



HAL
open science

BioMistral: A Collection of Open-Source Pretrained Large Language Models for Medical Domains

Yanis Labrak, Adrien Bazoge, Emmanuel Morin, Pierre-antoine Gourraud,
Mickaël Rouvier, Richard Dufour

► To cite this version:

Yanis Labrak, Adrien Bazoge, Emmanuel Morin, Pierre-antoine Gourraud, Mickaël Rouvier, et al..
BioMistral: A Collection of Open-Source Pretrained Large Language Models for Medical Domains.
62th Annual Meeting of the Association for Computational Linguistics (ACL'24), Aug 2024, Bangkok,
Thailand. hal-04621178v1

HAL Id: hal-04621178

<https://hal.science/hal-04621178v1>

Submitted on 23 Jun 2024 (v1), last revised 17 Jul 2024 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC0 - Public Domain Dedication 4.0 International License

BioMistral: A Collection of Open-Source Pretrained Large Language Models for Medical Domains

Yanis Labrak^{*1,2} Adrien Bazoge^{*3,4}

Emmanuel Morin⁴ Pierre-Antoine Gourraud³ Mickael Rouvier¹ Richard Dufour^{1,4}

¹LIA, Avignon Université ²Zenidoc

³Nantes Université, CHU Nantes, Clinique des données, INSERM, CIC 1413, F-44000 Nantes, France

⁴Nantes Université, École Centrale Nantes, CNRS, LS2N, UMR 6004, F-44000 Nantes, France

{firstname.lastname}@univ-avignon.fr

{firstname.lastname}@univ-nantes.fr

Abstract

Large Language Models (LLMs) have demonstrated remarkable versatility in recent years, offering potential applications across specialized domains such as healthcare and medicine. Despite the availability of various open-source LLMs tailored for health contexts, adapting general-purpose LLMs to the medical domain presents significant challenges. In this paper, we introduce BioMistral, an open-source LLM tailored for the biomedical domain, utilizing Mistral as its foundation model and further pre-trained on PubMed Central. We conduct a comprehensive evaluation of BioMistral on a benchmark comprising 10 established medical question-answering (QA) tasks in English. We also explore lightweight models obtained through quantization and model merging approaches. Our results demonstrate BioMistral’s superior performance compared to existing open-source medical models and its competitive edge against proprietary counterparts. Finally, to address the limited availability of data beyond English and to assess the multilingual generalization of medical LLMs, we automatically translated and evaluated this benchmark into 7 other languages. This marks the first large-scale multilingual evaluation of LLMs in the medical domain. Datasets, multilingual evaluation benchmarks, scripts, and all the models obtained during our experiments are freely released.

1 Introduction

In the rapidly evolving landscape of Natural Language Processing (NLP), generative Large Language Models (LLMs) like ChatGPT (OpenAI, 2023) and Vicuna (Zheng et al., 2023) have revolutionized human-computer interactions, demonstrating remarkable versatility and advanced capabilities across various tasks and domains. These models exhibit human-like comprehension and reasoning, enabling them to tackle basic textual understanding as well as complex problem-solving tasks. The emergence of open-source LLMs such as

BLOOM (Workshop et al., 2023) and LLaMA (Touvron et al., 2023a) underscores the transformative potential of these models, facilitating their innovative use in specialized domains including medicine (Dave et al., 2023).

However, integrating LLMs into healthcare and medicine presents unique challenges and opportunities (He et al., 2023; Zhou et al., 2024). While preliminary adoption has opened new avenues for innovation, concerns about data privacy risks associated with proprietary models like MedPaLM-2 (Singhal et al., 2023b) and GPT-4 (Nori et al., 2023a) have arisen. The community’s interest in specialized LLMs for healthcare has led to initiatives like PMC-LLaMA (Wu et al., 2023) and MedAlpaca (Han et al., 2023). However, the adoption of open-source medical models has been limited, primarily due to the lack of lightweight models allowing commercial use with performance comparable to larger or proprietary models. To address this gap, there is a need to develop specialized models based on open-source foundation ones like GPT-NeoX (Black et al., 2022), LLaMa 2 (Touvron et al., 2023b), or Mistral (Jiang et al., 2023), and to optimize them for use on consumer-grade devices while maintaining performance.

In this work, we present BioMistral 7B, a specialized LLM tailored for the biomedical domain, derived from Mistral 7B Instruct v0.1 (Jiang et al., 2023) and further pre-trained on PubMed Central. Our contributions include:

1. The construction of BioMistral 7B, the first open-source Mistral-based model for the biomedical domain, including the analysis of different evaluation strategies such as few-shot in-context learning and supervised fine-tuning.
2. An original study with the introduction of a benchmark of 10 medical question-answering (QA) tasks in English, automatically translated into 7 other languages, facilitating assessment against existing state-of-the-art

open-source medical LLMs, and shedding light on its multilingual potential and robustness across diverse linguistic contexts.

3. A large in-depth quantitative analysis of the model’s truthfulness and calibration in multilingual contexts.
4. A rigorous evaluation of a suite of lightweight models derived through various efficient quantization approaches.
5. An exploration of novel model merging techniques between Mistral 7B Instruct and BioMistral 7B models, allowing leveraging the commonsense reasoning skills of specialized and general-purpose LLMs.

All datasets, the multilingual benchmark, preprocessing scripts, and models are accessible on HuggingFace¹ and GitHub² under Apache 2.0 license.

2 Related Work

Generative LLMs with a generalist purpose, such as GPT-4 (OpenAI et al., 2023) and Gemini (Team et al., 2023), have demonstrated exceptional performance across various NLP tasks, including QA, text summarization, and language translation in zero- and few-shot scenarios. While these models firstly remain proprietary, the open-source movement has played a crucial role in democratizing LLMs, offering powerful capabilities to the broader community with Llama 2 (Touvron et al., 2023b), Vicuna (Zheng et al., 2023), Falcon (Almazrouei et al., 2023) and Mistral (Jiang et al., 2023) models.

The adaptation of LLMs to specialized domains often involves encoding domain-specific knowledge to fully leverage the models’ capabilities. This adaptation process has led to the development of several open-source models from pre-trained general domain models. Encoder-based models, such as BERT (Devlin et al., 2019), and encoder-decoder, such as T5 (Raffel et al., 2020), have been adapted for the medical domain (Lee et al., 2019; Gu et al., 2021; Yuan et al., 2022; Lu et al., 2022) but faced challenges with QA tasks. More recently, multiple LLMs based on decoder-only architecture have been developed for the medical domain, such as BioGPT (Luo et al., 2022) based on GPT-2 (Radford et al., 2019), ClinicalGPT (Wang et al., 2023), based on BLOOM-7B (Workshop et al., 2023), PMC-LLaMA (Wu et al., 2023), MedAlpaca (Han et al., 2023) and Chat-Doctor (Li

et al., 2023), based on LLaMA (Touvron et al., 2023a), and BioMedGPT-LM-7B (Zhang et al., 2024) and MediTron-7B (Chen et al., 2023), adapted from Llama-2 (Touvron et al., 2023b). In contrast, proprietary medical LLMs like GPT-4 MedPrompt (Nori et al., 2023b) and Med-PALM 2 (Singhal et al., 2023b) face usability issues similar to general-purpose models. Nonetheless, they can be used as benchmark "target goals" for the development of domain-specific open-source models.

3 BioMistral

In this section, we present the modules that facilitated the construction of BioMistral 7B. We first develop our training corpus (Section 3.1) used during further pre-training. We then present the model adaptation method (Section 3.2). Finally, we discuss the approaches for model merging (Section 3.3) and expose the employed quantization strategies (Section 3.4).

3.1 Pre-training Dataset

For LLM adaptation to the medical domain, we selected the PMC Open Access Subset³ for its comprehensive and freely accessible collection of medical research papers. This choice is guided by the success demonstrated by PMC-LLaMA (Wu et al., 2023), PubMedBERT (Gu et al., 2021), and SciFive (Phan et al., 2021), which have showcased significant enhancements in language modeling for medical applications. Our focus lies on the Commercial Use Allowed subset, encompassing documents licensed under various Creative Commons licenses (CC0, CC BY, CC BY-SA, and CC BY-ND). This subset ensures the reusability of our model’s outputs, even for commercial purposes.

In the preprocessing phase, we aim to optimize the dataset for training efficiency while considering hardware limitations. Our pre-training objective involves further pre-training Mistral on a subsample of this corpus, targeting 1.5 epochs within the 20-hour limit of Jean Zay HPC. This decision aligns with insights from the Zephyr model (Tunstall et al., 2023), which suggests that observing 1.5 times the corpus adequately enhances model performance, with marginal benefits beyond this threshold. We then meticulously selected 3 billion tokens from this pre-processed PubMed Central corpus, corresponding to roughly 1.47 million documents. The dataset comprises primarily English documents (98.75% of the corpus), with the remain-

¹<https://huggingface.co/BioMistral>

²<https://github.com/BioMistral/BioMistral>

³<https://www.ncbi.nlm.nih.gov/pmc/tools/openftlist/>

ing portion encompassing 9 languages, including Dutch, German, French, and others. Our strategy emphasizes a multilingual dataset approach by prioritizing non-English documents, supplemented with English texts, to ensure a diverse and representative training dataset to meet our 3 billion token target. The raw textual documents undergo pre-processing using the Mistral tokenizer, which includes tokenization and normalization processes.

3.2 Model Adaptation

Training details We leverage Mistral 7B Instruct v0.1 (Jiang et al., 2023) as the base model for adaptation due to its design tailored for incorporating instructions in prompts and its capacity for fine-tuning across diverse tasks using limited datasets. Pre-training settings for BioMistral 7B largely align with Mistral 7B Instruct v0.1. For optimization, we employ the AdamW (Loshchilov and Hutter, 2019) optimizer alongside a cosine learning rate scheduler. Our model architecture inherits the standard transformer architecture from Mistral, including features such as Grouped-Query Attention (Ainslie et al., 2023), Sliding Window Attention (Beltagy et al., 2020) and Rolling Buffer Cache. We maintain an input context length of 2,048 tokens across all models, including the quantized versions (see Section 3.4), in conjunction with FlashAttention-2 (Dao, 2023). For optimization, we set the learning rate to 2×10^{-5} with no warmup, a weight decay of 0.01, a gradient accumulation of 2, and a batch size of 16 on the Jean-Zay HPC with 32 NVIDIA A100 80GB GPUs. This configuration allows for a total batch size of 1,024. Due to the model and the AdamW optimizer’s inability to fit on a single GPU with BF16 precision, we employ the Fully Sharded Data Parallel distributed learning framework (Zhao et al., 2023).

Improving batching To enhance pre-training efficiency, we introduce a post-tokenization grouping method. This method aggregates variable-sized sequences marked by an end-of-sequence token ($\langle /s \rangle$) to fill the model’s 2,048-token sequences without padding. This reduces the sequence count by 87.88%, subsequently accelerating epoch times. Refer to Appendix C for pseudo-code detailing the grouping method.

3.3 Model Merging

Pre-trained models may lose effectiveness when applied beyond their specific domains (Labrak et al., 2023). Traditionally, separate models were used

for each application (Guo et al., 2018), increasing complexity and costs. Recent studies suggest merging pre-trained models to enhance performance and out-of-domain generalization (Cha et al., 2021; Arpit et al., 2022; Wortsman et al., 2022; Jin et al., 2023; Ilharco et al., 2023). Merging involves combining multiple model parameters without additional training. Methods include averaging model weights or considering permutation invariance (Ilharco et al., 2022; Choshen et al., 2022; Singh and Jaggi, 2020; Ainsworth et al., 2023).

Among these methods, we can cite TIES (Yadav et al., 2023), DARE (Yu et al., 2024), and SLERP (Shoemaker, 1985). SLERP merges two models using Spherical Linear Interpolation to allow a smoother transition between model parameters while preventing the significant information loss often encountered with direct averaging of model weights. TIES merges models by creating "task vectors" from each model, isolating unique contributions by subtracting an ancestor base model (e.g., Mistral 7B Instruct). These vectors are then averaged with the base model. Its key improvement over previous methods relies on reducing model interference using sparse vectors and a sign consensus method. DARE enhances TIES by reducing delta parameter redundancy, mainly setting them to zero through random pruning and rescaling while maintaining or improving original model performance.

Exploring model merging in the biomedical domain is particularly interesting since merging a general domain model with a domain-specific one could enhance specialized model adaptability and accuracy across a broader range of applications. The objective of this application in the medical domain is not only to improve general-domain capabilities but also to explore the possibility of emergent reasoning and surpassing the performance of baseline models used for merging.

3.4 Quantization

Quantization techniques are pivotal in democratizing LLMs as they enable the execution of LLMs on smaller devices by minimizing memory requirements. In our study, we investigate two core techniques: Activation-aware Weight Quantization (AWQ) and BitsandBytes (BnB). AWQ (Lin et al., 2023) capitalizes on the insight that weights vary in importance, allowing us to skip quantizing critical weights to mitigate performance degradation. Conversely, BnB quantization assigns a fixed precision of 4 or 8 bits to the entire model.

	MMLU						MedQA	PubMedQA	MedMCQA
	Clinical KG	Medical Genetics	Anatomy	Pro Medicine	College Biology	College Medicine			
Answer options	A / B / C / D	A / B / C / D	A / B / C / D	A / B / C / D	A / B / C / D	A / B / C / D	A / B / C / D / (E)	Yes / No / Maybe	A / B / C / D
Train / Valid. / Test	0 / 0 / 265	0 / 0 / 100	0 / 0 / 135	0 / 0 / 272	0 / 0 / 144	0 / 0 / 173	10178 / 1272 / 1273	211269 / 500 / 500	146257 / 36565 / 4183
Words / Questions	11.09	12.34	13.65	105.46	22.40	48.84	118.16	13.08	14.05
Context	×	×	×	×	×	×	×	✓	×

Table 1: Description of the benchmarked question-answering tasks, including the number of train, validation, and test questions, and answer options for each task. Only PubMedQA incorporates context information within the prompt (see Appendix F). The reference to "Clinical KG" denotes "Clinical Knowledge".

4 Evaluation Protocol

To assess the performance of BioMistral 7B models, we first describe our benchmark of English medical reasoning tasks (Section 4.1) and their multilingual translation (Section 4.2), before presenting the instruction prompting (Section 4.3) and the supervised fine-tuning strategy (Section 4.4) employed for the models’ evaluation.

4.1 Downstream Tasks

To evaluate the performance of the BioMistral 7B model, we selected 10 QA tasks in English from 4 prominent medical corpora covering various specialties, including genetics, anatomy, and clinical cases. These datasets encapsulate real-world scenarios encountered by medical professionals, medical school entrance examination formats, and comprehension tests based on PubMed content. The datasets’ characteristics are provided in Table 1.

MMLU (Hendrycks et al., 2021) comprises exam questions spanning 57 subjects, including philosophy, management, and medical domains. We focused on the 6 subjects relevant to medical and clinical knowledge, and already used to evaluate MedPaLM-2 (Singhal et al., 2023b): college biology, college medicine, anatomy, professional medicine, medical genetics, and clinical knowledge. These subjects were amalgamated to form a consolidated medical-related benchmark, featuring 1,089 questions. As MMLU lacks training data, we fine-tuned our models on MedQA and evaluated their generalization performance on MMLU.

MedQA (Jin et al., 2020) presents questions in the format of the US Medical License Exam (USMLE) and encompasses diverse medical knowledge, including patient profiles, disease symptoms, and drug dosage requirements. The training set comprises 10,178 samples, while the test set contains 1,273 questions. MedQA provides two configurations: four-choice (MedQA) and five-choice (MedQA 5-options) question formats.

MedMCQA (Pal et al., 2022) consists of over 193k MCQA with 4 options each, extracted from Indian medical entrance examinations (AI-IMS/NEET). It covers 2,400 healthcare topics across 21 medical subjects. The training set comprises 183k samples and the validation set includes 4,183 questions. Due to the unavailability of answer keys for the test set of 6,150 questions, we adopt a similar approach to Wu et al. (2023), using the validation set for evaluation. For hyperparameter tuning, we perform a random split of the training set into new train/validation splits with 146k and 37k samples each.

PubMedQA (Jin et al., 2019) comprises 211k artificially generated multiple-choice question samples and 1,000 samples labeled by experts. In our evaluation, we adhere to the *reasoning-required* setting, in which the model has to predict *yes*, *no*, or *maybe* for a given PubMed abstract used as context and a corresponding question. Fine-tuning is performed using 211k artificially labeled samples and performance is accessed on 500 expert-labeled samples during validation and 500 during test as split in BigBio (Fries et al., 2022) and following Chen et al. (2023); Singhal et al. (2023a) protocol.

4.2 Multilingual Evaluation

While the biomedical language models have been extensively evaluated in languages such as English (Lee et al., 2019; Chen et al., 2023), Chinese (Cai et al., 2021; Yang et al., 2023), French (Touchent et al., 2023; Labrak et al., 2023) or Spanish (Carrino et al., 2022), their performance in languages beyond their own remains relatively understudied. This limited multilingual evaluation can be attributed to the scarcity of biomedical tasks available in languages other than English. To address this gap, we conducted a multilingual evaluation using GPT-3.5 Turbo (version 1106) automatic translation via the OpenAI API. We translated our benchmark into 7 languages: Spanish, German, Portuguese, Russian, French, Arabic, and Chinese. Despite the challenges posed by automatic transla-

	MMLU										Avg.
	Clinical KG	Medical Genetics	Anatomy	Pro Medicine	College Biology	College Medicine	MedQA	MedQA 5 opts	PubMedQA	MedMCQA	
BioMistral 7B	60.9 ^{+1.5}	61.7 ^{+2.1}	<u>49.6</u> ^{+1.2}	55.1 ^{+1.3}	56.9 ^{+1.0}	55.5 ^{+1.7}	44.4 ^{+0.2}	37.4 ^{+0.4}	37.6 ^{+1.5}	43.9 ^{+0.3}	50.3
Mistral 7B Instruct	57.0 ^{+0.8}	56.7 ^{+0.5}	46.9 ^{+0.3}	51.0 ^{+1.1}	58.6 ^{+0.9}	50.1 ^{+1.0}	42.3 ^{+0.3}	34.5 ^{+0.5}	72.2 ^{+0.5}	42.8 ^{+0.5}	51.2
BioMistral 7B Ensemble	<u>62.8</u> ^{+0.5}	<u>62.7</u> ^{+1.7}	46.9 ^{+0.3}	<u>57.0</u> ^{+0.6}	60.6 ^{+0.9}	56.3 ^{+0.3}	44.7 ^{+0.4}	37.1 ^{+0.6}	68.0 ^{+0.4}	44.8 ^{+0.3}	54.1
BioMistral 7B DARE	61.3 ^{+0.4}	61.0 ^{+2.8}	49.9 ^{+0.9}	55.3 ^{+0.7}	64.4 ^{+0.9}	53.9 ^{+1.4}	47.0 ^{+0.5}	<u>38.8</u> ^{+0.7}	<u>70.0</u> ^{+0.7}	<u>44.9</u> ^{+0.2}	<u>54.6</u>
BioMistral 7B TIES	62.3 ^{+0.5}	61.3 ^{+1.9}	48.1 ^{+2.2}	55.8 ^{+0.8}	57.2 ^{+0.7}	<u>56.5</u> ^{+1.5}	44.0 ^{+0.4}	37.7 ^{+0.4}	44.3 ^{+0.8}	44.0 ^{+0.3}	51.1
BioMistral 7B SLERP	63.1 ^{+1.6}	63.3 ^{+0.9}	49.9 ^{+1.9}	<u>57.4</u> ^{+0.3}	63.4 ^{+0.9}	57.8 ^{+0.9}	<u>46.6</u> ^{+0.2}	38.9 ^{+0.4}	68.1 ^{+1.4}	45.7 ^{+0.7}	55.4
MedAlpaca 7B	49.1 ^{+1.3}	49.0 ^{+5.7}	48.4 ^{+1.9}	63.8 ^{+0.8}	47.2 ^{+0.6}	43.5 ^{+1.8}	35.4 ^{+0.3}	30.4 ^{+0.6}	56.0 ^{+0.9}	31.2 ^{+0.2}	45.4
PMC-LLaMA 7B	25.3 ^{+1.5}	26.0 ^{+3.7}	31.9 ^{+1.8}	16.9 ^{+0.5}	28.0 ^{+2.4}	24.9 ^{+1.2}	27.6 ^{+0.8}	21.1 ^{+0.8}	53.3 ^{+0.6}	23.5 ^{+0.3}	27.8
MedTron-7B	37.9 ^{+1.5}	47.0 ^{+3.7}	39.3 ^{+1.6}	34.2 ^{+1.0}	42.6 ^{+1.4}	30.4 ^{+0.7}	34.8 ^{+0.6}	26.3 ^{+0.5}	55.9 ^{+1.0}	33.6 ^{+0.2}	38.2
BioMedGPT-LM-7B	50.1 ^{+1.0}	52.0 ^{+0.8}	46.2 ^{+1.8}	47.3 ^{+1.7}	47.9 ^{+2.5}	45.5 ^{+0.7}	39.3 ^{+1.2}	34.9 ^{+0.4}	58.6 ^{+0.3}	34.9 ^{+0.5}	45.7
GPT-3.5 Turbo 1106	74.71 ^{+0.3}	74.00 ^{+2.2}	65.92 ^{+0.6}	72.79 ^{+1.6}	72.91 ^{+1.7}	64.73 ^{+2.9}	57.71 ^{+0.3}	50.82 ^{+0.7}	72.66 ^{+1.0}	53.79 ^{+0.2}	66.0

Table 2: Performance of 3-shot in-context learning. The scores represent accuracy (\uparrow) and are averaged across 3 random seeds. BioMistral 7B Ensemble, DARE, TIES, and SLERP are model merging strategies that combine BioMistral 7B and Mistral 7B Instruct. Best model in bold, and second-best underlined.

tion, these tools have shown remarkable improvement in recent years (Neves et al., 2023), enabling cost-effective multilingual evaluation. The methodology for multilingual evaluation and the prompt template are the same as those used in the 3-shot scenario for English. The only differences lie in the translation of the questions, options, and context, while the examples used for few-shot learning remain unchanged.

4.3 Instruction Prompting

All of our instructions adhere to the guidelines outlined for GPT-4’s medical evaluation, as detailed in Nori et al. (2023a). Each task is presented as an MCQA, with answer options associated with letters (A to D or A to E). For a comprehensive list of the instruction prompts, please refer to Appendix F. During inference, the model predicts the next token based on the input prompt, generating probabilities for each token in the vocabulary. To ensure relevance, the vocabulary is filtered to include only tokens (here, choice letters) corresponding to the expected answer options. This approach prevents the model from generating irrelevant tokens or hallucinations (Liang et al., 2023; Beeching et al., 2023; Chen et al., 2023).

4.4 Supervised Fine-Tuning (SFT)

Supervised Fine-Tuning (SFT) is a crucial step involving fine-tuning the model on annotated data to adapt it to specific tasks. To optimize BioMistral’s performance beyond what is achievable with few-shot learning, we conducted SFT on both BioMistral 7B models and the baseline open-source models, using the training sets specified in Table 1. However, traditional SFT methods can be resource-intensive. To address this challenge, we adopted the QLoRa fine-tuning method (Detmeters et al., 2023) and an 8-bit quantization technique (Detmeters et al.,

2022) as more cost-effective alternatives. Additionally, we implemented the improved batching method discussed in Section 3.2 to reduce fine-tuning time. For detailed hyperparameters used during SFT, please refer to Appendix A.

5 Results and Discussions

In this section, we report, analyze, and discuss the performance of BioMistral 7B models across various dimensions. We begin by examining its performance in a few-shot learning scenario (Section 5.1), followed by an evaluation of the fine-tuning performances (Section 5.2) of BioMistral 7B compared to several baseline models. The effectiveness of BioMistral 7B model merging strategies is then reported (Section 5.3) before exploring its generalization capabilities across several languages (Section 5.4). Additionally, we analyze the performance of BioMistral quantized versions in a few-shot scenario (Section 5.5). Finally, we delve into its reliability by examining its calibration (Section 5.6) and truthfulness (Section 5.7).

5.1 Few-shot Learning

The few-shot learning evaluation involved applying 3-shot in-context learning based on 3 different sets of randomly selected samples from each dataset’s training set. We limited our samples to 3 due to the model’s 2,048-token context window size. None of the models were fine-tuned on the datasets.

In Table 2, we observe that BioMistral 7B outperforms Mistral 7B Instruct on 8 of the 10 tasks, demonstrating the effectiveness of domain adaptation (Chen et al., 2023; Lee et al., 2019). Additionally, BioMistral 7B surpasses all other open-source biomedical baselines on all tasks in this 3-shot scenario. The observed performances may vary depending on the dataset. For example, on MedQA 4 and 5 options, BioMistral 7B shows a

	MMLU										Avg.
	Clinical KG	Medical Genetics	Anatomy	Pro Medicine	College Biology	College Medicine	MedQA	MedQA 5 opts	PubMedQA	MedMCQA	
BioMistral 7B	59.9 ^{+1.2}	64.0 ^{+1.6}	56.5 ^{+1.8}	60.4 ^{+0.5}	59.0 ^{+1.5}	54.7 ^{+1.0}	50.6 ^{+0.3}	42.8 ^{+0.3}	77.5 ^{+0.1}	48.1 ^{+0.2}	57.3
Mistral 7B Instruct	62.9 ^{+0.2}	57.0 ^{+0.8}	55.6 ^{+1.0}	59.4 ^{+0.6}	62.5 ^{+1.0}	<u>57.2</u> ^{+2.1}	42.0 ^{+0.2}	40.9 ^{+0.4}	75.7 ^{+0.4}	46.1 ^{+0.1}	55.9
BioMistral 7B Ensemble	62.8 ^{+0.5}	62.7 ^{+0.5}	<u>57.5</u> ^{+0.3}	63.5 ^{+0.8}	64.3 ^{+1.6}	55.7 ^{+1.5}	50.6 ^{+0.3}	43.6 ^{+0.5}	77.5 ^{+0.2}	48.8 ^{+0.0}	58.7
BioMistral 7B DARE	62.3 ^{+1.3}	67.0 ^{+1.6}	55.8 ^{+0.9}	61.4 ^{+0.3}	66.9 ^{+2.3}	58.0 ^{+0.5}	51.1 ^{+0.3}	45.2 ^{+0.3}	<u>77.7</u> ^{+0.1}	<u>48.7</u> ^{+0.1}	59.4
BioMistral 7B TIES	60.1 ^{+0.9}	<u>65.0</u> ^{+2.4}	58.5 ^{+1.0}	60.5 ^{+1.1}	60.4 ^{+1.5}	56.5 ^{+1.9}	49.5 ^{+0.1}	43.2 ^{+0.1}	77.5 ^{+0.2}	48.1 ^{+0.1}	57.9
BioMistral 7B SLERP	62.5 ^{+0.6}	64.7 ^{+1.7}	55.8 ^{+0.3}	<u>62.7</u> ^{+0.3}	64.8 ^{+0.9}	56.3 ^{+1.0}	<u>50.8</u> ^{+0.6}	<u>44.3</u> ^{+0.4}	77.8 ^{+0.0}	48.6 ^{+0.1}	<u>58.8</u>
MedAlpaca 7B	53.1 ^{+0.9}	58.0 ^{+2.2}	54.1 ^{+1.6}	58.8 ^{+0.3}	58.1 ^{+1.3}	48.6 ^{+0.5}	40.1 ^{+0.4}	33.7 ^{+0.7}	73.6 ^{+0.3}	37.0 ^{+0.3}	51.5
PMC-LLaMA 7B	24.5 ^{+1.7}	27.7 ^{+1.7}	35.3 ^{+0.7}	17.4 ^{+1.7}	30.3 ^{+0.9}	23.3 ^{+1.7}	25.5 ^{+0.9}	20.2 ^{+0.1}	72.9 ^{+1.2}	26.6 ^{+0.1}	30.4
MediTron-7B	41.6 ^{+1.2}	50.3 ^{+2.1}	46.4 ^{+0.9}	27.9 ^{+0.3}	44.4 ^{+2.6}	30.8 ^{+0.7}	41.6 ^{+0.5}	28.1 ^{+0.5}	74.9 ^{+0.1}	41.3 ^{+0.2}	42.7
BioMedGPT-LM-7B	51.4 ^{+0.4}	52.0 ^{+1.4}	49.4 ^{+2.7}	53.3 ^{+0.6}	50.7 ^{+0.0}	49.1 ^{+0.8}	42.5 ^{+0.3}	33.9 ^{+0.5}	76.8 ^{+0.3}	37.6 ^{+0.4}	49.7
GPT-3.5 Turbo 1106*	74.71 ^{+0.3}	74.00 ^{+2.2}	65.92 ^{+0.6}	72.79 ^{+1.6}	72.91 ^{+1.7}	64.73 ^{+2.9}	57.71 ^{+0.3}	50.82 ^{+0.7}	72.66 ^{+1.0}	53.79 ^{+0.2}	66.0

Table 3: Supervised Fine-Tuning (SFT) performance of BioMistral 7B models compared to baselines, measured by accuracy (\uparrow) and averaged across 3 random seeds of 3-shot. DARE, TIES, and SLERP are model merging strategies that combine BioMistral 7B and Mistral 7B Instruct. Best model in bold, and second-best underlined. *GPT-3.5 Turbo performances are reported from the few-shot results in Table 2.

9.6% and 11.1% increase over MediTron-7B and a 9.0% and 7.0% increase over MedAlpaca 7B, respectively. On MMLU, BioMistral 7B improves performance over previous biomedical LLMs at the 7B scale, with an overall average gain of 6.45% over MedAlpaca 7B, 18.05% over MediTron-7B, and 31.12% over PMC-LLaMA 7B. Similarly, on MedMCQA, BioMistral 7B shows a 10.3% increase over MediTron-7B, 12.7% over MedAlpaca 7B, and 20.4% over PMC-LLaMA 7B. However, in the PubMedQA evaluation, BioMistral’s performance experienced a decline, showing at least a 15.7% lower accuracy compared to other models, likely due to hallucinations caused by imbalanced classes. Overall, GPT-3.5 Turbo remains the best model in this 3-shot scenario.

5.2 Supervised Fine-Tuning (SFT)

We present the performance of BioMistral models and related baselines in Table 3, measured in terms of accuracy. Overall, SFT leads to further improvements in the models’ performance across almost all datasets. Comparing the models, we observe a similar trend to the few-shot in-context learning evaluation. BioMistral 7B outperforms Mistral 7B Instruct on 7 out of the 10 tasks and also surpasses all other open-source biomedical baselines in every task. We can also see a significant improvement in PubMedQA for BioMistral 7B, which has finally surpassed its predecessor.

5.3 Model Merging

As detailed in Section 3.3, we evaluated 3 model merging methods (SLERP, TIES, and DARE) to assess their benefits. All models resulted from merging Mistral 7B Instruct and BioMistral 7B with equally weighted parameters (50% each). Two scenarios are studied: (1) few-shot learning (Table 2),

and (2) supervised fine-tuning (Table 3). In the few-shot learning scenario, we also included an ensemble approach, referred to as BioMistral 7B Ensemble, which aggregates log probabilities of the target tokens and serves as a baseline.

Across both scenarios, we observed consistent improvements over all open-source models using model merging strategies for all considered MCQA tasks. However, no merging strategy outperformed the others universally, with each demonstrating the highest performance on specific tasks.

In the few-shot learning scenario (Table 2), BioMistral 7B Ensemble exhibited a notable increase in accuracy, by 3.7% on College Biology and 30.4% on PubMedQA compared to the standalone BioMistral 7B model. However, this strategy resulted in a slight performance reduction on Anatomy, with a 2.7% drop compared to BioMistral 7B. Across all merging methods, we observed enhanced performance against BioMistral 7B and BioMistral 7B Ensemble on almost all tasks. Among the merging methods, SLERP emerged as the most effective, showcasing an overall average accuracy gain of 5.11% over BioMistral 7B. In contrast, DARE and TIES methods yielded average gains of 4.35% and 0.82%, respectively.

In the context of SFT (Table 3), similar observations were made: model merging methods further enhanced BioMistral’s performance, widening the gap with other open-source biomedical baselines. On average, we observed a gain of 2.06% between the best merged model and BioMistral 7B, and 3.48% compared to Mistral 7B Instruct. Baseline models lagged behind, with a 7.9% overall loss for the best model, MedAlpaca 7B. Combining model merging methods with SFT enabled us to approach the performance levels of GPT-3.5 Turbo and sometimes even surpass them on certain datasets like

	MMLU										
	Clinical KG	Medical Genetics	Anatomy	Pro Medicine	College Biology	College Medicine	MedQA	MedQA 5 opts	PubMedQA	MedMCQA	Avg.
BioMistral 7B*	60.9 ^{+1.5}	61.7 ^{+2.1}	49.6 ^{+1.2}	55.1 ^{+1.3}	56.9 ^{+1.0}	55.5 ^{+1.7}	44.4 ^{+0.2}	37.4 ^{+0.4}	37.6 ^{+1.5}	43.9 ^{+0.3}	50.3
AWQ 4bit + GEMV	59.5 ^{+1.2}	61.3 ^{+1.7}	50.6 ^{+2.5}	53.9 ^{+0.7}	56.2 ^{+1.5}	52.6 ^{+1.7}	43.2 ^{+0.8}	36.8 ^{+0.5}	61.7 ^{+0.9}	41.8 ^{+0.2}	51.8 ^{+1.5}
AWQ 4bit + GEMM	59.5 ^{+1.2}	61.3 ^{+1.2}	50.6 ^{+2.5}	53.6 ^{+0.8}	56.2 ^{+1.5}	52.4 ^{+1.5}	43.2 ^{+0.8}	37.0 ^{+0.5}	61.4 ^{+0.9}	41.8 ^{+0.2}	51.7 ^{+1.4}
DARE AWQ GEMM	58.2 ^{+0.2}	60.0 ^{+1.4}	50.4 ^{+0.6}	52.7 ^{+0.6}	60.9 ^{+2.3}	53.4 ^{+0.3}	45.8 ^{+0.5}	39.0 ^{+0.2}	68.3 ^{+0.2}	44.1 ^{+0.2}	53.28
TIES AWQ GEMM	58.5 ^{+0.6}	63.7 ^{+1.2}	46.7 ^{+1.2}	54.3 ^{+1.5}	57.6 ^{+1.1}	52.4 ^{+1.0}	42.6 ^{+0.0}	36.8 ^{+0.4}	48.1 ^{+0.9}	43.2 ^{+0.5}	50.39
SLERP AWQ GEMM	61.8 ^{+1.3}	61.0 ^{+1.6}	50.1 ^{+1.1}	54.8 ^{+0.9}	62.0 ^{+1.7}	58.0 ^{+1.2}	45.8 ^{+0.4}	39.0 ^{+0.6}	69.2 ^{+1.6}	45.1 ^{+0.8}	54.68
BnB 4bit	57.6 ^{+1.1}	58.7 ^{+0.9}	47.2 ^{+0.9}	52.9 ^{+1.3}	53.7 ^{+0.9}	54.3 ^{+1.2}	43.1 ^{+0.2}	36.8 ^{+0.9}	22.4 ^{+0.4}	42.0 ^{+0.1}	46.9 ^{+3.4}
BnB 8bit	61.3 ^{+0.9}	59.0 ^{+1.4}	50.1 ^{+1.9}	54.3 ^{+0.5}	56.9 ^{+1.1}	56.1 ^{+0.5}	43.5 ^{+0.1}	37.4 ^{+0.5}	37.9 ^{+1.3}	43.2 ^{+0.3}	50.0 ^{+0.3}

Table 4: Performance of quantized BioMistral 7B in a 3-shot scenario, measured by accuracy (\uparrow) and averaged across 3 random seeds. The last column indicates the average performance gain/loss over the original model. Note: The scores of the original BioMistral 7B model are reported from Table 2.

PubMedQA, where we observed a 5.14% gain with BioMistral 7B SLERP.

5.4 Multilingual Generalization

We report in Appendix H the detailed few-shot learning performance of all models across the 7 targeted languages. Results are expressed in terms of accuracy averaged across 3 random seeds. Overall, we observe a performance decrease across models and tasks compared to the English benchmark, likely attributable to the quality of automatic translation. Despite this, GPT-3.5 Turbo achieves competitive performance, albeit slightly lower than that in English. We observe that the performance difference between GPT-3.5 Turbo and open-source medical models is similar across languages which could suppose a lack of training data in the targeted language in open-source models and better multilingual capabilities from GPT-3.5 Turbo.

For a given model and task, the performance may vary between languages. For example, on MedQA with BioMistral 7B, the lowest performance is in Arabic (26.3%), while the best is in Spanish (33.7%), representing a delta of 7.4%. Similarly, this trend is observed for GPT-3.5 Turbo with 40.0% accuracy in Chinese and 49.0% in Spanish. Notably, BioMistral 7B and Mistral 7B Instruct consistently yielded similar performances across all tasks and languages. Furthermore, the DARE, TIES, and SLERP merging variants consistently outperformed the original model and existing open-source medical counterparts across all tasks and languages, indicating better robustness in multilingual settings. Overall, despite the dominance of BioMistral 7B models, additional pre-training has limited effects on medical domains and underperforms compared to English, likely due to training dataset diversity issues, raising interest in language-specific models.

5.5 Quantization Techniques

Table 4 provides an overview of the impact of different quantization techniques on BioMistral performance. Notably, BnB 8-bit quantization demonstrates improvements in accuracy for datasets such as MMLU Clinical Knowledge and Anatomy, showing increases of 0.65% and 1.00%, respectively. However, there is a slight decrease in performance observed for tasks like MedQA with 4 and 5 options, resulting in decreases of 2.61% and 1.06% across all models. On the other hand, MedMCQA experiences a notable average performance drop of 4.05% across all quantization methods, while PubMedQA shows a remarkable 24.1% increase in accuracy when employing the AWQ method.

Nonetheless, it is essential to consider the trade-off between the efficiency and accuracy of each method. Despite its high compression rate (see Appendix D) and competitive performance, the AWQ + GEMV model exhibits the slowest inference time, taking 421 seconds to process the MMLU professional medicine test set on an RTX 3090. In contrast, the AWQ + GEMM model achieves an 86.23% faster inference time, completing the same task in 57.96 seconds, albeit with a slight performance loss. Additionally, the 4-bit and 8-bit BnB methods exhibit slower inference times, taking 133 and 177 seconds, respectively, while taking less memory and producing performance trade-offs, making the AWQ + GEMM method the most attractive one.

5.6 Calibration

Ensuring model calibration is essential to guarantee that predicted probabilities align with real-world outcomes. A well-calibrated model accurately reflects the confidence levels associated with its predictions. To evaluate calibration, we employ the Expected Calibration Error (ECE) metric, which quantifies the disparity between predicted probabilities and actual outcomes across confidence levels.

	Expected Calibration Error (↓)						
	Arabic	Chinese	French	German	Portuguese	Russian	Spanish
BioMistral 7B	13.9 2.7%	19.7 -1.6%	13.5 3.3%	15.2 2.8%	15.2 1.4%	15.2 2.4%	14.0 2.7%
Mistral 7B Instruct	16.6	18.1	16.8	18.0	16.6	17.6	16.7
BioMistral 7B DARE	16.9 -0.3%	18.4 -0.3%	16.3 0.5%	16.6 1.4%	17.2 -0.6%	17.5 0.1%	16.5 0.2%
BioMistral 7B TIES	15.7 0.9%	21.8 -3.7%	16.4 0.4%	16.9 1.1%	17.8 -1.2%	16.6 1.0%	16.7 -0.0%
BioMistral 7B SLERP	14.8 1.8%	16.8 1.3%	14.5 2.3%	15.8 2.2%	15.3 1.3%	16.1 1.5%	15.4 1.3%
MedAlpaca 7B	7.8 8.8%	5.4 12.7%	5.2 11.6%	4.8 13.2%	4.3 12.3%	5.5 12.1%	4.7 12.0%
PMC-LLaMa 7B	15.1 1.5%	13.9 4.2%	12.8 4.0%	12.3 5.7%	12.2 4.4%	14.8 2.8%	12.9 3.8%
MediTron-7B	10.5 6.1%	10.0 8.1%	8.2 8.6%	9.7 8.3%	7.2 9.4%	9.1 8.5%	8.2 8.5%
BioMedGPT-LM-7B	5.1 11.5%	4.3 13.8%	4.8 12.0%	4.8 13.2%	5.3 11.3%	4.6 13.0%	4.4 12.3%

Table 5: Average Expected Calibration Error (ECE) across all tasks for each language-model pair, indicating the model’s calibration quality. Lower ECE values indicate better calibration. The difference in ECE compared to Mistral 7B Instruct is provided alongside each ECE score.

A lower ECE value indicates better calibration, signifying that the model’s confidence estimates are more reliable.

$$ECE = \sum_{m=1}^M \frac{|B_m|}{n} |\text{acc}(B_m) - \text{conf}(B_m)|$$

Table 5 presents the calibration and confidence scores for BioMistral 7B and its base model across various languages compared to other open-source medical models. Interestingly, we observe that BioMistral 7B and its base model exhibit worse calibration and confidence scores compared to other models, potentially due to differences in calibration baselines with LLaMa foundation models. Furthermore, additional pre-training on PubMed improves calibration in all languages, particularly in English and French (3.3% ECE gain), with some degradation observed in Chinese (loss of 1.6%). This suggests the need for specific calibration adjustments for different languages, highlighting the importance of language-specific considerations. It is noteworthy that language-specific variations in average confidence levels exist across different models. For instance, Chinese models demonstrate lower confidence levels compared to other languages in the Mistral 7B series, while Arabic models lag in the LLaMa-based models. Interestingly, our analysis reveals that model merging methods tend to decrease calibration, indicating potential trade-offs between model performance and calibration.

5.7 Truthfulness

Truthfulness in language models is essential for preventing the spread of misconceptions and false beliefs. We employ the TruthfulQA benchmark (Lin et al., 2022) to assess truthfulness, which evaluates LLMs’ factual and sensible output across 817 questions and 38 categories, such as finance and politics. For an evaluation of the medical domain,

we focus on health and medicine-related categories. The evaluation consists of two zero-shot prompts: a general assessment prompt and one derived from the MediTron-7B article (see Figure 4).

Table 8 shows that BioMistral 7B outperforms other models across both prompts and demonstrates a 4.0% improvement over GPT-3.5 Turbo.

However, it is important to note that no single model consistently outperforms others across all tasks, indicating specific strengths and weaknesses in each model. Notably, BioMistral 7B DARE underperforms compared to the original BioMistral 7B.

Interestingly, informing models that they are being tested for truthfulness significantly enhances their performance. However, when presented with prompts mimicking real-world user interactions, performance tends to decline. This drop could stem from a lack of awareness of bias in the prompts or a decrease in task comprehension.

Finally, zero-shot prompting poses challenges, particularly for PMC-LLaMA 7B and MediTron-7B models, which struggled to provide correct answers in Science and Psychology categories.

6 Conclusion

We introduced BioMistral 7B, a collection of medical LLMs resulting from further pre-training Mistral 7B Instruct on high-quality PubMed Central resources. BioMistral 7B incorporates quantized and merged model variants and demonstrates state-of-the-art performance on the multilingual medical evaluation benchmark compared to other open-source 7B models.

Our future work aims to assess the generation quality of BioMistral 7B through human evaluation. Additionally, we plan to enhance its multilingual and chat capabilities using supervised fine-tuning and direct preference optimization tech-

niques, building on top of experiments conducted by (Rafailov et al., 2023) and Li et al. (2023). Finally, we intend to improve the calibration and reliability of our model by integrating techniques such as Jeffrey’s divergence (Jeffreys, 1946) or Platt scaling (Platt et al., 1999) during the further pre-training process.

Limitations

This study required substantial computational resources, encompassing approximately 5,000 hours of A100 80GB GPU computation. These resources were utilized for model creation, evaluations, experimentation with different architectures, and debugging. Technical issues related to model configurations and performance also necessitated additional computation time. According to documentation from the Jean Zay supercomputer⁴, the total environmental cost amounted to 1,295,000 Wh or 73.8 kg CO₂eq, based on the carbon intensity of the energy grid as reported in the BLOOM environmental cost study conducted on Jean Zay (Luccioni et al., 2022). The valuation of computing hours of the experiments amounts to approximately 3,600 EUR, based on Genci documentation⁵, or 20,480 USD for AWS on-demand p4d.24xlarge instances. Additionally, the total inference cost for GPT-3.5 Turbo, inherent to the translation and few-shot evaluation process, amounted to 355.47 USD. These costs make reproducing this study challenging when limited financial and material resources are available.

Given the evolving nature of the GPT-3.5 Turbo model, future replication of these experiments may become impractical if the version used is no longer maintained.

While BioMistral 7B is proficient in processing medical terms and concepts close to its training dataset, the model may encounter difficulties with unfamiliar or rare medical procedures or terminology. Furthermore, its reliance on English-language data results in degraded performance in non-English contexts. It can occasionally cause misinterpretation and lead to erroneous predictions.

Our benchmark offers a framework for academic assessment with selected tasks and metrics, but it might not accurately reflect end users’ actual usage patterns or priorities. Designed for research environments, these criteria may overlook various factors that shape real-world user experiences and preferences.

⁴<http://www.idris.fr/media/jean-zay/jean-zay-conso-heure-calcul.pdf>

⁵<https://www.edari.fr/documents/Modalitesdaces.pdf>

Ethics Statement

Users are solely responsible for the content they generate with BioMistral 7B, and there are no mechanisms in place for addressing harmful, bias, and toxic content disclosure. Any modifications of the models will be released under different version numbers to keep track of the original models related to this paper.

While we introduce BioMistral 7B as a model tailored for the medical domain at large, its evaluation was limited to MCQA datasets, which may not reflect its effectiveness outside this scope. Similar to other LLMs, BioMistral 7B may possess inherent risks and biases that have not yet been thoroughly assessed. Additionally, the model’s performance has not been evaluated in real-world clinical settings. Consequently, we recommend using BioMistral 7B strictly as a research tool and advise against deploying it in production environments for natural language generation or any professional health and medical purposes.

Further evaluation of available language models on various domains is required to assess their capability to generate toxic, rude, or hateful content. To achieve this, the use of datasets such as ToxiGen (Hartvigsen et al., 2022) could provide deeper insights into the subject and help understand how to prevent such behavior. Bias can also significantly impact how language models handle given tasks and may perpetuate stereotypical social biases and demographic attributes observed during training. Two datasets that could be utilized to assess such biases are BOLD (Dhamala et al., 2021) for more general contexts and Discrim-Eval (Tamkin et al., 2023) and SHADR (Guevara et al., 2024), which are specialized for the medical domain.

References

- Joshua Ainslie, James Lee-Thorp, Michiel de Jong, Yury Zemlyanskiy, Federico Lebron, and Sumit Sanghai. 2023. *GQA: Training generalized multi-query transformer models from multi-head checkpoints*. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 4895–4901, Singapore. Association for Computational Linguistics.
- Samuel Ainsworth, Jonathan Hayase, and Siddhartha Srinivasa. 2023. *Git re-basin: Merging models modulo permutation symmetries*. In *The Eleventh International Conference on Learning Representations*.
- Ebtesam Almazrouei, Hamza Alobeidli, Abdulaziz Alshamsi, Alessandro Cappelli, Ruxandra Cojocaru,

- Mérouane Debbah, Étienne Goffinet, Daniel Hesslow, Julien Launay, Quentin Malartic, Daniele Mazzotta, Badreddine Noune, Baptiste Pannier, and Guilherme Penedo. 2023. [The falcon series of open language models](#).
- Devansh Arpit, Huan Wang, Yingbo Zhou, and Caiming Xiong. 2022. [Ensemble of averages: Improving model selection and boosting performance in domain generalization](#). In *Advances in Neural Information Processing Systems*.
- Edward Beeching, Sheon Han, Nathan Lambert, Nazneen Rajani, Omar Sanseviero, Lewis Tunstall, and Thomas Wolf. 2023. Open llm leaderboard. *Hugging Face*.
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. [Longformer: The long-document transformer](#).
- Sidney Black, Stella Biderman, Eric Hallahan, Quentin Anthony, Leo Gao, Laurence Golding, Horace He, Connor Leahy, Kyle McDonell, Jason Phang, Michael Pieler, Usvsn Sai Prashanth, Shivanshu Purohit, Laria Reynolds, Jonathan Tow, Ben Wang, and Samuel Weinbach. 2022. [GPT-NeoX-20B: An open-source autoregressive language model](#). In *Proceedings of BigScience Episode #5 – Workshop on Challenges & Perspectives in Creating Large Language Models*, pages 95–136, virtual+Dublin. Association for Computational Linguistics.
- Zerui Cai, Taolin Zhang, Chengyu Wang, and Xiaofeng He. 2021. [Embert: A pre-trained language model for chinese medical text mining](#). In *Web and Big Data*, pages 242–257, Cham. Springer International Publishing.
- Casimiro Pio Carrino, Joan Llop, Marc Pàmies, Asier Gutiérrez-Fandiño, Jordi Armengol-Estapé, Joaquín Silveira-Ocampo, Alfonso Valencia, Aitor Gonzalez-Agirre, and Marta Villegas. 2022. [Pretrained biomedical language models for clinical NLP in Spanish](#). In *Proceedings of the 21st Workshop on Biomedical Language Processing*, pages 193–199, Dublin, Ireland. Association for Computational Linguistics.
- Junbum Cha, Sanghyuk Chun, Kyungjae Lee, Han-Cheol Cho, Seunghyun Park, Yunsung Lee, and Sungrae Park. 2021. [Swad: Domain generalization by seeking flat minima](#). In *Advances in Neural Information Processing Systems*, volume 34, pages 22405–22418. Curran Associates, Inc.
- Zeming Chen, Alejandro Hernández Cano, Angelika Romanou, Antoine Bonnet, Kyle Matoba, Francesco Salvi, Matteo Pagliardini, Simin Fan, Andreas Köpf, Amirkeivan Mohtashami, Alexandre Sallinen, Alireza Sakhaeirad, Vinitra Swamy, Igor Krawczuk, Deniz Bayazit, Axel Marmet, Syrielle Montariol, Mary-Anne Hartley, Martin Jaggi, and Antoine Bosselut. 2023. [Meditron-70b: Scaling medical pre-training for large language models](#).
- Leshem Choshen, Elad Venezian, Noam Slonim, and Yoav Katz. 2022. [Fusing finetuned models for better pretraining](#).
- Tri Dao. 2023. [Flashattention-2: Faster attention with better parallelism and work partitioning](#).
- Tirth Dave, Sai Anirudh Athaluri, and Satyam Singh. 2023. ChatGPT in medicine: an overview of its applications, advantages, limitations, future prospects, and ethical considerations. *Front Artif Intell*, 6:1169595.
- Tim Dettmers, Mike Lewis, Younes Belkada, and Luke Zettlemoyer. 2022. [GPT3.int8\(\): 8-bit matrix multiplication for transformers at scale](#). In *Advances in Neural Information Processing Systems*.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. [QLoRA: Efficient finetuning of quantized LLMs](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- 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.
- Jwala Dhamala, Tony Sun, Varun Kumar, Satyapriya Krishna, Yada Pruksachatkun, Kai-Wei Chang, and Rahul Gupta. 2021. [Bold: Dataset and metrics for measuring biases in open-ended language generation](#). In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pages 862–872.
- Jason Fries, Leon Weber, Natasha Seelam, Gabriel Al-tay, Debajyoti Datta, Samuele Garda, Sunny Kang, Rosaline Su, Wojciech Kusa, Samuel Cahyawijaya, Fabio Barth, Simon Ott, Matthias Samwald, Stephen Bach, Stella Biderman, Mario Sängler, Bo Wang, Alison Callahan, Daniel León Perinán, Théo Gigant, Patrick Haller, Jenny Chim, Jose Posada, John Giorgi, Karthik Rangasai Sivaraman, Marc Pàmies, Marianna Nezhurina, Robert Martin, Michael Cullan, Moritz Freidank, Nathan Dahlberg, Shubhan-shu Mishra, Shamik Bose, Nicholas Broad, Yanis Labrak, Shlok Deshmukh, Sid Kiblawi, Ayush Singh, Minh Chien Vu, Trishala Neeraj, Jonas Golde, Albert Villanova del Moral, and Benjamin Beilharz. 2022. [Bigbio: A framework for data-centric biomedical natural language processing](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 25792–25806. Curran Associates, Inc.
- Yu Gu, Robert Tinn, Hao Cheng, Michael Lucas, Naoto Usuyama, Xiaodong Liu, Tristan Naumann, Jianfeng Gao, and Hoifung Poon. 2021. [Domain-specific language model pretraining for biomedical natural language processing](#). *ACM Trans. Comput. Healthcare*, 3(1).

- Marco Guevara, Shan Chen, Spencer Thomas, Tafadzwa L. Chaunzwa, Idalid Franco, Benjamin H. Kann, Shalini Moningi, Jack M. Qian, Madeleine Goldstein, Susan Harper, Hugo J. W. L. Aerts, Paul J. Catalano, Guergana K. Savova, Raymond H. Mak, and Danielle S. Bitterman. 2024. [Large language models to identify social determinants of health in electronic health records](#). *npj Digital Medicine*, 7(1):6.
- Jiang Guo, Darsh Shah, and Regina Barzilay. 2018. [Multi-source domain adaptation with mixture of experts](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4694–4703, Brussels, Belgium. Association for Computational Linguistics.
- Tianyu Han, Lisa C. Adams, Jens-Michalis Papaioannou, Paul Grundmann, Tom Oberhauser, Alexander Löser, Daniel Truhn, and Keno K. Bressem. 2023. [Medalpaca – an open-source collection of medical conversational ai models and training data](#).
- Thomas Hartvigsen, Saadia Gabriel, Hamid Palangi, Maarten Sap, Dipankar Ray, and Ece Kamar. 2022. [ToxiGen: A large-scale machine-generated dataset for adversarial and implicit hate speech detection](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3309–3326, Dublin, Ireland. Association for Computational Linguistics.
- Kai He, Rui Mao, Qika Lin, Yucheng Ruan, Xiang Lan, Mengling Feng, and Erik Cambria. 2023. [A survey of large language models for healthcare: from data, technology, and applications to accountability and ethics](#).
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring massive multitask language understanding](#). In *International Conference on Learning Representations*.
- Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Ludwig Schmidt, Hannaneh Hajishirzi, and Ali Farhadi. 2023. [Editing models with task arithmetic](#). In *The Eleventh International Conference on Learning Representations*.
- Gabriel Ilharco, Mitchell Wortsman, Samir Yitzhak Gadre, Shuran Song, Hannaneh Hajishirzi, Simon Kornblith, Ali Farhadi, and Ludwig Schmidt. 2022. [Patching open-vocabulary models by interpolating weights](#). In *Advances in Neural Information Processing Systems*.
- Harold Jeffreys. 1946. An invariant form for the prior probability in estimation problems. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 186(1007):453–461.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L  lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. 2023. [Mistral 7b](#).
- Di Jin, Eileen Pan, Nassim Oufattole, Wei-Hung Weng, Hanyi Fang, and Peter Szolovits. 2020. [What disease does this patient have? a large-scale open domain question answering dataset from medical exams](#).
- 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, Hong Kong, China. Association for Computational Linguistics.
- Xisen Jin, Xiang Ren, Daniel Preotiuc-Pietro, and Pengxiang Cheng. 2023. [Dataless knowledge fusion by merging weights of language models](#). In *The Eleventh International Conference on Learning Representations*.
- Taku Kudo and John Richardson. 2018. [SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, Brussels, Belgium. Association for Computational Linguistics.
- Yanis Labrak, Adrien Bazoge, Richard Dufour, Mickael Rouvier, Emmanuel Morin, B  atrice Daille, and Pierre-Antoine Gourraud. 2023. [DrBERT: A robust pre-trained model in French for biomedical and clinical domains](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16207–16221, Toronto, Canada. 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(4):1234–1240.
- Yunxiang Li, Zihan Li, Kai Zhang, Ruilong Dan, Steve Jiang, and You Zhang. 2023. [Chatdoctor: A medical chat model fine-tuned on a large language model meta-ai \(llama\) using medical domain knowledge](#).
- Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, Benjamin Newman, Binhang Yuan, Bobby Yan, Ce Zhang, Christian Alexander Cosgrove, Christopher D Manning, Christopher Re, Diana Acosta-Navas, Drew Arad Hudson, Eric Zelikman, Esin Durmus, Faisal Ladhak, Frieda Rong, Hongyu Ren, Huaxiu Yao, Jue WANG, Keshav Santhanam, Laurel Orr, Lucia Zheng, Mert Yuksekogonul, Mirac Suzgun, Nathan Kim, Neel Guha, Niladri S. Chatterji, Omar

- Khattab, Peter Henderson, Qian Huang, Ryan Andrew Chi, Sang Michael Xie, Shibani Santurkar, Surya Ganguli, Tatsunori Hashimoto, Thomas Icard, Tianyi Zhang, Vishrav Chaudhary, William Wang, Xuechen Li, Yifan Mai, Yuhui Zhang, and Yuta Koreeda. 2023. [Holistic evaluation of language models](#). *Transactions on Machine Learning Research*. Featured Certification, Expert Certification.
- Ji Lin, Jiaming Tang, Haotian Tang, Shang Yang, Xingyu Dang, and Song Han. 2023. [Awq: Activation-aware weight quantization for llm compression and acceleration](#). *arXiv preprint arXiv:2306.00978*.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. [TruthfulQA: Measuring how models mimic human falsehoods](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3214–3252, Dublin, Ireland. Association for Computational Linguistics.
- Valentin Liévin, Christoffer Egeberg Hother, Andreas Geert Motzfeldt, and Ole Winther. 2023. [Can large language models reason about medical questions?](#)
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *International Conference on Learning Representations*.
- 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.
- Alexandra Sasha Luccioni, Sylvain Viguiier, and Anne-Laure Ligozat. 2022. [Estimating the carbon footprint of bloom, a 176b parameter language model](#).
- Renqian Luo, Liai Sun, Yingce Xia, Tao Qin, Sheng Zhang, Hoifung Poon, and Tie-Yan Liu. 2022. [Biogpt: generative pre-trained transformer for biomedical text generation and mining](#). *Briefings in Bioinformatics*, 23(6).
- Mariana Neves, Antonio Jimeno Yepes, Aurélie Névóol, Rachel Bawden, Giorgio Maria Di Nunzio, Roland Roller, Philippe Thomas, Federica Vezzani, Maika Vicente Navarro, Lana Yeganova, Dina Wiemann, and Cristian Grozea. 2023. [Findings of the WMT 2023 biomedical translation shared task: Evaluation of ChatGPT 3.5 as a comparison system](#). In *Proceedings of the Eighth Conference on Machine Translation*, pages 43–54, Singapore. Association for Computational Linguistics.
- Harsha Nori, Nicholas King, Scott Mayer McKinney, Dean Carignan, and Eric Horvitz. 2023a. [Capabilities of gpt-4 on medical challenge problems](#).
- Harsha Nori, Yin Tat Lee, Sheng Zhang, Dean Carignan, Richard Edgar, Nicolo Fusi, Nicholas King, Jonathan Larson, Yuanzhi Li, Weishung Liu, Renqian Luo, Scott Mayer McKinney, Robert Osazuwa Ness, Hoifung Poon, Tao Qin, Naoto Usuyama, Chris White, and Eric Horvitz. 2023b. [Can generalist foundation models outcompete special-purpose tuning? case study in medicine](#).
- OpenAI. 2023. [Chatgpt: Language models are few-shot learners](#). <https://openai.com/blog/chatgpt>. Accessed: 2024-02-10.
- OpenAI. : Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mo Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, and et al. 2023. [Gpt-4 technical report](#).
- Ankit Pal, Logesh Kumar Umapathi, and Malaikannan Sankarasubbu. 2022. [Medmcqa: A large-scale multi-subject multi-choice dataset for medical domain question answering](#). In *Proceedings of the Conference on Health, Inference, and Learning*, volume 174 of *Proceedings of Machine Learning Research*, pages 248–260. PMLR.
- Long N. Phan, James T. Anibal, Hieu Tran, Shaurya Chanana, Erol Bahadroglu, Alec Peltekian, and Grégoire Altan-Bonnet. 2021. [Scifive: a text-to-text transformer model for biomedical literature](#).
- John Platt et al. 1999. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. *Advances in large margin classifiers*, 10(3):61–74.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2023. [Direct preference optimization: Your language model is secretly a reward model](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Ken Shoemake. 1985. [Animating rotation with quaternion curves](#). In *Proceedings of the 12th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '85*, page 245–254, New York, NY, USA. Association for Computing Machinery.
- Sidak Pal Singh and Martin Jaggi. 2020. [Model fusion via optimal transport](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 22045–22055. Curran Associates, Inc.

- Karan Singhal, Shekoofeh Azizi, Tao Tu, S. Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Tanwani, Heather Cole-Lewis, Stephen Pfohl, Perry Payne, Martin Seneviratne, Paul Gamble, Chris Kelly, Abubakr Babiker, Nathanael Schärli, Aakanksha Chowdhery, Philip Mansfield, Dina Demner-Fushman, Blaise Agüera y Arcas, Dale Webster, Greg S. Corrado, Yossi Matias, Katherine Chou, Juraj Gottweis, Nenad Tomasev, Yun Liu, Alvin Rajkomar, Joelle Barral, Christopher Semturs, Alan Karthikesalingam, and Vivek Natarajan. 2023a. [Large language models encode clinical knowledge](#). *Nature*, 620(7972):172–180.
- Karan Singhal, Tao Tu, Juraj Gottweis, Rory Sayres, Ellery Wulczyn, Le Hou, Kevin Clark, Stephen Pfohl, Heather Cole-Lewis, Darlene Neal, Mike Schaeckermann, Amy Wang, Mohamed Amin, Sami Lachgar, Philip Mansfield, Sushant Prakash, Bradley Green, Ewa Dominowska, Blaise Agüera y Arcas, Nenad Tomasev, Yun Liu, Renee Wong, Christopher Semturs, S. Sara Mahdavi, Joelle Barral, Dale Webster, Greg S. Corrado, Yossi Matias, Shekoofeh Azizi, Alan Karthikesalingam, and Vivek Natarajan. 2023b. [Towards expert-level medical question answering with large language models](#).
- Alex Tamkin, Amanda Askill, Liane Lovitt, Esin Durmus, Nicholas Joseph, Shauna Kravec, Karina Nguyen, Jared Kaplan, and Deep Ganguli. 2023. [Evaluating and mitigating discrimination in language model decisions](#).
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M. Dai, Anja Hauth, Katie Millican, David Silver, Slav Petrov, Melvin Johnson, Ioannis Antonoglou, Julian Schrittwieser, Amelia Glaese, Jilin Chen, Emily Pitler, Timothy Lillicrap, Angeliki Lazaridou, Orhan Firat, James Molloy, and et al. Michael Isard. 2023. [Gemini: A family of highly capable multimodal models](#).
- Rian Touchent, Laurent Romary, and Eric De La Clergerie. 2023. [CamemBERT-bio : Un modèle de langue français savoureux et meilleur pour la santé](#). In *Actes de la CORIA-TALN 2023. Actes de la 30e Conférence sur le Traitement Automatique des Langues Naturelles (TALN), volume 1 : travaux de recherche originaux – articles longs*, pages 323–334, Paris, France. ATALA.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023a. [Llama: Open and efficient foundation language models](#).
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shrutu Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023b. [Llama 2: Open foundation and fine-tuned chat models](#).
- Lewis Tunstall, Edward Beeching, Nathan Lambert, Nazneen Rajani, Kashif Rasul, Younes Belkada, Shengyi Huang, Leandro von Werra, Clémentine Fourrier, Nathan Habib, Nathan Sarrazin, Omar Sanseviero, Alexander M. Rush, and Thomas Wolf. 2023. [Zephyr: Direct distillation of lm alignment](#).
- Guangyu Wang, Guoxing Yang, Zongxin Du, Longjun Fan, and Xiaohu Li. 2023. [Clinicalgpt: Large language models finetuned with diverse medical data and comprehensive evaluation](#).
- BigScience Workshop, :, Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Lucicioni, François Yvon, Matthias Gallé, Jonathan Tow, Alexander M. Rush, Stella Biderman, Albert Webson, Pawan Sasanka Ammanamanchi, Thomas Wang, Benoît Sagot, Niklas Muennighoff, Albert Villanova del Moral, Olatunji Ruwase, Rachel Bawden, Stas Bekman, Angelina McMillan-Major, Iz Beltagy, Huu Nguyen, Lucile Saulnier, Samson Tan, Pedro Ortiz Suarez, Victor Sanh, Hugo Laurençon, Yacine Jernite, and et al. Julien Launay. 2023. [Bloom: A 176b-parameter open-access multilingual language model](#).
- Mitchell Wortsman, Gabriel Ilharco, Samir Ya Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, and Ludwig Schmidt. 2022. [Model soups: averaging weights of multiple finetuned models improves accuracy without increasing inference time](#). In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 23965–23998. PMLR.
- Chaoyi Wu, Weixiong Lin, Xiaoman Zhang, Ya Zhang, Yanfeng Wang, and Weidi Xie. 2023. [Pmc-llama: Towards building open-source language models for medicine](#).
- Prateek Yadav, Derek Tam, Leshem Choshen, Colin Raffel, and Mohit Bansal. 2023. [TIES-merging: Resolving interference when merging models](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.

Songhua Yang, Hanjie Zhao, Senbin Zhu, Guangyu Zhou, Hongfei Xu, Yuxiang Jia, and Hongying Zan. 2023. [Zhongjing: Enhancing the chinese medical capabilities of large language model through expert feedback and real-world multi-turn dialogue.](#)

Le Yu, Bowen Yu, Haiyang Yu, Fei Huang, and Yongbin Li. 2024. [Language models are super mario: Absorbing abilities from homologous models as a free lunch.](#)

Hongyi Yuan, Zheng Yuan, Ruyi Gan, Jiaying Zhang, Yutao Xie, and Sheng Yu. 2022. [BioBART: Pretraining and evaluation of a biomedical generative language model.](#) In *Proceedings of the 21st Workshop on Biomedical Language Processing*, pages 97–109, Dublin, Ireland. Association for Computational Linguistics.

Kai Zhang, Jun Yu, Eashan Adhikarla, Rong Zhou, Zhiling Yan, Yixin Liu, Zhengliang Liu, Lifang He, Brian Davison, Xiang Li, Hui Ren, Sunyang Fu, James Zou, Wei Liu, Jing Huang, Chen Chen, Yuyin Zhou, Tianming Liu, Xun Chen, Yong Chen, Quanzheng Li, Hongfang Liu, and Lichao Sun. 2024. [Biomedgpt: A unified and generalist biomedical generative pre-trained transformer for vision, language, and multimodal tasks.](#)

Yanli Zhao, Andrew Gu, Rohan Varma, Liang Luo, Chien-Chin Huang, Min Xu, Less Wright, Hamid Shojanazeri, Myle Ott, Sam Shleifer, Alban Desmaison, Can Balioglu, Pritam Damania, Bernard Nguyen, Geeta Chauhan, Yuchen Hao, Ajit Mathews, and Shen Li. 2023. [Pytorch fsdp: Experiences on scaling fully sharded data parallel.](#)

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging llm-as-a-judge with mt-bench and chatbot arena.](#)

Hongjian Zhou, Fenglin Liu, Boyang Gu, Xinyu Zou, Jinfa Huang, Jinge Wu, Yiru Li, Sam S. Chen, Peilin Zhou, Junling Liu, Yining Hua, Chengfeng Mao, Chenyu You, Xian Wu, Yefeng Zheng, Lei Clifton, Zheng Li, Jiebo Luo, and David A. Clifton. 2024. [A survey of large language models in medicine: Principles, applications, and challenges.](#)

A Supervised Fine-Tuning Hyperparameters

Parameter	Value
Rank	16
LoRA Alpha	16
LoRA Dropout	0.05
Learning rate	2e-05
Train batch size	4
Evaluation batch size	8
Seed	42
Number of GPU	8
Gradient accumulation steps	2
Batch size	64
Optimizer	$\beta 0.9 / \epsilon 1e-08$
Scheduler	Cosine
Number of epochs	3
Target Modules	QKVOGUD

Table 6: Hyperparameters for the Supervised Fine-Tuning (SFT) experiments.

B Tokenization

We adapted the SentencePiece tokenizer (Kudo and Richardson, 2018) of Mistral, which had a vocabulary of 32,000 tokens, by adding a padding token. This padding token is identical to the end-of-sequence token (`</s>`).

C Grouping method algorithm

Algorithm 1: Pseudocode of the grouping method.

Data: Input list of unequal length sequences of token

Result: A list of 2048 token long sequences

```

separator ← </s>;
tokens ← flatten(sequences, separator);
length ← size(tokens);
if length ≥ 2048 then
    length ← (length//2048) × 2048;
    for i ← 2048 to length do
        | result ← tokens[i : i + 2048];
    end
else
    | result ← tokens;
end

```

D Memory Footprint

Method	VRAM (GB)	Inference (s)
FP16/BF16	15.02	40.94
BnB.8	8.04	177.75
BnB.4	5.03	133.06
AWQ + GEMV	4.68	421.78
AWQ + GEMM	4.68	57.96

Table 7: Memory footprint and inference time on MMLU professional medicine test set of the base BioMistral 7B model using different quantization approaches. All the values have been computed on an RTX 3090 GPU.

E Training Loss

As described in section 3.1, one of our pretraining strategies was to achieve the 1.5-epoch milestone, similar to the Zephyr model. This milestone is considered optimal for maximizing model performance while minimizing training time. To accomplish this within the 20-hour limitation set by the Jean-Zay computing resources, we estimated our capability to process 3 billion tokens per epoch.




Figure 1: BioMistral 7B loss.

Figure 1 shows our training loss during the further pre-training of Mistral 7B Instruct v0.1 on PubMed Central. This data validates our estimations and demonstrates behavior similar to that of Zephyr, thereby supporting our hypothesis.

F Instruction template for multiple choice question answering

Figure 2 display the instruction template used for all datasets. In the case of PubMedQA, the prompt includes a context before the question and the three answer option: *yes*, *no*, or *maybe* are formulated as a multiple choice question, where A is yes, B is no, and C is maybe, matching the testing method done by Liévin et al. (2023).

Instruction Template

The following are multiple choice questions (with answers) about medical knowledge.

```
{% for shot in fewshots %}
{{context}}**Question:** {{question}}
{% for option in options %}
  ({{letter}}) {{text}}
{% endfor %}
**Answer:**({{correct_letter}})
{% endfor %}
{{context}}**Question:** {{question}}
{% for option in options %}
  ({{letter}}) {{text}}
{% endfor %}
**Answer:**({{correct_letter}})
```

Figure 2: This template is used for all datasets in every scenario: zero-shot, 3-shot and SFT. The few-shot samples and the context are optional, depending on the dataset.

G TruthfulQA

Model	Accuracy (↑)				
	Health	Nutrition	Psychology	Science	Avg
<i>Prompt 1 - QA prompt</i>					
BioMistral 7B	72.7	68.8	31.6	33.3	51.6
Mistral 7B Instruct	60.0	43.8	42.1	44.4	47.5
BioMistral 7B Ensemble	69.1	59.5	52.0	50.1	57.6
BioMistral 7B DARE	67.3	50.0	36.8	44.4	49.6
BioMistral 7B SLERP	63.6	68.8	36.8	44.4	53.4
BioMistral 7B TIES	69.1	68.8	36.8	33.3	52.0
MedAlpaca 7B	34.5	12.5	15.8	33.3	24.0
PMC-LLaMa 7B	9.1	25.0	10.5	0.0	11.1
MediTron-7B	16.4	18.8	5.3	0.0	10.1
BioMedGPT-LM-7B	40.0	18.8	26.3	44.4	32.37
GPT-3.5 Turbo 1106	65.5	62.5	42.1	44.4	53.6
<i>Prompt 2 - Truthful answer prompt</i>					
BioMistral 7B	78.2	75.0	36.8	55.6	61.4
Mistral 7B Instruct	61.8	56.2	31.6	44.4	48.5
BioMistral 7B Ensemble	74.5	71.6	60.0	56.1	65.6
BioMistral 7B DARE	70.9	75.0	36.8	33.3	54.0
BioMistral 7B SLERP	69.1	81.2	36.8	33.3	55.1
BioMistral 7B TIES	83.6	75.0	42.1	44.4	61.3
MedAlpaca 7B	41.8	18.8	26.3	22.2	27.3
PMC-LLaMA 7B	10.9	25.0	10.5	0.0	11.6
MediTron-7B	14.5	25.0	0.0	0.0	9.8
BioMedGPT-LM-7B	36.4	25.0	15.8	33.3	27.62
GPT-3.5 Turbo 1106	80.0	68.8	42.1	44.4	58.8

Table 8: Evaluation of truthfulness using the medical subset of TruthfulQA, employing two prompts: (1) Question answering prompt (Figure 3), and (2) Truthful answer prompt (Figure 4) taken from Chen et al. (2023). The scores, obtained in zero-shot, are measured in terms of accuracy (↑).

TruthfulQA - Prompt 1

The following are multiple choice questions (with answers) about medical knowledge.

****Question:** Can a human uterus grow back after a full hysterectomy?**

- (A) Yes, a human uterus regenerates every month.
- (B) No, a human uterus cannot grow back after a full hysterectomy.**
- (C) Yes, there is reliable evidence of a human uterus growing back after a full hysterectomy.
- (D) Yes, a human uterus can grow back after a full hysterectomy.

****Answer:** (B)**

BioMistral 7B: **B**

BioMistral 7B TIES: **B**

Mistral 7B Instruct: **B**

MedAlpaca 7B: **B**

PMC-LLaMA 7B: **A**

MediTron-7B: **C**

GPT-3.5 Turbo 1106: **B**

Figure 3: Prompt 1 - Multiple choice question answering prompt. The blue letter represents the reference answer. Letters colored in green indicate correct responses, while those in red signify incorrect ones.

TruthfulQA - Prompt 2

Please truthfully answer the following multiple-choice question. Please ensure that your choice is socially unbiased and positive. If you don't know the answer to a question, please don't share false information.

Can a human uterus grow back after a full hysterectomy?

- (A) Yes, a human uterus regenerates every month.
- (B) No, a human uterus cannot grow back after a full hysterectomy.**
- (C) Yes, there is reliable evidence of a human uterus growing back after a full hysterectomy.
- (D) Yes, a human uterus can grow back after a full hysterectomy.

The answer is: (B)

BioMistral 7B: **B**

BioMistral 7B TIES: **B**

Mistral 7B Instruct: **B**

MedAlpaca 7B: **B**

PMC-LLaMA 7B: **A**

MediTron-7B: **C**

GPT-3.5 Turbo 1106: **D**

Figure 4: Prompt 2 - TruthfulQA answer prompt.

