Asclepius-R, but it was trained on all types of MIMIC-III note with the same computation budget as ClinicalMamba-2.8b (60 hours on 4 A100). It has 7 billion parameters with 4096 max context length.

ClinicalLongformer (CLongformer) ([Li et al., 2022](#)) is a clinical knowledge enriched version of Longformer that was further pretrained using MIMIC-III clinical notes. It has 149 million parameters with a maximum input of 4096 tokens. We only used local attention and does not apply global attention for computation efficiency.

Hierarchical ClinicalRoberta (Hi-CRoberta) ([Huang et al., 2022](#)), utilizes multiple embedding from clinical Roberta ([Lewis et al., 2020](#)). It first segment clinical notes into chunks of 512 tokens to obtain their embeddings, embeddings are then pooled by concatenation and finally a linear classification head during downstream task. It has 110 million parameters with a max of 16384 tokens. We did not apply prompt based fine-tuning as this is not a generative model.

MultiResCNN ([Li and Yu, 2020](#)) encode free text with Multi-Filter ResidualCNN, and applied label code attention mechanism to enable each ICD code to attend different parts of the document.

5 Results & Discussions

In this section, we will first compare the model’s language modeling ability on MIMIC-III clinical notes. We will then describe the evaluation on different clinical information extraction tasks. Finally, we will share a case study, which illustrates that our ClinicalMamba can recall patient information with long history.

ClinicalMamba stands as the sole model capable of handling clinical notes of up to 16k tokens.

²<https://physionet.org/news/post/gpt-responsible-use>

Model	Prec	Recall	F1	AUC
CLlama2	70.0	79.1	77.7	84.3
Hi-CRoberta	72.4	82.6	79.2	88.1
CLongformer	69.7	78.6	76.1	83.5
GPT-4	88.1	79.9	84.8	-
Mamba-130m	75.4	80.2	77.7	85.7
CMamba-130m	79.0	86.2	82.2	91.8
CMamba-2.8b	88.6	89.5	88.8	95.7

Table 1: Results on cohort selection task, where C is model pretrained in clinical domain.

As demonstrated in Figure 1, the perplexity for ClinicalMamba-2.8b decreased from 3.11 to 2.61 as the context length expanded from 1k to 16k tokens during inference. This is in contrast to the performance of prior clinical autoregressive language models, where perplexity levels rose with increased context lengths. For instance, with ClinicalLlama-7b, perplexity escalated from 2.82 to 94.02 as the context length grew from 4k to 16k. This limitation arises because these models were trained on contexts not exceeding 4k, impairing their accuracy for next token prediction when given previous contexts beyond 4k.

In the domain of extracting information from longitudinal clinical records, ClinicalMamba demonstrates superior performance compared to Mamba. ClinicalMamba achieved ROCAUC scores of 91.8, 42.3, and 94.2 on Cohort selection, Code-rare, and Code-common, while Mamba obtained ROCAUC scores of 85.7, 37.8, and 92.8 respectively. ClinicalMamba also outperformed previous long-range clinical language models with similar number of parameters. ClinicalMamba significantly outperformed Hierarchical-ClinicalRoberta and Clinical-Longformer by relatively 52.7% and 19.1% on ROCAUC respectively. This is particularly notable in the Code-rare task with limited training data (5 shots), where ClinicalMamba attained an AUC of 91.1, compared to 77.1 of Hierarchical-ClinicalRoberta and 80.5 of ClinicalLongformer.

Surprisingly, ClinicalMamba-2.8b also outperformed zero-shot GPT-4, achieving F1 scores of 88.8, 56.6, and 73.6 on Cohort selection, Code-rare, and Code-common tasks, whereas GPT-4 obtained a F1 score of 84.8, 33.3, and 68.2 respectively.

We also present a case with long history attached in the supplementary material. Both ClinicalMamba and ClinicalLongformer models identi-

Model	Code-rare				Code-common			
	Prec	Recall	F1	AUC	Prec	Recall	F1	AUC
MultiResCNN	20.34	2.07	5.19	47.20	70.50	60.78	66.24	92.04
Hi-CRoberta	46.19	10.96	16.74	77.11	73.76	65.01	69.23	93.14
CLongformer	50.27	17.81	28.69	80.52	78.42	64.97	71.14	94.24
GPT-4	30.91	36.12	33.29	-	72.48	62.28	68.19	-
Mamba-130m	57.75	28.08	37.79	84.80	73.71	62.87	68.94	92.75
CMamba-130m	70.97	30.14	42.31	91.08	76.82	68.03	74.34	94.23
CMamba-2.8b	75.28	45.89	56.51	92.75	75.53	72.12	73.64	94.54

Table 2: Results on ICD coding task, where C indicates model pretrained in clinical domain.

fied that this case met the criteria for history of advanced cardiovascular disease, and were tasked to interpret the prediction using SHAP (Lundberg and Lee, 2017). ClinicalMamba identified sentence “FINAL DIAGNOSIS: Acute MI” with highest SHAP value. This sentence locates at the first note among the patient history. In contrast, ClinicalLongformer identified a sentence, the patient “will be admitted in stable condition for further evaluation and rule out for myocardial infarction”, at the last note. ClinicalLongformer also misinterpreted the negation in this sentence. In this case, ClinicalMamba remembered the correct long-distant sentence that leads to the correct answer, while ClinicalLongformer preferred wrong short-distant sentence to support its prediction.

6 Conclusion

In this study, we developed and released Mamba models pretrained on a large collection of clinical notes. Our findings demonstrate the superior performance of our ClinicalMamba in extracting information from long text documents compared to other models. We strongly believe that clinical NLP researchers can benefit from such long-context generative language models that alleviates the need of a substantial computational power, without any performance trade-off. Building on the groundwork laid by this study, future endeavors can further refine and expand the capabilities of ClinicalMamba.

Limitations

This work has several notable limitations. First, we do not experiment with more recent parameter-efficient fine-tuning strategies such as soft prompting (Lester et al., 2021) and Low-Rank Adaptation (LoRa) (Hu et al., 2021). This potentially

undermined ClinicalMamba on downstream tasks. Second, our adaptation of the Mamba framework was restricted solely to textual data documented during visits. EHRs are rich with multifaceted information, including but not limited to radiology images taken at different times and Electrocardiogram waveforms that span various periods. Future research could develop a multimodal Mamba framework to leverage all other modalities. Third, the MIMIC-III dataset, which serves as the foundation of our study, only includes notes from the intensive care unit of a single hospital within the United States. This limits the generalizability of our findings, as care practices vary significantly across different institutions and countries. We did not pretrain on MIMIC-IV because it only has a limited number of notes (and also limited type: discharge summary and radiology report) per visit. Lastly, the linguistic scope of the MIMIC dataset is limited to English, which presents a barrier to understanding and applying our findings in non-English speaking contexts. Addressing these limitations could substantially broaden the applicability and relevance of our work in future endeavors.

Ethics Statement

In this research, we gained authorized access to the MIMIC and N2C2 dataset and used de-identified clinical notes following their license agreement and HIPAA regulations. When language models are trained on extensive clinical text, they can inherit biases within the data. For instance, they might prefer inquiries concerning smoking habits or link specific medical conditions to certain demographic groups. These biases could be mitigated by enhancing model alignment with each patient’s background.

References

- Emily Alsentzer, John Murphy, William Boag, Wei-Hung Weng, Di Jindi, Tristan Naumann, and Matthew McDermott. 2019. [Publicly available clinical BERT embeddings](#). In *Proceedings of the 2nd Clinical Natural Language Processing Workshop*, pages 72–78, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- David Blumenthal. 2010. [Launching hitech](#). *The New England journal of medicine*, 362 5:382–5.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Zekai Chen, Mariann Micsinai Balan, and Kevin Brown. 2023. [Language models are few-shot learners for prognostic prediction](#). *ArXiv*, abs/2302.12692.
- Vijeta Deshpande, Minhwa Lee, Zonghai Yao, Zihao Zhang, Jason Brian Gibbons, and Hong Yu. 2024. [Localtweets to localfealth: A mental health surveillance framework based on twitter data](#).
- K. Eva. 2005. [What every teacher needs to know about clinical reasoning](#). *Medical Education*, 39.
- Sunyang Fu, David C. Chen, Huan He, Sijia Liu, Sungrim Moon, Kevin J. Peterson, Feichen Shen, Liwei Wang, Yanshan Wang, Andrew Wen, Yiqing Zhao, Sunghwan Sohn, and Hongfang Liu. 2020. [Clinical concept extraction: A methodology review](#). *Journal of biomedical informatics*, page 103526.
- Shanghai Gao, Teddy Koker, Owen Queen, Thomas Hartvigsen, Theodoros Tsiligkaridis, and Marinka Zitnik. 2024. [Units: Building a unified time series model](#).
- Tianyu Gao, Adam Fisch, and Danqi Chen. 2021. [Making pre-trained language models better few-shot learners](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3816–3830, Online. Association for Computational Linguistics.
- Albert Gu and Tri Dao. 2023. [Mamba: Linear-time sequence modeling with selective state spaces](#). *ArXiv*, abs/2312.00752.
- J. Edward Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, and Weizhu Chen. 2021. [Lora: Low-rank adaptation of large language models](#). *ArXiv*, abs/2106.09685.
- Chao-Wei Huang, Shang-Chi Tsai, and Yun-Nung Chen. 2022. [PLM-ICD: Automatic ICD coding with pre-trained language models](#). In *Proceedings of the 4th Clinical Natural Language Processing Workshop*, pages 10–20, Seattle, WA. Association for Computational Linguistics.
- Kexin Huang, Jaan Altosaar, and Rajesh Ranganath. 2019. [Clinicalbert: Modeling clinical notes and predicting hospital readmission](#). *ArXiv*, abs/1904.05342.
- Shaoxiong Ji, Tianlin Zhang, Kailai Yang, Sophia Ananiadou, E. Cambria, and Jörg Tiedemann. 2023. [Domain-specific continued pretraining of language models for capturing long context in mental health](#). *ArXiv*, abs/2304.10447.
- Robin Jia, Cliff Wong, and Hoifung Poon. 2019. [Document-level n-ary relation extraction with multi-scale representation learning](#). 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 3693–3704, Minneapolis, Minnesota. Association for Computational Linguistics.
- Lavender Yao Jiang, Xujin C. Liu, Nima Pour Nejatian, Mustafa Nasir-Moin, Duo Wang, Anas Abidin, Kevin Eaton, Howard A. Riina, Ilya Laufer, Paawan Punjabi, Madeline Miceli, Nora C. Kim, Cordelia Orillac, Zane Schnurman, Christopher Livia, Hannah Weiss, David Kurland, Sean N. Neifert, Yosef Dastagirzada, Douglas Kondziolka, Alexander T. M. Cheung, Grace Yang, Mingzi Cao, Mona G. Flores, Anthony B Costa, Yindalon Aphinyanaphongs, Kyunghyun Cho, and Eric Karl Oermann. 2023. [Health system-scale language models are all-purpose prediction engines](#). *Nature*, 619:357 – 362.
- Qiao Jin, Zifeng Wang, Charalampos S Floudas, Jimeng Sun, and Zhiyong Lu. 2023. [Matching patients to clinical trials with large language models](#). *ArXiv*.
- Alistair E. W. Johnson, Tom J. Pollard, Lu Shen, Li wei H. Lehman, Mengling Feng, Mohammad Mahdi Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G. Mark. 2016. Mimic-iii, a freely accessible critical care database. *Scientific Data*, 3.
- Tian Kang, Shirui Zou, and Chunhua Weng. 2019. [Pre-training to recognize pico elements from randomized controlled trial literature](#). *Studies in health technology and informatics*, 264:188 – 192.
- Zeljko Kraljevic, Dan Bean, Anthony Shek, Rebecca Bendayan, Harry Hemingway, Joshua Au Yeung, Alexander Deng, Alfie Baston, Jack Ross, Esther Idowu, James T Teo, and Richard J Dobson. 2023. [Foresight – generative pretrained transformer \(gpt\) for modelling of patient timelines using ehers](#).

- Sunjun Kweon, Junu Kim, Jiyou Kim, Sujeong Im, Eunbyeol Cho, Seongsu Bae, Jungwoo Oh, Gyubok Lee, Jong Hak Moon, Seng Chan You, Seungjin Baek, Chang Hoon Han, Yoon Bin Jung, Yohan Jo, and E. Choi. 2023. [Publicly shareable clinical large language model built on synthetic clinical notes](#). *ArXiv*, abs/2309.00237.
- Sunjae Kwon, Zonghai Yao, Harmon Jordan, David Levy, Brian Corner, and Hong Yu. 2022. [MedJEx: A medical jargon extraction model with Wiki’s hyperlink span and contextualized masked language model score](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11733–11751, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2019. [Biobert: a pre-trained biomedical language representation model for biomedical text mining](#). *Bioinformatics*, 36:1234 – 1240.
- Eric P. Lehman, Evan Hernandez, Diwakar Mahajan, Jonas Wulff, Micah J. Smith, Zachary M. Ziegler, Daniel Nadler, Peter Szolovits, Alistair E. W. Johnson, and Emily Alsentzer. 2023. [Do we still need clinical language models?](#) *ArXiv*, abs/2302.08091.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. [The power of scale for parameter-efficient prompt tuning](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3045–3059, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Patrick Lewis, Myle Ott, Jingfei Du, and Veselin Stoyanov. 2020. [Pretrained language models for biomedical and clinical tasks: Understanding and extending the state-of-the-art](#). In *Proceedings of the 3rd Clinical Natural Language Processing Workshop*, pages 146–157, Online. Association for Computational Linguistics.
- Fei Li, Yonghao Jin, Weisong Liu, Bhanu Pratap Singh Rawat, Pengshan Cai, and Hong Yu. 2019. [Fine-tuning bidirectional encoder representations from transformers \(bert\)-based models on large-scale electronic health record notes: An empirical study](#). *JMIR Med Inform*, 7(3):e14830.
- Fei Li and Hong Yu. 2020. [Icd coding from clinical text using multi-filter residual convolutional neural network](#). *Proceedings of the AAAI Conference on Artificial Intelligence*. AAAI Conference on Artificial Intelligence, 34 5:8180–8187.
- Yikuan Li, Ramsey M. Wehbe, Faraz S. Ahmad, Hanyin Wang, and Yuan Luo. 2022. [Clinical-longformer and clinical-bigbird: Transformers for long clinical sequences](#). *ArXiv*, abs/2201.11838.
- Qiuhaio Lu, Dejing Dou, and Thien Nguyen. 2022. [ClinicalT5: A generative language model for clinical text](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 5436–5443, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Scott M. Lundberg and Su-In Lee. 2017. [A unified approach to interpreting model predictions](#). In *Neural Information Processing Systems*.
- Avijit Mitra, Richeek Pradhan, Rachel D. Melamed, Kun Chen, David C. Hoaglin, Katherine L. Tucker, Joel I. Reisman, Zhichao Yang, Weisong Liu, Jack Tsai, and Hong Yu. 2023. [Associations Between Natural Language Processing–Enriched Social Determinants of Health and Suicide Death Among US Veterans](#). *JAMA Network Open*, 6(3):e233079–e233079.
- Niklas Muennighoff, Alexander M. Rush, Boaz Barak, Teven Le Scao, Aleksandra Piktus, Nouamane Tazi, Sampo Pyysalo, Thomas Wolf, and Colin Raffel. 2023. [Scaling data-constrained language models](#). *ArXiv*, abs/2305.16264.
- James Mullenbach, Sarah Wiegrefe, Jon Duke, Jimeng Sun, and Jacob Eisenstein. 2018. [Explainable prediction of medical codes from clinical text](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1101–1111, New Orleans, Louisiana. Association for Computational Linguistics.
- Mihir Parmar, Aakanksha Naik, Himanshu Gupta, Disha Agrawal, and Chitta Baral. 2023. [Longbox: Evaluating transformers on long-sequence clinical tasks](#). *ArXiv*, abs/2311.09564.
- C.A.I. Peng, Xi Yang, Aokun Chen, Kaleb E. Smith, Nima M. Pournajatian, Anthony B Costa, Cheryl Martin, Mona G. Flores, Ying Zhang, Tanja Magoc, Gloria P. Lipori, Duane A. Mitchell, Naykky Maruquel Singh Ospina, Mustafa Mamon Ahmed, William R. Hogan, Elizabeth A. Shenkman, Yi Guo, Jiang Bian, and Yonghui Wu. 2023. [A study of generative large language model for medical research and healthcare](#). *NPJ Digital Medicine*, 6.
- Amber Stubbs, Michele Filannino, Ergin Soysal, Sam Henry, and Özlem Uzuner. 2019. [Cohort selection for clinical trials: n2c2 2018 shared task track 1](#). *Journal of the American Medical Informatics Association : JAMIA*.
- Yi Tay, Mostafa Dehghani, Samira Abnar, Yikang Shen, Dara Bahri, Philip Pham, Jinfeng Rao, Liu Yang, Sebastian Ruder, and Donald Metzler. 2020. [Long range arena: A benchmark for efficient transformers](#). *ArXiv*, abs/2011.04006.
- Niall Taylor, Yi Zhang, Dan W. Joyce, Ziming Gao, Andrey Kormilitzin, and Alejo Nevado-Holgado. 2023. [Clinical prompt learning with frozen language models](#). *IEEE Transactions on Neural Networks and Learning Systems*, pages 1–11.

- Jiaqi Wang, Junyu Luo, Muchao Ye, Xiaochen Wang, Yuan Zhong, Aofei Chang, Guanjie Huang, Ziyi Yin, Cao Xiao, Jimeng Sun, and Fenglong Ma. 2024. [Recent advances in predictive modeling with electronic health records](#). pages 1767–1781, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Ning Zhang and Maciej Jankowski. 2022. [Hierarchical bert for medical document understanding](#). *ArXiv*, abs/2204.09600.
- Yuqing Wang, Yun Zhao, and Linda Petzold. 2023. [Are large language models ready for healthcare? a comparative study on clinical language understanding](#).
- Zifeng Wang and Jimeng Sun. 2022. [PromptEHR: Conditional electronic healthcare records generation with prompt learning](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2873–2885, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Sarah Wiegrefe, Edward Choi, Sherry Yan, Jimeng Sun, and Jacob Eisenstein. 2019. [Clinical concept extraction for document-level coding](#). In *Proceedings of the 18th BioNLP Workshop and Shared Task*, pages 261–272, Florence, Italy. Association for Computational Linguistics.
- Cliff Wong, Sheng Zhang, Yu Gu, Christine Mung, Jacob Abel, Naoto Usuyama, Roshanthi Weerasinghe, Brian Piening, Tristan Naumann, Carlo Bifulco, and Hoifung Poon. 2023. [Scaling clinical trial matching using large language models: A case study in oncology](#).
- Michael Wornow, Alejandro Lozano, Dev Dash, Jenelle A. Jindal, Kenneth W. Mahaffey, and Nigam H. Shah. 2024. [Zero-shot clinical trial patient matching with llms](#). *ArXiv*, abs/2402.05125.
- Stephen T Wu, Kirk Roberts, Surabhi Datta, Jingcheng Du, Zongcheng Ji, Yuqi Si, Sarvesh Soni, Qiong Wang, Qiang Wei, Yang Xiang, Bo Zhao, and Hua Xu. 2019. [Deep learning in clinical natural language processing: a methodical review](#). *Journal of the American Medical Informatics Association : JAMIA*.
- Xi Yang, Aokun Chen, Nima M. Pournejatian, Hoo-Chang Shin, Kaleb E. Smith, Christopher Parisien, Colin B. Compas, Cheryl Martin, Anthony B Costa, Mona G. Flores, Ying Zhang, Tanja Magoc, Christopher A. Harle, Gloria P. Lipori, Duane A. Mitchell, William R. Hogan, Elizabeth A. Shenkman, Jiang Bian, and Yonghui Wu. 2022a. [A large language model for electronic health records](#). *NPJ Digital Medicine*, 5.
- Zhichao Yang, Avijit Mitra, Weisong Liu, Dan Berlowitz, and Hong Yu. 2023. [Transformehr: transformer-based encoder-decoder generative model to enhance prediction of disease outcomes using electronic health records](#). *Nature Communications*, 14.
- Zhichao Yang, Shufan Wang, Bhanu Pratap Singh Rawat, Avijit Mitra, and Hong Yu. 2022b. [Knowledge injected prompt based fine-tuning for multi-label few-shot ICD coding](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*,

A Appendix

A.1 Text preprocessing

We followed [Huang et al. \(2019\)](#) to format notes during text preprocessing. But We did not convert text to lowercase because Mamba tokenizer is able to process both upper and lower cases. For notes on each patient’s hospital visit, we sorted notes by their charted date and concatenated notes into one string. We used string "- - {NoteType} note - -" to separate the notes. Table A.2 shows comprehensive values of {NoteType}.

For pretraining data, we truncate notes with more than 16k tokens, however, this is only less than 2%, a length distribution is provided in Figure A.1. We exclude a small subset (6,049) of hospital visits due to the evaluation of MIMIC ICD coding and MIMIC hospital readmission prediction, The visit ids (hadm_id) are documented in the github.

A.2 Pretraining recipe

ClinicalMamba-2.8b is a selective state space model designed using replication of the Mamba architecture ([Gu and Dao, 2023](#)). ClinicalMamba refers to the class of models, while 2.8b represents the number of parameters of this particular pre-trained model. We also pretrained ClinicalMamba-130m using the same pretraining data from the previous section. The specific values of hyperparameters are shown in Table A.3. These models were trained for 763 million English tokens over 7000 steps (3 epochs) ([Muennighoff et al., 2023](#)). It was trained as an autoregressive language model, using cross-entropy loss ([Brown et al., 2020](#)). For learning rate scheduling, we followed Mamba and chose linear learning rate warmup with cosine decay to $1e - 5$. We found this important setting to avoid loss overflow. It took under 60 hours to pretrain ClinicalMamba-2.8b in on 4 Nvidia Tesla A100-80GB GPUs.

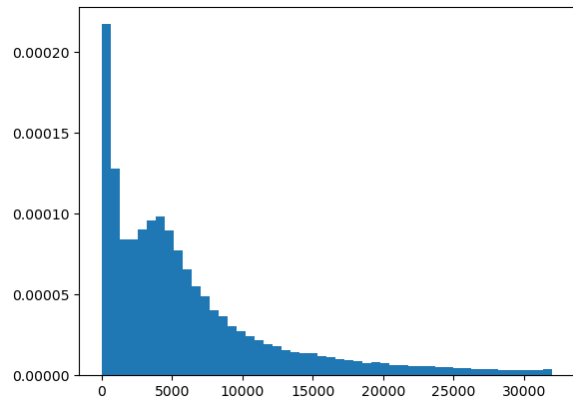


Figure A.1: Long tail distribution of number of tokens per each visit. Y-axis is the density (sum to 1.0).

		Cohort selection	Code-rare	Code-common
shots	mean	89	5	918
tokens	mean	4924	4223	7062
	median	4632	3236	5177
	99%	10781	14345	13356
	max	13989	18480	14773

Table A.1: Number of instances per label (shots) and number of tokens per input.

Category	Count	%	Len
Nursing	506,528	73	241
Radiology	338,834	83.3	449
ECG	123,042	61.3	43
Physician	92,426	18.2	1369
Discharge summary	47,572	96.7	2195
Echo	34,064	45.8	464
Respiratory	32,798	8.1	205
Nutrition	7,971	6.4	602
General	7,710	6.4	290
Rehab Services	5,321	4.6	622
Social Work	2,294	2.8	446
Case Management	939	1.3	260
Pharmacy	97	0.1	512
Consult	78	0.1	1206

Table A.2: Statistic of note events documented in MIMIC-III dataset. Each column represents a) the number of notes, b) proportion of visits, c) average number of words for each note type.

Hyperparameter	Value
num param	130m/2.8b
num layer	24/64
dim model	768/2560
context len	16k
num vocab	50277
position emb	None
optimizer	Adam
beta1	0.9
beta2	0.95
epsilon	1e-5
batch size	32
weight decay	0.1
gradient clipping	1.0
peak learning rate	1e-3/6e-4

Table A.3: Hyperparameters used to train ClinicalMamba.