

# A Continued Pretrained LLM Approach for Automatic Medical Note Generation

Dong Yuan\* Eti Rastogi\* Gautam Naik Sree Prasanna Rajagopal

Sagar Goyal Fen Zhao Bharath Chintagunta Jeff Ward

*DeepScribe Inc.*

*San Francisco, California, USA*

*{dong, eti, gautam, sree, sagar, fen, jai, jeff}@deepscribe.tech*

## Abstract

LLMs are revolutionizing NLP tasks. However, the use of the most advanced LLMs, such as GPT-4, is often prohibitively expensive for most specialized fields. We introduce HEAL, the first continuously trained 13B LLaMA2-based LLM that is purpose-built for medical conversations and measured on automated scribing. Our results demonstrate that HEAL outperforms GPT-4 and PMC-LLaMA in PubMedQA, with an accuracy of 78.4%. It also achieves parity with GPT-4 in generating medical notes. Remarkably, HEAL surpasses GPT-4 and Med-PaLM 2 in identifying more correct medical concepts and exceeds the performance of human scribes and other comparable models in correctness and completeness.

## 1 Introduction

The emergence of large language model (LLM) has brought revolutionary changes to natural language processing and understanding tasks, paving the way for practical applications of AI across multiple domains such as law, finance, and healthcare. Private LLMs such as GPT-4 (OpenAI, 2023) and Med-PaLM 2 (Singhal et al., 2023) and open-source LLMs like LLaMA2 (Meta, 2023) have shown strong performance on general NLP benchmarks. However, recent studies have shown promise that with continued training on more targeted datasets, e.g. smaller LLMs like Orca (Mukherjee et al., 2023; Mitra et al., 2023) and Phi-2 (Mojan Javeri, 2023), can surpass much larger LLMs on general tasks. Despite the success of LLM in general capabilities, they often fall short in niche domains like healthcare, where precision and profound understanding are crucial. Hence, several models such as Meditron-70B (Chen et al., 2023b), PMC-LLaMA (Wu et al., 2023) have emerged.

Transcribing medical conversations is a challenging task for both humans and machines due to po-

tential transcription errors and the innate complexity of spoken language, an issue unaddressed by existing medical LLMs. Existing LLMs trained on medical data largely do well on problems like medical Q&A but struggle to produce a comprehensive EHR-compatible medical note. Some domain-adapted LLMs (Van Veen et al., 2023) can write some components of the note, but they leave out the crucial "Subjective" section. Some fine-tuned models (Zhang et al., 2021) can generate notes from medical conversations but need human overview.

Overall, we developed a new medical LLM proficient in interpreting medical conversation. By using techniques like continued pretraining on diverse data and explanation tuning, including medical and general web corpora, GPT-4 task instructions, EHRs, the model was capable of producing medical SOAP notes approved by physicians.

Our main contributions include:

To the best of our knowledge, we are the first to build a small-size (13B) medical LLM that can produce medical notes without any human intervention from doctor-patient conversations that bypass human quality and are accepted by physicians.

HEAL surpasses Med-PaLM 2 and other publicly available models of the same size, matches GPT-4's performance in medical notes generation, and excels with the highest completeness.

Despite having a smaller model size, we achieved an accuracy of 78.4% on PubMedQA, outperforming GPT-4 and within 5% of Med-PaLM 2's performance.

## 2 Continued Pretraining

### 2.1 Dataset

We collected our training data from three major sources to enable the model to generate coherent English sentences, comprehend medical content, and execute complex instructions required for generating medical notes. (see Table 1)

<sup>1</sup>\*Core Contributors and Corresponding Authors

Dataset	Number of tokens (in billions)	Percentage of total data
Non-medical public	5.33	35.79
Medical public	5.68	38.14
Medical proprietary	3.88	26.07
<b>Total</b>	<b>14.89</b>	<b>100.00</b>

Table 1: Pretraining datasets.

**Non-medical public datasets.** To ensure that the new model doesn’t lose the generative capabilities of the pretrained LLaMA2 model, we added general domain datasets such as C4 (Raffel et al., 2019). Continued pretraining on them was crucial for generational tasks, enhancing the model’s grammar and phrase composition skills. Initially, we also included filtered subtitle data from open-subtitle and youtube. However, we decided to exclude these datasets due to their poor quality negatively impacting the model’s performance.

**Medical public datasets.** We filtered data from medical web domains such as nih.gov to cover different aspects of medical concept understanding and replay medical knowledge to the model, so the model won’t forget the medical knowledge after continued training. MedDialog (Chen et al., 2020) taught medical language conversation while reading materials such as PubMed articles (Gao et al., 2020) provided the model with an overall medical context. PubMed and filtered web medical corpus were two major sources, each contributed around 2.5B tokens each in the final training dataset.

**Proprietary medical datasets.** We also curated a deidentified proprietary medical dataset that consists of real-world doctor-patient conversations from the United States, Electronic Health Records (EHR), SOAP (Subjective, Objective, Assessment, and Plan) notes, and ROS (Review of System) templates. We also created a synthetic dataset comprising of medical instructions, like extraction of medications from a medical conversation and grammar correction of a generated medical note, respectively. These instructions were generated with the help of both humans and GPT-3.5/GPT-4. For some of the instructions, we also included detailed explanation as shown in (Mukherjee et al., 2023). Training on such instructions with explanations, helped the model better comprehend the medical notes and understand the reasoning behind it, which was especially needed for the downstream medical documentation task. For example, we created a medical instruction that asks the model to retrieve information from a conversation as shown below:

You specialize in summarizing medical conversations, providing clear and thorough explanations so that people can trust your summary with evidence. I have part of a transcript from a conversation between my doctor and myself.

Task: Summarize the *<targeted content>* from this conversation.

Requirements: *<requirements>*

Transcript: *<transcript>*

Then we further created instructions about reviewing the generated note:

Your job is to review a given medical note and generate an updated note.

Rules: *<rules on how to review>*.

List all the needed updates for the medical note as Updates. Return the updated medical note as Updated Medical Note.

Transcript: *<transcript>*

Medical Note: *<medical note>*

Finally, both of them were used for training the model to improve the model’s understanding of the summarization task.

While we developed a much larger high-quality custom dataset including more than 60B tokens, currently only 14.89B tokens were used for this training exercise.

## 2.2 Training Details

We performed training using FSDP (Zhao et al., 2023) pipeline parallelism with hybrid sharding and flash attention 2 on 32 A100 80 GB GPUs. We continued training LLaMA2 13B using learning rate of  $5e-5$  which decays to  $1e-5$  following a cosine schedule. We chose a relatively small batch size of 256, to achieve more than 10K effective gradient update steps. A medical conversation can exceed 30 minutes and surpass 4K in context length. Therefore, we used 8K context length by applying positional interpolation (Chen et al., 2023a) to the base model. We set the weight decay at 0.1 and a warm-up step count to 50.

**Robust Training.** To be tolerant of machine and experiment related mishaps, we used fixed seed, checkpoints, and implemented phased training where we divided the training data into  $n$  subsets. If the loss of a particular validation subset started to stabilize, we reduced the sampling rate in the next phase for efficiency.

**Data Packing & Dedup.** We packed data by sentence to fit into max sequence length. We also

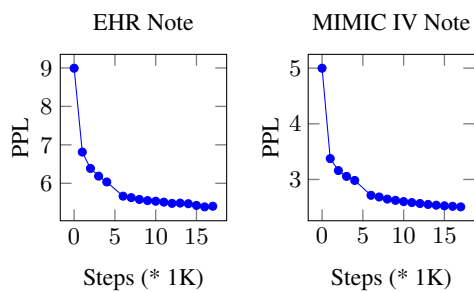


Figure 1: Pretraining validation perplexity.

deduplicated our data to improve data quality (Lee et al., 2021).

**Loss.** For the general corpus including C4, public medical materials, we calculated the gradient on every token. However, on proprietary instruction data, the loss was only calculated on response tokens like (Mukherjee et al., 2023).

### 3 Evaluations

This section shows some of our continued pretraining results and evaluation methodology.

#### 3.1 Pretraining

We employed two evaluation methods to monitor pertaining. Firstly, we measured the perplexity across all the data sources. We used a validation set to track how efficiently the model learns from each source. Figure 1 is a subset of evaluations on EHR and MIMIC IV Note. EHR Note is 1K notes sampled from our proprietary dataset, which is the doctors’ written notes from real clinic visits. MIMIC IV Note is 1K sampled deidentified critical care notes from the public dataset (Johnson et al., 2020). The Figure 1 shows that as the training continues, the model progressively increases its understanding of both data sets. However, MIMIC IV has a much lower perplexity suggesting that the base LLaMA2 model might have been trained on this dataset during the initial pertaining process.

Secondly, for a holistic understanding of the generation quality, we used several few-shot (3-shot) generative tasks for validation, that included:

**1) Long text generation:** This task is associated with summarizing different categories of the subjective section of SOAP notes from medical transcripts between doctor and patient. For example:

**Prompt** Summarize the patient’s *chief complaint* from the given text.  
**Transcript:** <transcript>  
**Output** <response>

Training data	ROS (multi-choice) (Acc %)	Long Text Rouge-1 (f1 %)	Long Text Rouge-cls (f1 %)
1B Total	<b>47.36</b>	<b>44.81</b>	41.53
MED	37.85	39.44	35.91
PUB	36.81	<b>44.49</b>	<b>42.35</b>

Table 2: **Training data ablation results.** The **MED** dataset is derived from the 1B training dataset by excluding all the public datasets. Similarly, the **PUB** dataset is produced by removing all medical datasets.

**2) Medium text generation:** This is a question answering task on medical transcript. We curated this data by modifying the Alpaca (Rohan et al., 2023) pipeline on the collected transcription dataset. We queried GPT-4 to generate questions prompting responses ranging from a few words to a full sentence based on the transcription. For example:

**Prompt** Identify the patient’s current medication.  
**Transcript:** <transcript>  
**Output** <response>

**3) Short text generation:** This comprises of ROS (Review of System) - related classification tasks, including questions about body system identification (multi-choice), and absence or presence of symptoms (single-choice). For example:

**Prompt** Is the patient showing signs of depression, like persistent sadness, lack of interest, or appetite changes?  
**Transcript:** <transcript>  
**Output** <response>

We measured Rouge-cls for tasks 1, 2 and accuracy for task 3, to monitor pretraining performance. Each of evaluation dataset has 1000 examples.

Figure 2 demonstrates that our model’s performance consistently improved in generating long and medium texts, and in multi-choice classification. However, no significant improvement was observed in single-choice classification. We attribute this to the already high accuracy numbers and the fact that further improvement was noted when the model was separately trained on a smaller related dataset, indicating potential enhancements with scaled-up training.

#### 3.2 Pretraining Ablation

Table 2 shows our examination of the effects of varying data proportions using a 1B token dataset, derived from a scaled-down version of our custom

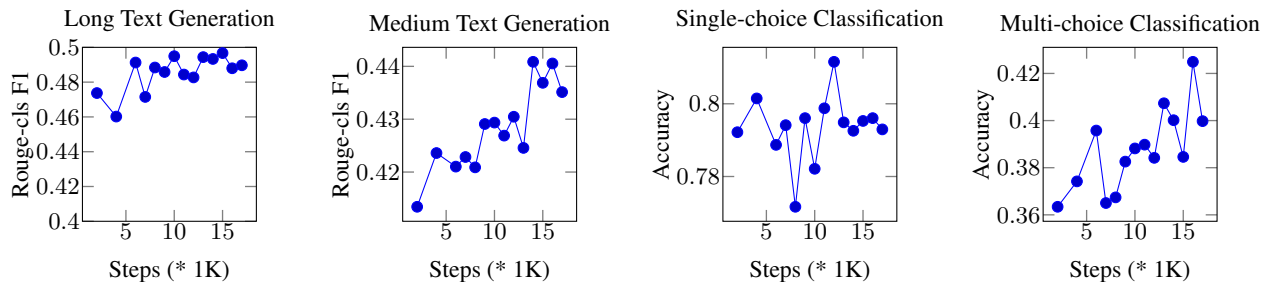


Figure 2: Pretraining validation generation capability monitoring.

Model	#Incorrect	#Irrelevant	#Missed
Human	1.20	0	11.20
GPT-4	<b>0.80</b>	0.20	6.75
Med-PaLM 2	1.36	0	10.50
GPT-3.5	2.00	1.71	8.50
†LLaMA2-chat-13B	4.14	4.71	11.21
†PMC-LLaMA-13B	1.57	0.43	15.14
*LLaMA2-13B	1.50	0.14	9.86
*MedLLaMA-13B	2.07	0.71	11.57
*Meditron-7B	3.00	0.57	10.64
<b>HEAL</b>	<b>0.85</b>	<b>0.30</b>	<b>4.30</b>

Table 3: **Average entity errors comparison.** Both \* and † are fine-tuned models. \* indicates a pretrained model was used as the base, † denotes a fine-tuned instruction model was used as the base.

15B dataset on the 7B LLaMA2 model. The ablation study revealed that removing general datasets from the mix detrimentally impacted the model’s generative abilities, resulting in decreased summarization quality. We were also able to conclude that the medical datasets indeed improved the model’s understanding of the medical context. Consequently, we decided to use equal proportions of these datasets during training to maintain the model’s generative abilities while improving its understanding of medical contexts.

### 3.3 Medical Note Generation

**Evaluation Dataset and Setup.** We compared the HEAL model to several general and medical SOTA models, including the high-end GPT-4, GPT-3.5, and Med-PaLM 2 (Singhal et al., 2023) and other similarly sized open-source medical LLMs, as shown in Table 3. We meticulously fine-tuned LLaMA2-Chat-13B (Meta, 2023) and the PMC-LLaMA-13B (Wu et al., 2023) on medical generative tasks of varying lengths, detailed in Section 3.1 using 10K instruction samples. Pre-trained models like LLaMA2-13B (Meta, 2023), MedLLaMA (base model of PMC-LLaMA), and Meditron-7B (Chen et al., 2023b) were explanation-tuned on our proprietary dataset of 500K examples to enhance their instruction-following capabilities.

We also compared these models to human scribes from our production system (medical students who underwent internal scribe training and received monetary compensation for their services). All the models and scribes were evaluated on generating the Subjective and Plan sections of the SOAP medical note using 10 doctor-patient dialogue-style conversations averaging 12 minutes each.

**Evaluation Metric.** We leveraged human medical experts to evaluate these models. They developed a rubric note for each transcript, highlighting all essential medical information as separate medical entities. Every entity symbolized a significant sentence or phrase that a healthcare provider needed to approve the note. On average, our experts identified 35 medical entities per transcript. We evaluated the generated notes on three key parameters: **Completeness**, **Correctness**, and **Conciseness** as outlined in (Van Veen et al., 2023) using the following metrics:

- 1) **Missed Information** refers to the entities omitted in the test note relative to the rubric note. This metric reflects the test note’s completeness.
- 2) **Incorrect Information** implies the entities inaccurately captured by the test note. This metric is critical in healthcare where information accuracy is essential, as misinformation can erode trust in AI.
- 3) **Irrelevant information** refers to extraneous elements in the test note not linked to the rubric note. As lengthy medical notes require more time for review, it’s crucial to reduce irrelevant information.

**Results and Analysis.** Table 3 compares the performance of our HEAL model, other models, and human scribes. Notably, HEAL surpasses all other models in the Missed Information metric, indicating a superior ability to identify and summarize critical medical information. We attribute this improved performance to our continued pretraining approach using complex medical instructions. We also observed some inaccuracies due to ASR (Automatic Speech Recognition) errors, yet both our model and GPT-4 excelled at correcting these



Dataset	LLaMA2 13B	PMC- LLaMA 13B	GPT-4 (5-shot)	Med- PaLM 2 (best)	HEAL 13B
PubMedQA	76.40	77.90	75.2	<b>81.8</b>	78.4
MedQA	45.48	56.36	81.4	<b>86.5</b>	47.2

Table 4: Accuracy (%) on PubMedQA and MedQA.

mistakes. Human scribes and Med-PaLM 2 created concise notes but missed vital medical details. Other models, such as GPT 3.5, MedLLaMa, and LLaMA2-chat, struggled to grasp real-world conversation nuances, as shown by their high Incorrect and Missed Information scores. Overall, our model shows exceptional performance in all metrics of the task, outperforming both human scribes and other fine-tuned models.

In our detailed quality evaluation, we found that a human scribe takes about 1.67 times longer than the audio recording to create a medical note. However, AI models can generate the same note almost instantly, demonstrating the efficiency and time-saving capabilities of AI in medical transcription.

### 3.4 Public Benchmark

Although HEAL is specifically designed for medical note summarization, we also tested its performance against other LLMs on two popular medical benchmarks to evaluate its efficiency in other medical tasks.

*PubMedQA* (Jin et al., 2019) A biomedical QA task to answer research questions with yes/no/maybe using the corresponding PubMed paper snippets.

*MedQA* (Jin et al., 2021) Multi-choice questions extracted from US Medical License Exams.

In PubMedQA, Med-PaLM 2 with the best prompting strategy (Singhal et al., 2023) took advantage of its huge size and further tuning on PubMedQA data to achieve the highest score. As shown in Table 4, HEAL achieved 78.4% accuracy after tuning, which surpasses GPT-4’s performance (Nori et al., 2023), fine-tuned LLaMA2 and even PMC-LLaMA (Wu et al., 2023) which is further tuned on 75B PubMed data. Our improved performance can be attributed to our proprietary medical instruction data on conversational data which focuses more on medical understanding.

In MedQA, we attained a 47.2% accuracy rate, surpassing the LLaMA2 13B model yet falling short of PMC-LLaMA. MedQA focuses on medical reasoning, requiring the model to recall medical knowledge and derive diagnoses or solutions from specified problems. Larger models like GPT-4, Med-PaLM 2, or those trained with vast amounts of

data hold an inherent advantage in this task. HEAL, which is geared towards interpreting medical conversations, does not align with this task, yielding suboptimal performance on this dataset.

## 4 Conclusion

This paper presents our work of developing a medical LLM capable of comprehending and summarizing medical conversation. As a result, this is the first model, with significantly fewer parameters, to outperform humans, existing medical LLMs including Med-PaLM 2, PMC-LLaMA and perform on par with GPT-4. Our evaluation shows that even small-scale continued pretraining of smaller LLMs can show impressive gains. We believe that scaling up our training can further improve results. Our work presents a promising development in healthcare documentation and other medical areas.

## 5 Related Work

**Medical LLMs.** Various medical LLMs such as MedGPT (Kraljevic et al., 2021), and Med-PaLM 2 (Singhal et al., 2023) show how training on various medical datasets, improves model’s performance on medical knowledge understanding tasks. MEDITRON-70B (Chen et al., 2023b), the state-of-the-art open-source LLM and PMC-LLaMA (Wu et al., 2023) demonstrates the effectiveness of task-specific fine-tuning and instruction tuning.

**Domain adaption LLM.** As demonstrated by (Gururangan et al., 2020), (Beltagy et al., 2019), continued pretraining on unlabeled, domain-specific data boosts model performance on domain tasks, providing a practical solution when resources for scratch domain-adaptive pretraining are limited.

**Medical Note Generation.** Prior work by (Zhang et al., 2021), (Van Veen et al., 2023) demonstrated the feasibility of using Language Models to generate medical summaries from dialogues. However, they primarily aimed at producing partial notes or semi-automated methods requiring human involvement, rather than comprehensive, provider-ready reports.

**Explanation tuning.** Orca (Mukherjee et al., 2023; Mitra et al., 2023) models showcased that smaller Language Models capable of sound reasoning can efficiently perform complex tasks. They were trained by explanation tuning a LLaMA2 13B model (Touvron et al., 2023) using bigger models like GPT4 as a teacher.

## 6 Ethical Considerations

All the data processing and experiments were done in HIPAA-compliant environment. We deidentified clinical data to remove any PHI information as per our data compliance agreement. HEAL is only used for internal medical tasks like summarization, transcription based Q&A, and note review. All prompts are audited to prevent unintentional usage.

## 7 Limitations

Our design focuses on contextual comprehension and summarization of transcripts, and can be further improved on MedQA or similar benchmarks with training on more medical data. Future projects could explore utilizing more sophisticated base models, curating higher quality data with a balanced mix of medical knowledge and reasoning content, and scaling up the experiment.

## References

- Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. Scibert: A pretrained language model for scientific text. *arXiv preprint arXiv:1903.10676*.
- Shouyuan Chen, Sherman Wong, Liangjian Chen, and Yuandong Tian. 2023a. Extending context window of large language models via positional interpolation. *arXiv preprint arXiv:2306.15595*.
- Shu Chen, Zeqian Ju, Xiangyu Dong, Hongchao Fang, Sicheng Wang, Yue Yang, Jiaqi Zeng, Ruisi Zhang, Ruoyu Zhang, Meng Zhou, Penghui Zhu, and Pengtao Xie. 2020. Meddialog: a large-scale medical dialogue dataset. *arXiv preprint arXiv:2004.03329*.
- Zeming Chen, Alejandro Hernández Cano, Angelika Romanou, Antoine Bonnet, Kyle Matoba, Francesco Salvi, Matteo Pagliardini, Simin Fan, Andreas Köpf, Amirkeivan Mohtashami, et al. 2023b. Meditron-70b: Scaling medical pretraining for large language models. *arXiv preprint arXiv:2311.16079*.
- Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Ho-race He, Anish Thite, Noa Nabeshima, et al. 2020. The pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*.
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. [Don't stop pretraining: Adapt language models to domains and tasks](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8342–8360, Online. Association for Computational Linguistics.
- Di Jin, Eileen Pan, Nassim Oufattole, Wei-Hung Weng, Hanyi Fang, and Peter Szolovits. 2021. What disease does this patient have? a large-scale open domain question answering dataset from medical exams. *Applied Sciences*, 11(14):6421.
- Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William Cohen, and Xinghua Lu. 2019. Pubmedqa: A dataset for biomedical research question answering. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2567–2577.
- A Johnson, L Bulgarelli, T Pollard, S Horng, LA Celi, and R Mark. 2020. Mimic-iv (version 1.0).
- Zeljko Kraljevic, Anthony Shek, Daniel Bean, Rebecca Bendayan, James Teo, and Richard Dobson. 2021. Medgpt: Medical concept prediction from clinical narratives. *arXiv preprint arXiv:2107.03134*.
- Katherine Lee, Daphne Ippolito, Andrew Nystrom, Chiyuan Zhang, Douglas Eck, Chris Callison-Burch, and Nicholas Carlini. 2021. Deduplicating training data makes language models better. *arXiv preprint arXiv:2107.06499*.
- Meta. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *ArXiv*, abs/2307.09288.
- Arindam Mitra, Luciano Del Corro, Shweti Mahajan, Andres Cudas, Clarisse Simoes, Sahaj Agarwal, Xuxi Chen, Anastasia Razdaibiedina, Erik Jones, Kriti Agarwal, Hamid Palangi, Guoqing Zheng, Corby Rosset, Hamed Khanpour, and Ahmed Awadallah. 2023. [Orca 2: Teaching small language models how to reason](#).
- Sébastien Bubeck Mojan Javaheripi. 2023. [Phi-2: The surprising power of small language models](#).
- Subhabrata Mukherjee, Arindam Mitra, Ganesh Jawahar, Sahaj Agarwal, Hamid Palangi, and Ahmed Awadallah. 2023. [Orca: Progressive learning from complex explanation traces of gpt-4](#).
- 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](#). *ArXiv*, abs/2303.08774.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2019. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *arXiv e-prints*.
- Taori Rohan, Gulrajani Ishaan, Zhang Tianyi, Dubois Yann, Li Xuechen, Guestrin Carlos, Liang Percy, and B. Hashimoto Tatsunori. 2023. [Alpaca: A strong, replicable instruction-following model](#).

- Karan Singhal, Tao Tu, Juraj Gottweis, Rory Sayres, Ellery Wulczyn, Le Hou, Kevin Clark, Stephen Pfohl, Heather Cole-Lewis, Darlene Neal, et al. 2023. Towards expert-level medical question answering with large language models. *arXiv preprint arXiv:2305.09617*.
- 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*.
- Dave Van Veen, Cara Van Uden, Louis Blankemeier, Jean-Benoit Delbrouck, Asad Aali, Christian Bluethgen, Anuj Pareek, Malgorzata Polacin, William Collins, Neera Ahuja, et al. 2023. Clinical text summarization: adapting large language models can outperform human experts. *arXiv preprint arXiv:2309.07430*.
- Chaoyi Wu, Weixiong Lin, Xiaoman Zhang, Ya Zhang, Yanfeng Wang, and Weidi Xie. 2023. Pmc-llama: Towards building open-source language models for medicine. *arXiv preprint arXiv:2305.10415*.
- Longxiang Zhang, Renato Negrinho, Arindam Ghosh, Vasudevan Jagannathan, Hamid Reza Hassanzadeh, Thomas Schaaf, and Matthew R Gormley. 2021. Leveraging pretrained models for automatic summarization of doctor-patient conversations. *arXiv preprint arXiv:2109.12174*.
- Yanli Zhao, Andrew Gu, Rohan Varma, Liang Luo, Chien-Chin Huang, Min Xu, Less Wright, Hamid Shojanazeri, Myle Ott, Sam Shleifer, et al. 2023. Pytorch fsdp: experiences on scaling fully sharded data parallel. *arXiv preprint arXiv:2304.11277*.