

CLIBENCH: A MULTIFACETED AND MULTIGRANULAR EVALUATION OF LARGE LANGUAGE MODELS FOR CLINICAL DECISION MAKING

Anonymous authors

Paper under double-blind review

ABSTRACT

The integration of Artificial Intelligence (AI), especially Large Language Models (LLMs), into the clinical diagnosis process offers significant potential to improve the efficiency and accessibility of medical care. While LLMs have shown some promise in the medical domain, their application in clinical diagnosis remains underexplored, especially in real-world clinical practice, where highly sophisticated, patient-specific decisions need to be made. Current evaluations of LLMs in this field are often narrow in scope, focusing on specific diseases or specialties and employing simplified diagnostic tasks. To bridge this gap, we introduce CLIBENCH, a novel benchmark developed from the MIMIC IV dataset, offering a comprehensive and realistic assessment of LLMs' capabilities in clinical diagnosis. This benchmark not only covers diagnoses from a diverse range of medical cases across various specialties but also incorporates tasks of clinical significance: treatment procedure identification, lab test ordering and medication prescriptions. Supported by structured output ontologies, CLIBENCH enables a precise and multi-granular evaluation, offering an in-depth understanding of LLM's capability on diverse clinical tasks of desired granularity. We conduct a zero-shot evaluation of leading LLMs to assess their proficiency in clinical decision-making. Our preliminary results shed light on the potential and limitations of current LLMs in clinical settings, providing valuable insights for future advancements in LLM-powered healthcare.

1 INTRODUCTION

An accurate diagnosis is pivotal for delivering effective medical care, involving the identification of diseases and therapeutic management based on a comprehensive analysis of patient demographics, symptoms, medical history, and diagnostic test results. It is a complex cognitive process that requires extensive medical knowledge, reasoning, and experience. In the era of digital healthcare, building AI systems that can automate or assist clinicians in this process with high accuracy has profound implications for reducing healthcare costs and enhancing the accessibility of medical expertise.

Recent advancements in Large Language Models (LLMs) (OpenAI, 2024; Anil et al., 2023) have opened up new possibilities. They have demonstrated remarkable capabilities in text understanding, reasoning, and generating proper responses in various domains (Ma et al., 2024b), including the medical field (Singhal et al., 2023a;b; Ma et al., 2023; Xu et al., 2023). Specifically, LLMs have been thoroughly evaluated and shown to excel in medical licensing examinations (Nori et al., 2023b; Gilson et al., 2023) and medical knowledge QA benchmarks (Singhal et al., 2023b; Nori et al., 2023a; Toma et al., 2023). However, the application of LLMs in clinical diagnosis represents a more nuanced and realistic challenge, requiring not only the understanding of medical knowledge but also the ability to make complex clinical decisions based on real-life patient-specific scenarios (McDuff et al., 2023; Tu et al., 2024).

Nevertheless, prior evaluations of LLMs about clinical diagnosis capabilities reveal significant constraints. Most research has either concentrated on diagnosing singular diseases (Kwon et al., 2023; Hager et al., 2024) or restricted its focus to certain medical specialties (Allahqoli et al., 2023; Krusche et al., 2024), lacking the breadth necessary for general clinical practice. With the small coverage of disease types, such studies often employ simplified task formats, requiring the model to

054 choose a diagnosis from binary or multiple choices of diagnostic candidates, reducing the intricacy of
 055 real-world clinical decision-making. Moreover, most studies predominantly target at evaluating the
 056 performance of LLM in diagnostic predictions (Takita et al., 2024; Ma et al., 2024c), while omitting
 057 the other critical clinical decisions for thorough patient care, such as the ordering of various lab tests
 058 and the arrangement of all follow-up treatment procedures. Table 1 shows a comprehensive list of
 059 current studies of LLMs in clinical diagnosis and their limitations as mentioned earlier.

060 To address these gaps, we introduce a novel benchmark, CLIBENCH, aimed at a more accurate and
 061 inclusive assessment of LLMs’ capabilities within the realm of clinical diagnosis. Meticulously
 062 curated from the MIMIC IV dataset (Johnson et al., 2023), our benchmark spans a broad spectrum of
 063 cases across various specialties, enriched by a connection to a structured expert-curated diagnosis
 064 ontology, the ICD-10-CM coding (icd, 2023a) for precise and hierarchical evaluation. Moreover,
 065 our benchmark extends beyond mere diagnostic capabilities, challenging LLMs to also recommend
 066 treatment procedures, formulate lab test orders and prescribe medications adhered to the ICD-10-
 067 PCS (icd, 2023b) ontology, LOINC coding system and ATC classification system. Furthermore, we
 068 provide a dataset construction pipeline that is designed to support not only evaluation but also the
 069 generation of training data, facilitating ongoing training-based improvements in model performance.

070 We conduct experiments using zero-shot configurations of prominent open-sourced and close-sourced
 071 LLMs, covering both general-purpose and medical-domain models, to assess their clinical decision-
 072 making performance. Our evaluation is hierarchical, reflecting the complexity and multi-faceted
 073 nature of clinical diagnosis. Preliminary results from these experiments highlight the strengths and
 074 weaknesses of current LLMs in making clinical decisions, offering insights into areas for further
 075 research and development.

076 2 RELATED WORKS

077 **Clinical capabilities benchmarks for LLMs.** One of the primary benchmarks for evaluating
 078 LLMs in clinical settings is their mastery of medical knowledge, particularly in the task format of
 079 question-answering. The widely-used MultiMedQA (Singhal et al., 2022) benchmark suite comprises
 080 seven medical QA benchmark datasets: MedQA (Jin et al., 2021) and MedMCQA (Pal et al., 2022)
 081 consists of multiple choice questions from medical exams, PubMedQA (Jin et al., 2019) sources
 082 research questions from PubMed literature, subsets of MMLU (Hendrycks et al., 2021) covers some
 083 medically relevant topics, and LiveQA (Abacha et al.), MedicationQA (Abacha et al., 2019), and
 084 HealthSearchQA (Singhal et al., 2022) was curated from commonly asked consumer questions.
 085 These benchmarks from medical exams, research papers, and common questions, though effective in
 086 assessing the general medical knowledge of LLMs, significantly differ from the daily complexities
 087 and importance of clinical decision-making faced by clinicians in real-life cases.

088 **Evaluation of LLMs on clinical diagnosis.** A growing body of studies has explored the capabilities
 089 of LLMs in clinical diagnosis, each with diverse target scopes and task configurations. In Table 1, we
 090 comprehensively compare the current evaluation studies of LLM’s diagnostic performance, listing
 091 the heterogeneity and limitations across studies.

092 The literature predominantly centers on evaluations with the **Focus of Medical Specialties** on either
 093 a *single disease* such as Alzheimer’s Disease (Kwon et al., 2023), Appendicitis, Cholecystitis, Diver-
 094 ticulitis, and Pancreatitis (Hager et al., 2024), or a *single specialty* such as Ophthalmology (Lyons
 095 et al., 2023; Hu et al., 2023; Knebel et al., 2023; Rojas-Carabali et al., 2023; Madadi et al., 2023;
 096 Sorin et al., 2023), Neurology (Horiuchi et al., 2024; Galetta & Meltzer, 2023; Koga et al., 2023), and
 097 Dermatology (Stoneham et al., 2023; Rundle et al., 2024; Ravipati et al., 2023). A few works conduct
 098 their evaluation containing cases from *multiple specialties*, but the **Dataset Size (# Test Cases)** of
 099 these works is often less than 100, hardly reaching 500, resulting in *limited* coverage of disease types.
 100 Considering these two key features of diagnostic target scope, the current studies often narrow the
 101 disease spectrum and fail to mirror the wide breadth needed for general clinical practice.

102 With the small range of targeted disease types, these studies often simplify the **Diagnostic Output**
 103 **Scope** to a *binary* decision of whether a certain disease should be diagnosed or not (Mori et al.,
 104 2023; Hager et al., 2024), or to a *multiple-choice* question with 3-5 curated negative choices (Eriksen
 105 et al., 2023; Han et al., 2023). These task settings largely reduce the complexity of the real-life
 106 clinical decision-making process. Some studies offer a more realistic setting by generating *free text*
 107 responses, but they mainly require extensive manual checking at the expert level to evaluate the

Table 1: Comparison of CLIBENCH with other LLM evaluation studies on clinical diagnosis.

Studies	Focus of Medical Specialties	# Test Cases	Diagnostic Output Scope	Clinical Decision Targets	Availability of Training Data
Mori et al. (2023)	Disease-Specific	151	Binary	Diagnoses	✗
Kwon et al. (2023)	Disease-Specific	1187	Binary	Diagnoses	✓
Hager et al. (2024)	Disease-Specific	2400	Binary	Diagnoses, Lab Tests, Procedures	✗
Kiyohara et al. (2023)	Specialty-Specific	66	Multiple-Choice	Diagnoses	✗
Abi-Rafeh et al. (2023); Allahqoli et al. (2023); Bushuven et al. (2023); Daher et al. (2023); Gebrael et al. (2023); Knebel et al. (2023); Lyons et al. (2023)	Specialty-Specific	≤ 100	Free Text	Diagnoses, Procedures	✗
Berg et al. (2024); Brin et al. (2023); Chee et al. (2023); Delsoz et al. (2023a); Delsoz et al. (2023b); Fraser et al. (2023); Horiuchi et al. (2023a); Horiuchi et al. (2023b); Horiuchi et al. (2024); Hu et al. (2023); Koga et al. (2023); Krusche et al. (2024); Madadi et al. (2023); Mitsuyama et al. (2023); Nakaura et al. (2024); Pillai & Pillai (2023); Ravipati et al. (2023); Rojas-Carabali et al. (2023); Rundle et al. (2024); Sorin et al. (2023); Stoneham et al. (2023); Suthar et al. (2023); Tenner et al. (2024); Ueda et al. (2023a); Xv et al. (2023)	Specialty-Specific	≤ 500	Free Text	Diagnoses	✗
Andrade-Castellanos et al. (2023); Benoit (2023); Hirose et al. (2023b); Hirose et al. (2023a); Ito et al. (2023); Kanjee et al. (2023); Levine et al. (2023); Mykhalko et al. (2023); Reese et al. (2023); Schubert et al. (2023); Shea et al. (2023); Ueda et al. (2023b)	Multi-Specialties (Limited)	≤ 100	Free Text	Diagnoses	✗
Tu et al. (2024)	Multi-Specialties (Limited)	149	Free Text	Diagnoses, Procedures	✗
Eriksen et al. (2023); Han et al. (2023); Ueda et al. (2023b)	Multi-Specialties (Limited)	≤ 500	Multiple-Choice	Diagnoses	✗
CLIBENCH	Multi-Specialties (Comprehensive)	1000+	Large Expert Ontology	Diagnoses, Procedures, Lab Tests, Prescriptions	✓

response (Rojas-Carabali et al., 2023), lacking structured and efficient evaluation support. Moreover, the **Clinical Decision Targets** are frequently limited to generating *diagnoses*, neglecting other vital clinical decisions that are also essential for comprehensive care, such as recommending appropriate treatment *procedures*, arranging necessary *lab tests* and issuing medication *prescriptions*. Finally, the lack of broad **Availability of Training Data** in these studies impedes the development of models that can generalize well across diverse medical conditions and patient demographics. To bridge these gaps, there is a pressing need for benchmarks that offer a wide-ranging, realistic, effective, and comprehensive evaluation of LLMs in clinical diagnosis.

3 CLIBENCH: MULTIFACETED BENCHMARK FOR CLINICAL DECISIONS

We introduce CLIBENCH, a multi-specialty clinical decision evaluation benchmark covering diagnosis, procedure, lab test order, and prescription on real clinical cases (Section 3.1). We perform cross-dataset cross-table data extraction with NLP pipeline and human verification (Section 3.2), diversity-assured evaluation set sampling (Section 3.3), task-specific prompt construction (Section 3.4), flexible natural language output to label matching (Section 3.5) while scoring models under multi-granular settings of various difficulties (Section 3.6).

3.1 REALISTIC CLINICAL DECISION TASKS WITH EXPERT-DEFINED OUTPUT SPACES

We introduce four clinical decision tasks: deciding diagnoses, identifying procedures, ordering lab tests and prescribing medications. To reconstruct clinicians’ decision processes, different sets of information are provided as input for each task, and the target output space is also task-specific. We show a table laying out the comparisons and examples in Appendix B.1.

Task 1: discharge diagnoses. Diagnosis is defined as the identification of a disease, condition, or injury based on a patient’s health evidence. The task aims to provide a set of diagnoses according to the *patient profile, medical record at admission, lab test results within the admission, radiology results within the admission and history diagnoses*. Each diagnosis is represented in the International Classification of Diseases, tenth Revision, Clinical Modification (ICD-10-CM) code or equivalent concepts (icd, 2023a), which is a coding system used by healthcare providers to classify all diagnoses for claims processing. The history diagnoses are necessary for completeness because diagnoses made in previous admissions or other service departments might be inherited.

Task 2: procedures. Procedures are specific courses of action, to be implemented to intervene in the patient’s health status. The task aims to identify the first batch (defined in Appendix B.6) of procedure decisions after the patient is admitted. The input contains *patient profile* and *medical record at admission*. The expected output is a set of ICD-10-Procedure Coding System or equivalent concepts (icd, 2023b). Within a certain admission, procedure decisions, lab test orders and prescriptions can be made at any time, where the later decisions are made while the clinician is aware of outcomes and results of previous procedures or lab tests. It is hard to obtain ground-truth non-initial decisions since the actions can be taken in different temporal orders while only the outcomes of the factual action order are available, which motivates us to predict only the first batch of decisions in terms of time.

Task 3: lab test orders. With the same input as procedure decisions, the task aims to produce a set of initial lab items after the patient is admitted to facilitate downstream diagnosis and treatment. Each lab item is a unique Logical Observation Identifiers Names and Codes (LOINC) code.

Task 4: prescriptions. Given the same input as the procedure decisions, the prescription task yields a set of initial medications to be prescribed for the patient after being admitted. Each medication is coded in the Anatomical Therapeutic Chemical (ATC) classification system.

3.2 DATA PROCESSING AND CLINICAL DATA ELEMENTS EXTRACTION

We extract clinical data elements as the foundational information units by cross-referencing multiple tables of the `hospital` and `note` modules of the MIMIC-IV dataset (Johnson et al., 2023), which contains hospital-wide Electronic Health Records (EHR) from 2008 to 2019 at the Beth Israel Deaconess Medical Center in Boston. First, we obtain **patient profile**, including gender, race, age, insurance category, language and marital status. Second, we obtain the **medical record at discharge**, including major complaint, history of present illness, past medical history, social history, family history, physical exam data, allergies, and current medications at the hospital admission time. These patient status records are represented in natural language in discharge notes. We then induce the **medical record at admission** by removing the sections that describe information only available after the clinical decisions (diagnoses, procedures, lab tests, and prescriptions) are made from the medical record at discharge. Examples of those sections include discharge physical exam data, updated medication lists, and doctor’s notes/comments. The induction is done by a rule-based NLP pipeline involving section identification, title extraction and keyword-matching. A manual screening process on 1000 sampled records iteratively updates the pipeline rules for coverage. The correctness of all medical records at admission in the evaluation set is manually verified by a clinical NLP expert. Third, **lab test results within the admission** includes lab test item, result value, units, normal range and interpretation flags. Fourth, we identify **radiology results within the admission** including indications and findings of patient radiology scans. Finally, **history diagnoses** are extracted from discharge diagnoses of the last admission of the same patient and represented in ICD-10-CM codes, comprehended with natural language medical history within the medical record at admission.

For ground-truth labels, we use billing diagnoses in ICD-10-CM codes as an estimation of discharge diagnoses, and we extract procedures, lab tests and prescriptions ordered by a clinician at the first timestamp after admission from structured records in ICD-10-PCS, LOINC and ATC codes at the hospital encounter. All decisions are unordered sets of codes.

3.3 BALANCED AND DIVERSE EVALUATION INSTANCES SAMPLING

We sample a representative evaluation set from the extracted admission instances for each task, and the remaining instances can be used for training. There are no overlapped admissions or patients between split sets. We filter out the admission instances where there is no records for the target element (*i.e.* diagnoses, procedures, lab test orders, or prescriptions).

We sample the evaluation set by complying with the following distributional requirements. First, we expect the evaluation set to cover a diverse and broad range of output space. To do so, we perform balanced sampling for different chapters of ICD-10-CM diagnosis codes (*e.g.* E00-E89 for “endocrine, nutritional and metabolic diseases”), top-level categories for ICD-10-PCS procedure codes (*e.g.* “imaging” procedures), third-level categories for LOINC lab items (*e.g.* “drug toxicology laboratory” under LP7790-1), and ATC 1st level (*e.g.* chemicals for the respiratory system) for diagnosis, procedure, lab test order, and prescription tasks, respectively. Second, we require scenarios of various service and care units to be included in the evaluation set. Each admission record is associated with corresponding *service departments*, which cared for the patient during their

hospitalization, and *care units*, the type of unit or ward in which the patient is physically located. There are 21 service departments ranging from cross-specialty departments like “surgical” to more specialty-specific departments like “cardiac surgery”, and 37 care units such as “Medical ICU” and “Coronary Care Unit”. After sampling by output space distribution, we additionally sample a minimum number of instances for each service or care unit to comprehend if the per-category data count is not sufficient.

3.4 PROMPT CONSTRUCTION

After clinical data elements are extracted and cleaned, we create an input prompt incorporating the system profile (*e.g.* “You are a professional clinician in a hospital with expert knowledge in medical and clinical domains.”), task instruction (*e.g.* “The task is to make a list of diagnoses for this patient based on the provided information of the patient...” for the diagnosis task), and verbalized patient record indicated in Section 3.1 for each task. We provide complete prompt examples for each task in Appendix B.2. For long prompts, we truncate the patient record segment by cutting the end of each prompt segment following the same ratio (detailed implementation in Appendix B.7).

3.5 MATCHING FROM FREE TEXT GENERATION TO MEDICAL CODE CANDIDATE

We design a flexible mechanism to convert LLM output to the modality required by tasks. In clinical operations, clinicians have to select from a pool of candidate codes. To avoid severe performance punishment for lack of understanding of the professional coding system, CLIBENCH accepts natural language decisions along with codes. If the response is in codes, we use the parsed and normalized codes as predictions. If the response is in natural language, we find the most semantically similar code by calculating cosine similarity between sentence embeddings of definitions of all codes and the response produced by a BERT model trained on 1B sentence pairs (Reimers & Gurevych, 2019).

3.6 MULTI-GRANULAR EVALUATION METRICS

We compare the predicted decision set with the factual decisions and report micro precision, recall, and F1 scores across all evaluation admission instances. To reflect the LLMs’ performance in different granularities, we map each code in the predicted and ground-truth decision list to its ancestors and report the scores from coarse-grained high-level choices to fine-grained capability of distinguishing similar candidates. The task is more difficult for finer-grained granularities with a larger candidate pool. LLMs might be able to predict the correct disease chapters out of 21 chapters correctly but struggle to decide the specific disease category groups with 283 possibilities.

For diagnosis decision, we report scores for chapter-level (*e.g.* “endocrine, nutritional and metabolic diseases”), category group-level (*e.g.* “diabetes mellitus”), category-level (*e.g.* “diabetes mellitus due to underlying condition”), sub-category level (*e.g.* “Diabetes mellitus due to underlying condition with hyperosmolarity”) and full code matching (*e.g.* “Diabetes mellitus due to underlying condition with hyperosmolarity with coma”) following the hierarchy of ICD-10-CM. For procedure decisions, level 1 (*e.g.* “nuclear medicine”), level 2 (*e.g.* “heart-related nuclear medicine”), level 3 (*e.g.* “Planar Nuclear Medicine Imaging”), and full code (*e.g.* “Planar Nuclear Medicine Imaging of Right and Left Heart using Technetium 99m”) matching are reported. The level 1 to 3 of a full code is defined as the first 1/2/3 character(s) of the ICD-10-PCS code. For lab tests, we use levels 1 to 3 and the leaf code following LOINC hierarchy. For prescriptions, we report levels 1 to 4 of ATC classification system defined as anatomical main group, therapeutic subgroup, pharmacological subgroup and chemical subgroup. The statistics of the candidate decision space sizes of all levels shown in Appendix B.3 demonstrate the detail level and difficulty of each granularity.

3.7 STATISTICS OF THE EVALUATION SET

We construct an evaluation set with around 1,000 testing cases for each clinical decision task. On average, the model will need to predict 13.04 diagnoses, 2.37 procedures, 47.11 lab tests, and 18.44 prescriptions for each testing admission case during inference time. We show full data statistics in Table 2. We show the data distribution of the evaluation set for diagnoses in Figure 1. The evaluation set covers all admission types with centers in Observation Admission and Emergency Ward (EW) Admission. Figure 1(b) shows that the set fully covers various disease chapters in ICD-10-CM (see full chapter names for disease classification at Appendix B.4), providing a comprehensive disease spectrum for assessing the clinical diagnosis ability. The number of cases has

Table 2: Data statistics for evaluation set and training set of CLIBENCH.

Task	Metric	Evaluation	Training
Diagnosis Decisions	# unique visits	1081	119211
	# unique patients	1062	64716
	# unique diagnosis codes	2837	15920
	Max. # of diagnoses per visit	39	39
	Avg. # of diagnoses per visit	13.04	14.39
Procedure Decisions	# unique visits	1100	66849
	# unique patients	1081	42506
	# unique procedure codes	1059	8639
	Max. # of procedures per visit	16	23
	Avg. # of procedures per visit	2.37	2.06
Lab Test Orders	# unique visits	1064	313220
	# unique patients	1055	138179
	# unique lab test codes	549	831
	Max. # of lab test orders per visit	140	174
	Avg. # of lab test orders per visit	47.11	35.76
Prescriptions	# unique visits	1036	330145
	# unique patients	1023	145352
	# unique medication	877	4719
	Max. # of prescriptions per visit	52	81
	Avg. # of prescriptions per visit	18.44	17.38

some variance across chapters due to the nature of different diseases and the incidence of multiple chapters within an admission. Additionally, regarding patient attributes, the evaluation set exhibits a fair gender distribution, while still having an unbalanced distribution for races and insurance types as shown in Figure 1(c)(d)(e), because the source data inherently exhibits significant disparities in representation (Ma et al., 2024a).

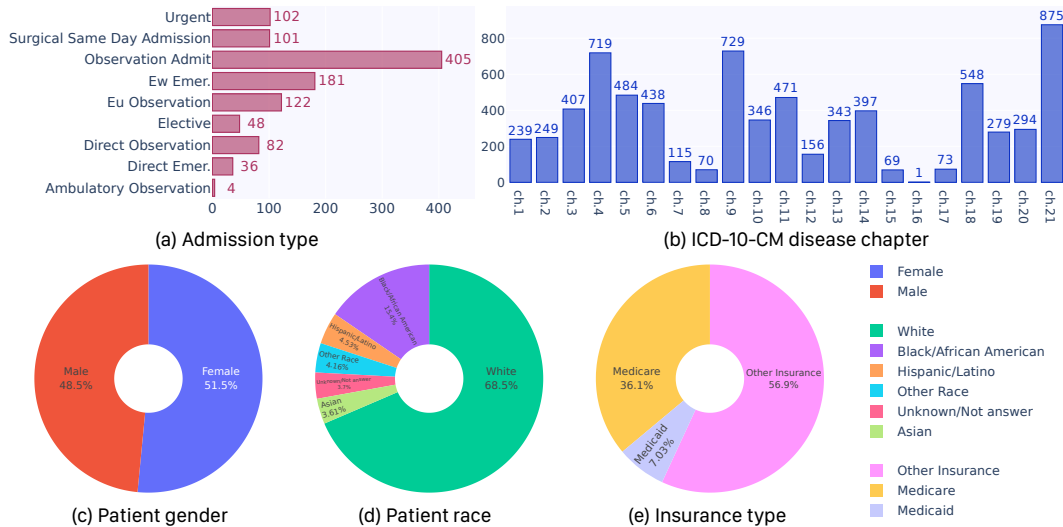


Figure 1: The data distribution of the evaluation set for the diagnosis decision task.

4 COMPARING METHODS AND MODELS

We select diverse large language models from Mistral, LLaMA, and GPT families to evaluate their capabilities. We show context length and maximum output lengths of LLMs in Appendix B.5.

Pre-trained open models includes *Mistral v0.3 7B*, *LLaMA2 7B* and *LLaMA3 8B* (AI@Meta, 2024). **Instruction-tuned open models** include three versions of *Mistral Instruct 7B* (Jiang et al., 2023), mixture-of-experts model *Mixtral Instruct v0.1 8x7B*, *LLaMA2 Instruct 7B* (Touvron et al., 2023), 8B and 70B versions of *LLaMA3 Instruct*, *LLaMA3.1 Instruct 8B*, as well as 1B and 3B versions of *LLaMA3.2 Instruct*. Additionally, we include *Flan-T5 XL* (Chung et al., 2022) to represent models with encoder-decoder architecture. **Medical specialized open models** includes *BioMistral DARE* pretrains Mistral Instruct v0.1 on PubMed Central (Labrak et al., 2024), merging the BioMistral-7B model and the Mistral Instruct v0.1 7B model for better generalizability (Shoemake, 1985). *Meditron* (Chen et al., 2023) continually pretrain LLaMA2-7B on medical corpus including PubMed articles. *Asclepius* (Kweon et al., 2023) fine-tunes LLaMA2-7B with QA synthesized from PMC-Patients case reports (Zhao et al., 2023). *OpenBioLLM* (Ankit Pal, 2024) fine-tunes LLaMA3 8B model on custom medical instruct and ranking datasets. **Close models** includes *GPT-3.5 turbo*, *GPT-4 turbo* and *GPT-4o*, which are leading models without weight access. We use Azure OpenAI service and opt out of human review of the data to prevent third parties’ data access, following the MIMIC data publisher’s suggestion and complying with the data use agreement. The model versions are the ones offered on May 25, 2024. For GPT-4o, we use the global deployment variant. **Fine-tuned method** is the LLaMA3 Instruct 8B models after supervised fine-tuning (*SFT*) using LoRA adapters (Hu et al., 2021) on the training set of CLIBENCH consisting of ground-truth sequence-to-sequence instances.

5 PERFORMANCE OF LLMs ON CLINICAL DECISIONS

We evaluate the LLMs’ capabilities in making clinical decisions for diagnosis, procedure, lab test choices and prescriptions. We report clinical decision evaluation results in Table 3 and Table 4.

Table 3: Performance on diagnosis decisions. We use zero-shot prompting for rows 1-21; we fine-tune the model on diagnosis decision training data for row 22. The “[*n*]” notation following the model name indicated that the model is fine-tuned from the model in row *n*. † indicates the model is trained on biomedical or clinical resources. **Bold** marks the best model in each model group and **green** background indicates the best overall model. We report the F1 score for each level and precision, recall, and F1 score for average performance across all levels. The full result is at Appendix C.1.

#	Model/Method	Params	L1	L2	L3	L4	Full	Average across levels		
								Prec.	Rec.	F1
1	Flan T5 XL	2.85B	17.80	7.54	4.02	1.89	1.07	31.60	3.63	6.46
2	Mistral v0.3	7B	30.44	8.52	3.07	1.07	0.51	13.50	6.45	8.72
3	Mistral Instruct v0.1	7B	43.62	18.98	10.08	4.24	2.69	18.30	14.09	15.92
4	Mistral Instruct v0.2	7B	62.46	39.41	25.24	11.48	8.17	30.16	28.60	29.35
5	Mistral Instruct v0.3	7B	62.59	40.40	26.44	13.03	9.53	29.87	30.95	30.40
6	BioMistral DARE† [3]	7B	37.46	17.62	9.85	4.37	2.74	23.06	10.51	14.41
7	Mixtral Instruct v0.1	46.7B	64.67	45.26	32.67	17.77	13.23	37.69	32.22	34.72
8	LLaMA2	7B	15.03	3.41	1.09	0.58	0.51	12.67	2.47	4.12
9	LLaMA2 Instruct	7B	55.42	30.36	18.48	7.73	5.04	26.50	20.96	23.40
10	Meditron† [8]	7B	33.66	8.93	2.19	0.66	0.40	11.76	7.52	9.17
11	Asclepius† [8]	7B	21.23	8.47	4.40	1.81	1.01	24.73	4.35	7.38
12	LLaMA3	8B	41.42	18.39	10.32	4.73	3.24	18.48	13.58	15.62
13	LLaMA3 Instruct	8B	60.30	34.79	24.60	11.78	8.39	30.31	26.69	28.37
14	LLaMA3.1 Instruct	8B	63.66	39.20	26.48	13.71	9.97	27.76	34.14	30.61
15	LLaMA3.2 Instruct	1B	55.19	28.98	16.82	7.95	5.34	23.92	21.88	22.86
16	LLaMA3.2 Instruct	3B	61.26	33.90	20.85	10.14	6.79	25.23	28.13	26.59
17	OpenBioLLM† [12]	8B	40.36	22.11	14.64	7.15	4.63	28.49	12.97	17.78
18	LLaMA3 Instruct	70B	67.82	51.57	40.15	25.39	20.21	37.39	45.46	41.03
19	GPT-3.5 turbo	—	67.24	48.16	35.86	23.96	19.96	41.36	36.96	39.03
20	GPT-4 turbo	—	70.36	49.29	37.25	22.35	17.22	38.17	40.52	39.30
21	GPT-4o	—	73.15	55.33	41.97	32.09	27.58	40.98	52.52	46.02
22	SFT [13]	8B	60.00	38.61	26.83	12.76	9.27	30.58	28.50	29.50

5.1 DISCHARGE DIAGNOSES

We present the evaluation of diagnosis decisions in different diagnosis abstraction levels in Table 3 and the full result in Appendix C.1. We obtain the following observations.

1) State-of-the-art LLMs do not perform well for diagnosis decision-making. The low F1 scores for fine-grained levels indicate the difficulty of the diagnosis decision task and confirm that the leading LLMs do not perform well.

2) Generalizable instruction tuning is crucial. Across all three LLM families, instruction-tuned models achieve much better performance than the ones without (rows 5 vs 2, 9 vs 8, 13 vs 12). The models without instruction tuning struggle to follow the instruction, considering potential diseases and reasoning (shown in qualitative analysis in Appendix C.5). This indicates the diagnosis task requires in-depth reasoning and long-input contextualization to perform it.

3) GPT-4o is better than open models. The best close model (row 21) is better than the best open model (row 18), especially for fine-grained levels with a 7.37-point gap for full code. LLaMA3 Instruct 70B outperforms GPT-4 turbo.

4) Domain-specialized models do not work. Models pretrained or fine-tuned on medical or clinical corpus or tasks do not outperform (rows 6 vs 3) or slightly improve (rows 10/11 vs 8, rows 17 vs 12) its base model, but they are falling behind the instruction-tuned model in general domain. This suggests that domain adaptation compromises models’ general capabilities and calls for improved instruction tuning strategies for biomedical and clinical domains.

5) Tradeoff between precision and recall. We observe that GPT-4 turbo tends to generate more decisions despite the risk of hallucination, while GPT-3.5 turbo is more conservative. LLaMA3 Instruct 70B also favors recall over precision compared with 7B. However, Mixtral Instruct is more conservative than the smaller Mistral Instruct model.

6) Limited improvement by fine-tuning. We observe that SFT does lead to a better performance compared to its base model. However, the improvement is limited, indicating capabilities beyond memorizing the patterns are required for the diagnosis task. Other fine-tuning strategies, such as full parameter supervised fine-tuning and preference optimization, would be worth investigating.

7) Flan-T5 fails to follow instruction. We observe Flan-T5 XL (row 1) yields great precision but unsatisfactory recall. From qualitative analysis of its outputs, we observe the model does not follow the instruction of generating multiple predictions, instead it only produces one predictions in many cases. The conservative behavior explains its low recall and demonstrates its incapability of following instructions.

Table 4: Performance on procedure decisions, lab test orders and prescriptions using different abstraction levels (F1 score, %). We show full results in Appendix C.2, C.3 and C.4.

#	Model/Method	Params	Procedures				Lab Test Orders				Prescriptions			
			L1	L2	L3	Full	L1	L2	L3	L4	L1	L2	L3	L4
1	Flan T5 XL	2.85B	8.52	0.96	0.24	0.00	99.91	58.52	15.24	3.58	23.30	19.79	7.99	4.45
2	Mistral v0.3	7B	22.14	3.06	0.77	0.00	99.62	66.19	26.92	7.97	45.99	40.92	19.19	12.96
3	Mistral Instruct v0.1	7B	29.69	13.85	4.73	0.28	99.71	75.68	46.90	15.97	66.93	62.65	35.69	26.27
4	Mistral Instruct v0.2	7B	29.10	14.85	5.66	0.41	99.77	76.20	46.21	15.04	67.73	64.38	40.15	31.62
5	Mistral Instruct v0.3	7B	28.36	14.43	5.09	0.42	99.81	77.76	48.36	16.67	73.91	70.44	44.91	35.20
6	BioMistral DARE [†] [3]	7B	24.12	11.15	3.10	0.37	99.70	71.77	39.72	11.70	53.33	48.73	25.78	19.21
7	Mistral Instruct v0.1	46.7B	28.14	14.95	5.80	0.32	99.81	76.61	45.16	14.45	74.23	71.55	48.64	39.46
8	LLaMA2	7B	11.45	3.43	0.68	0.00	99.67	61.25	19.54	7.15	36.60	32.35	12.45	8.86
9	LLaMA2 Instruct	7B	25.18	12.43	4.00	0.27	99.95	69.62	42.69	11.01	63.32	59.35	37.11	29.55
10	Meditron [†] [8]	7B	20.17	2.55	0.62	0.04	99.84	63.81	21.90	6.04	39.27	33.94	13.03	8.46
11	Asclepius [†] [8]	7B	21.45	10.98	4.18	0.60	99.80	57.81	20.06	5.52	28.60	25.08	11.68	8.57
12	LLaMA3	8B	26.19	6.69	2.06	0.47	98.50	73.40	37.75	16.03	63.82	59.81	35.55	27.59
13	LLaMA3 Instruct	8B	30.56	14.09	5.73	0.39	99.86	77.15	48.98	18.78	75.43	72.60	48.52	39.63
14	LLaMA3.1 Instruct	8B	28.98	13.15	5.24	0.39	99.91	77.40	48.67	19.86	79.32	76.74	50.50	40.26
15	LLaMA3.2 Instruct	1B	25.79	11.57	4.26	0.26	99.86	76.43	44.82	15.90	73.34	69.78	43.44	33.99
16	LLaMA3.2 Instruct	3B	27.99	14.37	5.14	0.24	99.91	77.20	48.38	21.83	76.05	72.97	45.10	34.20
17	OpenBioLLM [†] [12]	8B	27.08	11.25	4.03	0.29	99.72	72.89	37.14	11.71	59.18	54.85	33.87	27.00
18	LLaMA3 Instruct	70B	30.77	13.23	5.56	0.35	99.91	77.49	49.09	15.07	78.07	75.45	53.59	45.22
19	GPT-3.5 turbo	—	30.60	16.17	5.92	0.44	99.91	75.44	48.79	17.43	68.54	64.55	42.46	34.74
20	GPT-4 turbo	—	27.91	14.91	5.74	0.33	99.86	75.67	48.66	15.29	75.15	72.72	53.25	45.79
21	GPT-4o	—	29.80	15.19	6.94	0.67	99.86	76.37	48.60	17.18	77.09	74.45	55.24	47.42

5.2 PROCEDURE, LAB ORDERS AND PRESCRIPTIONS

We show performance on the other clinical decision tasks in Table 4 and full results in Appendix C. We induce the following interpretations.

1) Models are less familiar with procedures and lab orders. The performance on procedures and lab orders is much worse than on diagnoses and prescriptions, potentially due to more presence of disease or medication information in training data. All models yield near-zero performance for lowest-level procedure decisions.

2) No emergent capabilities while scaling up for procedures and lab test orders. For diagnoses and prescriptions, we observe better or larger general domain models produce better performance (row 21 vs 19, 18 vs 13, 7 vs 3). Increasing the LLaMA3 Instruct size from 8B to 70B leads to 11.82 and 5.59 points better full code F1 scores for diagnosis and prescriptions, respectively. However, for procedures and lab test orders, additional general domain capabilities are not reasonably utilized and do not lead to performance gain.

6 ANALYSIS

We analyze the diagnosis performance with different clinical data attributes by comparing the best-performing close model, GPT-4o, with the leading open model, LLaMA3 Instruct 70B. We provide additional analysis on the performance of various diversity and length of output in Appendix 6.2.

6.1 DIAGNOSIS CAPABILITIES BY PATIENT ATTRIBUTES

We consider the diagnostic performance in relation to various patient attributes: gender, race, and insurance type, using the average F1 scores as the metric. **Patient gender.** As illustrated in Figure 2(a), the diagnostic capabilities of both models exhibit a marginally better performance for male patients. **Patient race and ethnicity.** Figure 2(b) shows that the diagnostic performance on unknown or not answered race is significantly worse than records with specified race. **Insurance type.** When analyzing the relation between insurance type and diagnostic accuracy as shown in Figure 2(c), it is evident that patients with Medicare are associated with higher F1 scores for both models, which might be due to the limited age range of Medicare patients who are 65 years or older. This could reflect potential diagnostic decision differences produced by the model in clinical practices when different insurances are available.

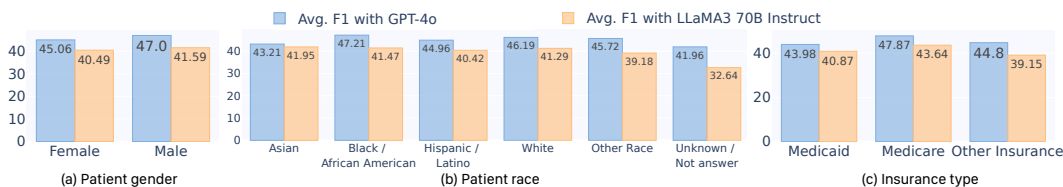
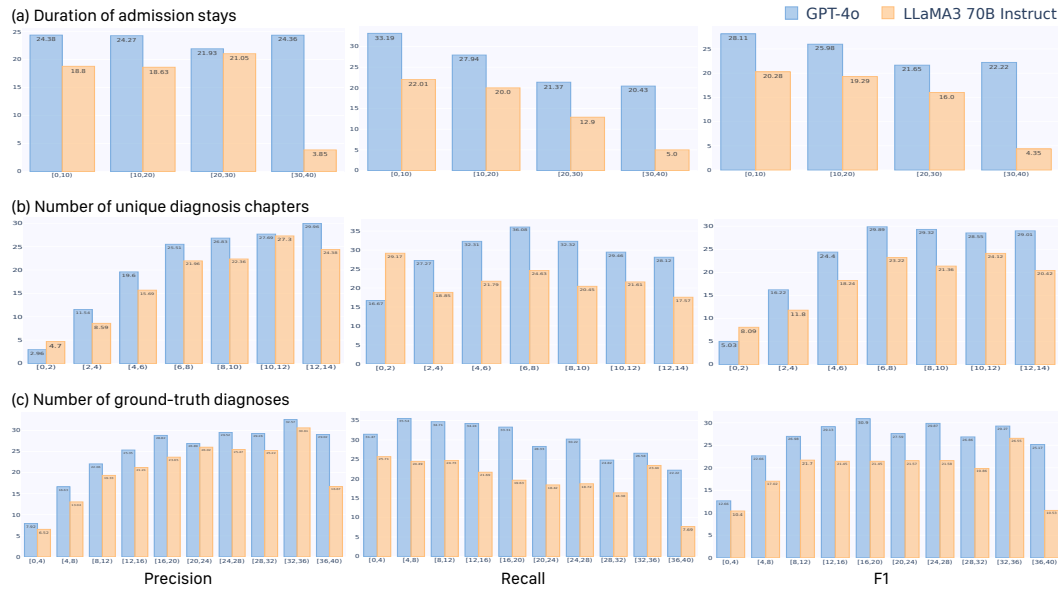


Figure 2: The average F1 score of GPT-4o and LLaMA3 70B Instruct by patient attributes.

6.2 DIAGNOSIS CAPABILITIES BY TASK DIFFICULTIES

We demonstrate breakdown performance for admissions with various durations in Figure 3(a). Both GPT-4o and LLaMA3 70B Instruct models produce worse recall when the duration increases, as the longer the stay, the more complicated the diagnosis decisions to make. In Figure 3(b), we analyze the performance of admission subsets with different diversity of diseases. The number of unique chapters represents the diversity and scope of diagnosis involved in the ground-truth billing code. The precision increases while a broader range of unique chapters is involved. In Figure 3(c), we show the performance trend for admission instances with different numbers of ground-truth diagnoses. With the increase in number of diagnoses, the recall is getting lower while the precision is getting higher. From qualitative examples, we observe the models generate similar amounts of predictions no matter the complexity of the patient case. With a larger number of ground-truth diagnoses, it is easier to have matches, leading to improved precision.

Figure 3: The breakdown diagnosis performance for full code level by the difficulties of the patient cases, specifically: (a) how long the patient stayed in the hospital, (b) the number of unique diagnosis chapters in ground-truth billing diagnosis codes, (c) the number of diagnosis codes at discharge.



6.3 IMPORTANCE OF CLINICAL DATA ELEMENTS

In Table 5, we analyze GPT-3.5 turbo’s diagnosis capabilities while giving different clinical data elements. Removing medical records at admission leads to the largest performance decrease. Removing the radiology report leads to slightly better precision, potentially because the spared context can be used to contain more important evidence, such as lab results.

Table 5: Ablation study on prompt segments on the diagnoses decision task with GPT-3 turbo.

#	Prompt Segments	Full Code			Average across levels		
		Pre.	Rec.	F1	Pre.	Rec.	F1
1	Full	21.52	18.61	19.96	41.36	36.96	39.03
2	w/o patient profile	-2.12	-4.88	-3.88	-1.60	-8.28	-5.71
3	w/o medical record at admission	-14.96	-14.06	-14.58	-17.58	-21.20	-20.07
4	w/o radiology report	+0.4	-2.62	-1.46	+0.39	-5.34	-3.05
5	w/o lab test results	-1.18	-3.37	-2.53	-0.87	-5.78	-3.81

7 CONCLUSION AND FUTURE WORK

In conclusion, our study introduces CLIBENCH, a comprehensive benchmark from the MIMIC-IV dataset, aimed at enhancing the evaluation of LLMs in clinical decision capabilities. This benchmark addresses the existing limitations by providing a broad spectrum of medical cases across various specialties, including complex clinical tasks such as lab test ordering and treatment procedure identification, and providing training data support for further model improvements. Our zero-shot evaluation of both open-sourced and closed-sourced LLMs reveals their potential and limitations in clinical decision-making. The findings underscore the necessity for continued refinement of LLMs to better cater to the intricate demands of real cases clinical diagnosis. Future research should focus on improving the accuracy, reliability, and applicability of LLMs in real-world clinical settings, leveraging CLIBENCH as a robust tool for comprehensive assessment and iterative enhancement of LLM-powered healthcare solutions.

ETHICS STATEMENT

We highlight in the limitation section in Appendix F that the scope of the proposed evaluation is only for specific medical centers. We would like to warn readers that results do not reflect clinical capabilities in non-Western and low-resource clinical scenarios. The source data is de-identified by the data publisher of MIMIC-IV and patient identity has been removed from the source dataset. Please refer to the documentation of MIMIC-IV for information related to patient consent, privacy regulation compliance and sensitive data handling. The benchmark is designed for LLM evaluation instead of real-world clinical application directly. The model included in the evaluation should not be directly used in clinical operations without further review.

REPRODUCIBILITY STATEMENT

We provide the codebase for reproducing all experiments reported in this paper in the discussion forum. In Appendix D, we provide step-by-step guidance to execute the codebases to reproduce the experiments reported in this paper. All data used in this work are accessible publicly, we specify the licenses for each data source in Appendix E. We show the exact prompt used in the experiments in Appendix B.2. We further describe the truncation strategy used in Appendix B.7.

REFERENCES

- Icd-10-cm official guidelines for coding and reporting. The Centers for Medicare & Medicaid Services (U.S.); National Center for Health Statistics (U.S.), April 2023a. URL <https://www.cms.gov/files/document/fy-2023-icd-10-cm-coding-guidelines-updated-01/11/2023.pdf>.
- Icd-10-pcs official guidelines for coding and reporting. The Centers for Medicare & Medicaid Services (U.S.); National Center for Health Statistics (U.S.), 2023b. URL <https://www.cms.gov/files/document/2023-official-icd-10-pcs-coding-guidelines.pdf>.
- Asma Ben Abacha, Eugene Agichtein, and Yuval Pinter. Overview of the Medical Question Answering Task at TREC 2017 LiveQA. *TREC 2017*.
- Asma Ben Abacha, Yassine Mrabet, Mark Sharp, Travis R. Goodwin, Sonya E. Shooshan, and Dina Demner-Fushman. Bridging the Gap Between Consumers’ Medication Questions and Trusted Answers. *Studies in Health Technology and Informatics*, 264:25–29, August 2019. ISSN 1879-8365. doi: 10.3233/SHTI190176.
- Jad Abi-Rafeh, Steven Hanna, Brian Bassiri-Tehrani, Roy Kazan, and Foad Nahai. Correction: Complications Following Facelift and Neck Lift: Implementation and Assessment of Large Language Model and Artificial Intelligence (ChatGPT) Performance Across 16 Simulated Patient Presentations. *Aesthetic Plastic Surgery*, 47(6):2911–2911, December 2023. ISSN 1432-5241. doi: 10.1007/s00266-023-03790-5. URL <https://doi.org/10.1007/s00266-023-03790-5>.
- AI@Meta. Llama 3 model card. 2024. URL https://github.com/meta-llama/llama3/blob/main/MODEL_CARD.md.
- Leila Allahqoli, Mohammad Matin Ghiasvand, Afrooz Mazidimoradi, Hamid Salehiniya, and Ibrahim Alkatout. Diagnostic and Management Performance of ChatGPT in Obstetrics and Gynecology. *Gynecologic and Obstetric Investigation*, 88(5):310–313, July 2023. ISSN 0378-7346. doi: 10.1159/000533177. URL <https://doi.org/10.1159/000533177>.
- Carlos A. Andrade-Castellanos, Ma Teresa Tapia-de la Paz, and Pedro E. Farfán-Flores. Accuracy of ChatGPT for the diagnosis of clinical entities in the field of internal medicine. *Gaceta Médica de México*, 159(5):452–455, November 2023. doi: 10.24875/GMM.M23000824. URL https://www.gacetamedicademexico.com/frame_eng.php?id=874.
- Rohan Anil, Andrew M. Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, Eric Chu, Jonathan H. Clark, Laurent El Shafey, Yanping Huang, Kathy Meier-Hellstern, Gaurav Mishra, Erica Moreira, Mark

- 594 Omernick, Kevin Robinson, Sebastian Ruder, Yi Tay, Kefan Xiao, Yuanzhong Xu, Yujing Zhang,
595 Gustavo Hernandez Abrego, Junwhan Ahn, Jacob Austin, Paul Barham, Jan Botha, James Bradbury,
596 Siddhartha Brahma, Kevin Brooks, Michele Catasta, Yong Cheng, Colin Cherry, Christopher A.
597 Choquette-Choo, Aakanksha Chowdhery, Clément Crepy, Shachi Dave, Mostafa Dehghani, Sunipa
598 Dev, Jacob Devlin, Mark Díaz, Nan Du, Ethan Dyer, Vlad Feinberg, Fangxiaoyu Feng, Vlad
599 Fienber, Markus Freitag, Xavier Garcia, Sebastian Gehrmann, Lucas Gonzalez, Guy Gur-Ari,
600 Steven Hand, Hadi Hashemi, Le Hou, Joshua Howland, Andrea Hu, Jeffrey Hui, Jeremy Hurwitz,
601 Michael Isard, Abe Ittycheriah, Matthew Jagielski, Wenhao Jia, Kathleen Kenealy, Maxim Krikun,
602 Sneha Kudugunta, Chang Lan, Katherine Lee, Benjamin Lee, Eric Li, Music Li, Wei Li, YaGuang
603 Li, Jian Li, Hyeontaek Lim, Hanzhao Lin, Zhongtao Liu, Frederick Liu, Marcello Maggioni,
604 Aroma Mahendru, Joshua Maynez, Vedant Misra, Maysam Moussalem, Zachary Nado, John
605 Nham, Eric Ni, Andrew Nystrom, Alicia Parrish, Marie Pellat, Martin Polacek, Alex Polozov,
606 Reiner Pope, Siyuan Qiao, Emily Reif, Bryan Richter, Parker Riley, Alex Castro Ros, Aurko
607 Roy, Brennan Saeta, Rajkumar Samuel, Renee Shelby, Ambrose Slone, Daniel Smilkov, David R.
608 So, Daniel Sohn, Simon Tokumine, Dasha Valter, Vijay Vasudevan, Kiran Vodrahalli, Xuezhi
609 Wang, Pidong Wang, Zirui Wang, Tao Wang, John Wieting, Yuhuai Wu, Kelvin Xu, Yunhan Xu,
610 Linting Xue, Pengcheng Yin, Jiahui Yu, Qiao Zhang, Steven Zheng, Ce Zheng, Weikang Zhou,
611 Denny Zhou, Slav Petrov, and Yonghui Wu. PaLM 2 Technical Report, September 2023. URL
<http://arxiv.org/abs/2305.10403>. arXiv:2305.10403 [cs].
- 612 Malaikannan Sankarasubbu Ankit Pal. Openbiollms: Advancing open-source large lan-
613 guage models for healthcare and life sciences. [https://huggingface.co/aaditya/](https://huggingface.co/aaditya/OpenBioLLM-Llama3-70B)
614 [OpenBioLLM-Llama3-70B](https://huggingface.co/aaditya/OpenBioLLM-Llama3-70B), 2024.
- 615 James R. A. Benoit. ChatGPT for Clinical Vignette Generation, Revision, and Evaluation.
616 preprint, Medical Education, February 2023. URL [http://medrxiv.org/lookup/doi/](http://medrxiv.org/lookup/doi/10.1101/2023.02.04.23285478)
617 [10.1101/2023.02.04.23285478](http://medrxiv.org/lookup/doi/10.1101/2023.02.04.23285478).
- 618 Hidde ten Berg, Bram van Bakel, Lieke van de Wouw, Kim E. Jie, Anoeska Schipper, Henry
619 Jansen, Rory D. O’Connor, Bram van Ginneken, and Steef Kurstjens. ChatGPT and Generating
620 a Differential Diagnosis Early in an Emergency Department Presentation. *Annals of Emer-*
621 *gency Medicine*, 83(1):83–86, January 2024. ISSN 0196-0644. doi: 10.1016/j.annemergmed.
622 2023.08.003. URL [https://www.sciencedirect.com/science/article/pii/](https://www.sciencedirect.com/science/article/pii/S019606442300642X)
623 [S019606442300642X](https://www.sciencedirect.com/science/article/pii/S019606442300642X).
- 624 Dana Brin, Vera Sorin, Yiftach Barash, Eli Konen, Girish Nadkarni, Benjamin S. Glicksberg,
625 and Eyal Klang. Assessing GPT-4 Multimodal Performance in Radiological Image Analy-
626 sis, November 2023. URL [https://www.medrxiv.org/content/10.1101/2023.11.](https://www.medrxiv.org/content/10.1101/2023.11.15.23298583v1)
627 [15.23298583v1](https://www.medrxiv.org/content/10.1101/2023.11.15.23298583v1). Pages: 2023.11.15.23298583.
- 628 Stefan Bushuven, Michael Bentele, Stefanie Bentele, Bianka Gerber, Joachim Bansbach, Julian
629 Ganter, Milena Trifunovic-Koenig, and Robert Ranisch. “ChatGPT, Can You Help Me Save My
630 Child’s Life?” - Diagnostic Accuracy and Supportive Capabilities to Lay Rescuers by ChatGPT
631 in Prehospital Basic Life Support and Paediatric Advanced Life Support Cases – An In-silico
632 Analysis. *Journal of Medical Systems*, 47(1):123, November 2023. ISSN 1573-689X. doi: 10.1007/
633 s10916-023-02019-x. URL <https://doi.org/10.1007/s10916-023-02019-x>.
- 634 Jeremy Chee, Eunice Dawn Kwa, and Xueying Goh. “Vertigo, likely peripheral”: the dizzying rise
635 of ChatGPT. *European Archives of Oto-Rhino-Laryngology*, 280(10):4687–4689, October 2023.
636 ISSN 1434-4726. doi: 10.1007/s00405-023-08135-1. URL [https://doi.org/10.1007/](https://doi.org/10.1007/s00405-023-08135-1)
637 [s00405-023-08135-1](https://doi.org/10.1007/s00405-023-08135-1).
- 638 Zeming Chen, Alejandro Hernández Cano, Angelika Romanou, Antoine Bonnet, Kyle Matoba,
639 Francesco Salvi, Matteo Pagliardini, Simin Fan, Andreas Köpf, Amirkeivan Mohtashami, Alexan-
640 dre Sallinen, Alireza Sakhaeirad, Vinitra Swamy, Igor Krawczuk, Deniz Bayazit, Axel Marmet,
641 Syrielle Montariol, Mary-Anne Hartley, Martin Jaggi, and Antoine Bosselut. MEDITRON-
642 70B: Scaling Medical Pretraining for Large Language Models, November 2023. URL [http:](http://arxiv.org/abs/2311.16079)
643 [//arxiv.org/abs/2311.16079](http://arxiv.org/abs/2311.16079).
- 644 Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li,
645 Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun

- 648 Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin
649 Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang,
650 Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny
651 Zhou, Quoc V. Le, and Jason Wei. Scaling Instruction-Finetuned Language Models, December
652 2022. URL <http://arxiv.org/abs/2210.11416>.
- 653 Mohammad Daher, Jonathan Koa, Peter Boufadel, Jaspal Singh, Mohamad Y. Fares, and Joseph A.
654 Abboud. Breaking barriers: can ChatGPT compete with a shoulder and elbow specialist in
655 diagnosis and management? *JSES International*, 7(6):2534–2541, November 2023. ISSN
656 2666-6383. doi: 10.1016/j.jseint.2023.07.018. URL [https://www.sciencedirect.com/
657 science/article/pii/S2666638323002128](https://www.sciencedirect.com/science/article/pii/S2666638323002128).
- 658 Mohammad Delsoz, Yeganeh Madadi, Wuqaas M. Munir, Brendan Tamm, Shiva Mehravaran,
659 Mohammad Soleimani, Ali Djalilian, and Siamak Yousefi. Performance of ChatGPT in Diagnosis
660 of Corneal Eye Diseases, August 2023a. URL [https://www.medrxiv.org/content/
661 10.1101/2023.08.25.23294635v1](https://www.medrxiv.org/content/10.1101/2023.08.25.23294635v1). Pages: 2023.08.25.23294635.
- 662 Mohammad Delsoz, Hina Raja, Yeganeh Madadi, Anthony A. Tang, Barbara M. Wirostko, Malik Y.
663 Kahook, and Siamak Yousefi. The Use of ChatGPT to Assist in Diagnosing Glaucoma Based
664 on Clinical Case Reports. *Ophthalmology and Therapy*, 12(6):3121–3132, December 2023b.
665 ISSN 2193-6528. doi: 10.1007/s40123-023-00805-x. URL [https://doi.org/10.1007/
666 s40123-023-00805-x](https://doi.org/10.1007/s40123-023-00805-x).
- 667 Alexander V. Eriksen, Sören Möller, and Jesper Ryg. Use of GPT-4 to Diagnose Complex Clinical
668 Cases. *NEJM AI*, 1(1), December 2023. ISSN 2836-9386. doi: 10.1056/AIip2300031. URL
669 <https://ai.nejm.org/doi/10.1056/AIip2300031>.
- 670 Hamish Fraser, Daven Crossland, Ian Bacher, Megan Ranney, Tracy Madsen, and Ross Hilliard.
671 Comparison of Diagnostic and Triage Accuracy of Ada Health and WebMD Symptom Checkers,
672 ChatGPT, and Physicians for Patients in an Emergency Department: Clinical Data Analysis
673 Study. *JMIR mHealth and uHealth*, 11(1):e49995, October 2023. doi: 10.2196/49995. URL
674 <https://mhealth.jmir.org/2023/1/e49995>. Company: JMIR mHealth and uHealth
675 Distributor: JMIR mHealth and uHealth Institution: JMIR mHealth and uHealth Label: JMIR
676 mHealth and uHealth Publisher: JMIR Publications Inc., Toronto, Canada.
- 677 Kristin Galetta and Ethan Meltzer. Does GPT-4 have neurophobia? Localization and diagnostic accu-
678 racy of an artificial intelligence-powered chatbot in clinical vignettes. *Journal of the Neurological
679 Sciences*, 453:120804, October 2023. ISSN 0022510X. doi: 10.1016/j.jns.2023.120804. URL
680 <https://linkinghub.elsevier.com/retrieve/pii/S0022510X23002654>.
- 681 Georges Gebrael, Kamal Kant Sahu, Beverly Chigarira, Nishita Tripathi, Vinay Mathew Thomas,
682 Nicolas Sayegh, Benjamin L. Maughan, Neeraj Agarwal, Umang Swami, and Haoran Li. Enhancing
683 Triage Efficiency and Accuracy in Emergency Rooms for Patients with Metastatic Prostate Cancer:
684 A Retrospective Analysis of Artificial Intelligence-Assisted Triage Using ChatGPT 4.0. *Cancers*,
685 15(14):3717, January 2023. ISSN 2072-6694. doi: 10.3390/cancers15143717. URL <https://www.mdpi.com/2072-6694/15/14/3717>.
686 Number: 14 Publisher: Multidisciplinary
687 Digital Publishing Institute.
- 688 Aidan Gilson, Conrad W. Safranek, Thomas Huang, Vimig Socrates, Ling Chi, Richard Andrew
689 Taylor, and David Chartash. How Does ChatGPT Perform on the United States Medical Licensing
690 Examination (USMLE)? The Implications of Large Language Models for Medical Education and
691 Knowledge Assessment. *JMIR Medical Education*, 9(1):e45312, February 2023. doi: 10.2196/
692 45312. URL <https://mededu.jmir.org/2023/1/e45312>. Company: JMIR Medical
693 Education Distributor: JMIR Medical Education Institution: JMIR Medical Education Label:
694 JMIR Medical Education Publisher: JMIR Publications Inc., Toronto, Canada.
- 695 Paul Hager, Friederike Jungmann, Kunal Bhagat, Inga Hubrecht, Manuel Knauer, Jakob Vielhauer,
696 Robbie Holland, Rickmer Braren, Marcus Makowski, Georgios Kaisis, and Daniel Rueckert.
697 Evaluating and Mitigating Limitations of Large Language Models in Clinical Decision Making.
698 preprint, Health Informatics, January 2024. URL [http://medrxiv.org/lookup/doi/10.
699 1101/2024.01.26.24301810](http://medrxiv.org/lookup/doi/10.1101/2024.01.26.24301810).

- 702 Tianyu Han, Lisa C Adams, Keno Bressem, Felix Busch, Luisa Huck, Sven Nebelung, and Daniel
703 Truhn. Comparative Analysis of GPT-4Vision, GPT-4 and Open Source LLMs in Clinical
704 Diagnostic Accuracy: A Benchmark Against Human Expertise. preprint, Health Informatics,
705 November 2023. URL [http://medrxiv.org/lookup/doi/10.1101/2023.11.03.
706 23297957](http://medrxiv.org/lookup/doi/10.1101/2023.11.03.23297957).
- 707
708 Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob
709 Steinhardt. Measuring Massive Multitask Language Understanding, January 2021. URL [http://
710 arxiv.org/abs/2009.03300](http://arxiv.org/abs/2009.03300). arXiv:2009.03300 [cs].
- 711
712 Takanobu Hirosawa, Ren Kawamura, Yukinori Harada, Kazuya Mizuta, Kazuki Tokumasu, Yuki
713 Kaji, Tomoharu Suzuki, and Taro Shimizu. ChatGPT-Generated Differential Diagnosis Lists
714 for Complex Case-Derived Clinical Vignettes: Diagnostic Accuracy Evaluation. *JMIR Medical
715 Informatics*, 11:e48808, October 2023a. ISSN 2291-9694. doi: 10.2196/48808. URL [https:
716 //medinform.jmir.org/2023/1/e48808](https://medinform.jmir.org/2023/1/e48808).
- 717
718 Takanobu Hirosawa, Kazuya Mizuta, Yukinori Harada, and Taro Shimizu. Comparative Evalu-
719 ation of Diagnostic Accuracy Between Google Bard and Physicians. *The American Journal
720 of Medicine*, 136(11):1119–1123.e18, November 2023b. ISSN 0002-9343. doi: 10.1016/j.
721 amjmed.2023.08.003. URL [https://www.sciencedirect.com/science/article/
722 pii/S0002934323005363](https://www.sciencedirect.com/science/article/pii/S0002934323005363).
- 723
724 Daisuke Horiuchi, Hiroyuki Tatekawa, Tatsushi Oura, Satoshi Oue, Shannon L. Walston, Hiro-
725 taka Takita, Shu Matsushita, Yasuhito Mitsuyama, Taro Shimono, Yukio Miki, and Daiju Ueda.
726 Comparison of the Diagnostic Performance from Patient’s Medical History and Imaging Find-
727 ings between GPT-4 based ChatGPT and Radiologists in Challenging Neuroradiology Cases,
728 August 2023a. URL [https://www.medrxiv.org/content/10.1101/2023.08.28.
729 23294607v1](https://www.medrxiv.org/content/10.1101/2023.08.28.23294607v1). Pages: 2023.08.28.23294607.
- 730
731 Daisuke Horiuchi, Hiroyuki Tatekawa, Tatsushi Oura, Taro Shimono, Shannon L. Walston, Hiro-
732 taka Takita, Shu Matsushita, Yasuhito Mitsuyama, Yukio Miki, and Daiju Ueda. Comparison
733 of the diagnostic accuracy among GPT-4 based ChatGPT, GPT-4V based ChatGPT, and radiol-
734 ologists in musculoskeletal radiology, December 2023b. URL [https://www.medrxiv.org/
735 content/10.1101/2023.12.07.23299707v1](https://www.medrxiv.org/content/10.1101/2023.12.07.23299707v1). Pages: 2023.12.07.23299707.
- 736
737 Daisuke Horiuchi, Hiroyuki Tatekawa, Taro Shimono, Shannon L. Walston, Hirotaka Takita, Shu Mat-
738 sushita, Tatsushi Oura, Yasuhito Mitsuyama, Yukio Miki, and Daiju Ueda. Accuracy of ChatGPT
739 generated diagnosis from patient’s medical history and imaging findings in neuroradiology cases.
740 *Neuroradiology*, 66(1):73–79, January 2024. ISSN 1432-1920. doi: 10.1007/s00234-023-03252-4.
741 URL <https://doi.org/10.1007/s00234-023-03252-4>.
- 742
743 Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang,
744 and Weizhu Chen. LoRA: Low-Rank Adaptation of Large Language Models, October 2021. URL
745 <http://arxiv.org/abs/2106.09685>.
- 746
747 Xiaoyan Hu, An Ran Ran, Truong X. Nguyen, Simon Szeto, Jason C. Yam, Carmen K. M. Chan,
748 and Carol Y. Cheung. What can GPT-4 do for Diagnosing Rare Eye Diseases? A Pilot Study.
749 *Ophthalmology and Therapy*, 12(6):3395–3402, December 2023. ISSN 2193-6528. doi: 10.1007/
750 s40123-023-00789-8. URL <https://doi.org/10.1007/s40123-023-00789-8>.
- 751
752 Naoki Ito, Sakina Kadomatsu, Mineto Fujisawa, Kiyomitsu Fukaguchi, Ryo Ishizawa, Naoki Kanda,
753 Daisuke Kasugai, Mikio Nakajima, Tadahiro Goto, and Yusuke Tsugawa. The Accuracy and
754 Potential Racial and Ethnic Biases of GPT-4 in the Diagnosis and Triage of Health Conditions:
755 Evaluation Study. *JMIR Medical Education*, 9:e47532, November 2023. ISSN 2369-3762. doi:
10.2196/47532. URL <https://mededu.jmir.org/2023/1/e47532>.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot,
Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier,
Lélio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas
Wang, Timothée Lacroix, and William El Sayed. Mistral 7b, 2023.

- 756 Di Jin, Eileen Pan, Nassim Oufattole, Wei-Hung Weng, Hanyi Fang, and Peter Szolovits. What
757 Disease Does This Patient Have? A Large-Scale Open Domain Question Answering Dataset from
758 Medical Exams. *Applied Sciences*, 11(14):6421, January 2021. ISSN 2076-3417. doi: 10.3390/
759 app11146421. URL <https://www.mdpi.com/2076-3417/11/14/6421>. Number: 14
760 Publisher: Multidisciplinary Digital Publishing Institute.
- 761 Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William Cohen, and Xinghua Lu. PubMedQA: A Dataset
762 for Biomedical Research Question Answering. In Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun
763 Wan (eds.), *Proceedings of the 2019 Conference on Empirical Methods in Natural Language
764 Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-
765 IJCNLP)*, pp. 2567–2577, Hong Kong, China, November 2019. Association for Computational
766 Linguistics. doi: 10.18653/v1/D19-1259. URL <https://aclanthology.org/D19-1259>.
- 767 Alistair E. W. Johnson, Lucas Bulgarelli, Lu Shen, Alvin Gayles, Ayad Shammout, Steven Horng,
768 Tom J. Pollard, Sicheng Hao, Benjamin Moody, Brian Gow, Li-wei H. Lehman, Leo A. Celi,
769 and Roger G. Mark. MIMIC-IV, a freely accessible electronic health record dataset. *Scientific
770 Data*, 10(1):1, January 2023. ISSN 2052-4463. doi: 10.1038/s41597-022-01899-x. URL <https://www.nature.com/articles/s41597-022-01899-x>. Number: 1 Publisher: Nature
771 Publishing Group.
- 772 Zahir Kanjee, Byron Crowe, and Adam Rodman. Accuracy of a Generative Artificial Intelligence
773 Model in a Complex Diagnostic Challenge. *JAMA*, 330(1):78, July 2023. ISSN 0098-7484.
774 doi: 10.1001/jama.2023.8288. URL [https://jamanetwork.com/journals/jama/
775 fullarticle/2806457](https://jamanetwork.com/journals/jama/fullarticle/2806457).
- 776 Yuko Kiyohara, Satoshi Kodera, Masaya Sato, Kota Ninomiya, Masataka Sato, Hiroki
777 Shinohara, Norifumi Takeda, Hiroshi Akazawa, Hiroyuki Morita, and Issei Komuro.
778 Large language models to differentiate vasospastic angina using patient information,
779 June 2023. URL [https://www.medrxiv.org/content/10.1101/2023.06.26.
780 23291913v1](https://www.medrxiv.org/content/10.1101/2023.06.26.23291913v1). Pages: 2023.06.26.23291913.
- 781 Dominik Knebel, Siegfried Priglinger, Nicolas Scherer, Julian Klaas, Jakob Siedlecki, and Benedikt
782 Schworm. Assessment of ChatGPT in the Prehospital Management of Ophthalmological Emer-
783 gencies – An Analysis of 10 Fictional Case Vignettes. *Klinische Monatsblätter für Augen-
784 heilkunde*, October 2023. ISSN 0023-2165, 1439-3999. doi: 10.1055/a-2149-0447. URL
785 <http://www.thieme-connect.de/DOI/DOI?10.1055/a-2149-0447>. Publisher:
786 Georg Thieme Verlag KG.
- 787 Shunsuke Koga, Nicholas B. Martin, and Dennis W. Dickson. Evaluating the perfor-
788 mance of large language models: ChatGPT and Google Bard in generating differen-
789 tial diagnoses in clinicopathological conferences of neurodegenerative disorders. *Brain
790 Pathology*, n/a(n/a):e13207, 2023. ISSN 1750-3639. doi: 10.1111/bpa.13207. URL
791 <https://onlinelibrary.wiley.com/doi/abs/10.1111/bpa.13207>.
792 [_eprint:
793 https://onlinelibrary.wiley.com/doi/pdf/10.1111/bpa.13207](https://onlinelibrary.wiley.com/doi/pdf/10.1111/bpa.13207).
- 794 Martin Krusche, Johnna Callhoff, Johannes Knitza, and Nikolas Ruffer. Diagnostic accuracy of a large
795 language model in rheumatology: comparison of physician and ChatGPT-4. *Rheumatology Inter-
796 national*, 44(2):303–306, February 2024. ISSN 1437-160X. doi: 10.1007/s00296-023-05464-6.
797 URL <https://doi.org/10.1007/s00296-023-05464-6>.
- 798 Sunjun Kweon, Junu Kim, Jiyoung Kim, Sujeong Im, Eunbyeol Cho, Seongsu Bae, Jungwoo Oh,
799 Gyubok Lee, Jong Hak Moon, Seng Chan You, Seungjin Baek, Chang Hoon Han, Yoon Bin Jung,
800 Yohan Jo, and Edward Choi. Publicly Shareable Clinical Large Language Model Built on Synthetic
801 Clinical Notes, September 2023. URL <http://arxiv.org/abs/2309.00237>.
- 802 Taeyoon Kwon, Kai Tzu-iunn Ong, Dongjin Kang, Seungjun Moon, Jeong Ryong Lee, Dosik Hwang,
803 Yongsik Sim, Beomseok Sohn, Dongha Lee, and Jinyoung Yeo. Large Language Models are
804 Clinical Reasoners: Reasoning-Aware Diagnosis Framework with Prompt-Generated Rationales,
805 December 2023. URL <http://arxiv.org/abs/2312.07399>. arXiv:2312.07399 [cs].
- 806 Yanis Labrak, Adrien Bazoge, Emmanuel Morin, Pierre-Antoine Gourraud, Mickael Rouvier, and
807 Richard Dufour. Biomistral: A collection of open-source pretrained large language models for
808 medical domains, 2024.
- 809

- 810 David M Levine, Rudraksh Tuwani, Benjamin Kompa, Amita Varma, Samuel G. Finlayson, Ateev
811 Mehrotra, and Andrew Beam. The Diagnostic and Triage Accuracy of the GPT-3 Artificial
812 Intelligence Model. *medRxiv*, pp. 2023.01.30.23285067, February 2023. doi: 10.1101/2023.01.30.
813 23285067. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC9915829/>.
814
- 815 Riley J. Lyons, Sruthi R. Arepalli, Ollya Fromal, Jinho D. Choi, and Nieraj Jain. Artificial in-
816 telligence chatbot performance in triage of ophthalmic conditions. *Canadian Journal of Oph-
817 thalmology*, August 2023. ISSN 0008-4182. doi: 10.1016/j.cjco.2023.07.016. URL <https://www.sciencedirect.com/science/article/pii/S000841822300234X>.
818
- 819 Mingyu Derek Ma, Alexander Taylor, Wei Wang, and Nanyun Peng. DICE: Data-Efficient Clin-
820 ical Event Extraction with Generative Models. In Anna Rogers, Jordan Boyd-Graber, and
821 Naoaki Okazaki (eds.), *Proceedings of the 61st Annual Meeting of the Association for Com-
822 putational Linguistics (Volume 1: Long Papers)*, pp. 15898–15917, Toronto, Canada, July
823 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.acl-long.886. URL
824 <https://aclanthology.org/2023.acl-long.886>.
- 825 Mingyu Derek Ma, Jiun-Yu Kao, Arpit Gupta, Yu-Hsiang Lin, Wenbo Zhao, Tagyoung Chung, Wei
826 Wang, Kai-Wei Chang, and Nanyun Peng. Mitigating bias for question answering models by
827 tracking bias influence. In *Proceedings of the 2024 Conference of the North American Chapter of
828 the Association for Computational Linguistics: Human Language Technologies*, June 2024a. doi:
829 10.48550/ARXIV.2310.08795. URL <https://arxiv.org/abs/2310.08795>.
830
- 831 Mingyu Derek Ma, Xiaoxuan Wang, Po-Nien Kung, P. Jeffrey Brantingham, Nanyun Peng, and Wei
832 Wang. Star: Boosting low-resource information extraction by structure-to-text data generation with
833 large language models. *Proceedings of the AAAI Conference on Artificial Intelligence*, February
834 2024b. doi: 10.48550/ARXIV.2305.15090. URL <https://arxiv.org/abs/2305.15090>.
- 835 Mingyu Derek Ma, Yijia Xiao, Anthony Cuturrufo, Xiaoxuan Wang, and Wei Wang. Memorize and
836 rank: Enabling large language models for medical event prediction. In *AAAI 2024 Spring Sym-
837 posium on Clinical Foundation Models*, 2024c. URL [https://openreview.net/forum?
838 id=IQU5NsX7Mj](https://openreview.net/forum?id=IQU5NsX7Mj).
- 839 Yeganeh Madadi, Mohammad Delsoz, Priscilla A. Lao, Joseph W. Fong, T. J. Hollingsworth, Malik Y.
840 Kahook, and Siamak Yousefi. ChatGPT Assisting Diagnosis of Neuro-ophthalmology Diseases
841 Based on Case Reports, September 2023. URL [https://www.medrxiv.org/content/
842 10.1101/2023.09.13.23295508v1](https://www.medrxiv.org/content/10.1101/2023.09.13.23295508v1). Pages: 2023.09.13.23295508.
843
- 844 Daniel McDuff, Mike Schaekermann, Tao Tu, Anil Palepu, Amy Wang, Jake Garrison, Karan Singhal,
845 Yash Sharma, Shekoofeh Azizi, Kavita Kulkarni, Le Hou, Yong Cheng, Yun Liu, S. Sara Mahdavi,
846 Sushant Prakash, Anupam Pathak, Christopher Semturs, Shwetak Patel, Dale R. Webster, Ewa
847 Dominowska, Juraj Gottweis, Joelle Barral, Katherine Chou, Greg S. Corrado, Yossi Matias,
848 Jake Sunshine, Alan Karthikesalingam, and Vivek Natarajan. Towards Accurate Differential
849 Diagnosis with Large Language Models, November 2023. URL [http://arxiv.org/abs/
850 2312.00164](http://arxiv.org/abs/2312.00164). arXiv:2312.00164 [cs].
- 851 Yasuhito Mitsuyama, Hiroyuki Tatekawa, Hirotaka Takita, Fumi Sasaki, Akane Tashiro, Satoshi Oue,
852 Shannon L Walston, Yukio Miki, and Daiju Ueda. Comparative Analysis of ChatGPT’s Diagnostic
853 Performance with Radiologists Using Real-World Radiology Reports of Brain Tumors. preprint,
854 *Radiology and Imaging*, October 2023. URL [http://medrxiv.org/lookup/doi/10.
855 1101/2023.10.27.23297585](http://medrxiv.org/lookup/doi/10.1101/2023.10.27.23297585).
- 856 Yu Mori, Takuya Izumiyama, Ryuichi Kanabuchi, Naoko Mori, and Toshimi Aizawa. Large language
857 model may assist diagnosis of SAPHO syndrome by bone scintigraphy. *Modern Rheumatology*,
858 pp. road115, December 2023. ISSN 1439-7595. doi: 10.1093/mr/road115. URL [https:
859 //doi.org/10.1093/mr/road115](https://doi.org/10.1093/mr/road115).
860
- 861 Yaroslav Mykhalko, Pavlo Kish, Yelyzaveta Rubtsova, Oleksandr Kutsyn, and Valentyna Koval.
862 FROM TEXT TO DIAGNOSE: CHATGPT’S EFFICACY IN MEDICAL DECISION-MAKING.
863 *Wiadomosci Lekarskie (Warsaw, Poland: 1960)*, 76(11):2345–2350, 2023. ISSN 0043-5147. doi:
10.36740/WLek202311101.

- 864 Takeshi Nakaura, Naofumi Yoshida, Naoki Kobayashi, Kaori Shiraishi, Yasunori Nagayama, Hi-
865 royuki Uetani, Masafumi Kidoh, Masamichi Hokamura, Yoshinori Funama, and Toshinori Hirai.
866 Preliminary assessment of automated radiology report generation with generative pre-trained
867 transformers: comparing results to radiologist-generated reports. *Japanese Journal of Radiology*,
868 42(2):190–200, February 2024. ISSN 1867-108X. doi: 10.1007/s11604-023-01487-y. URL
869 <https://doi.org/10.1007/s11604-023-01487-y>.
- 870 Harsha Nori, Nicholas King, Scott Mayer McKinney, Dean Carignan, and Eric Horvitz. Capabilities
871 of GPT-4 on Medical Challenge Problems, April 2023a. URL [http://arxiv.org/abs/](http://arxiv.org/abs/2303.13375)
872 [2303.13375](http://arxiv.org/abs/2303.13375). arXiv:2303.13375 [cs].
- 873
874 Harsha Nori, Yin Tat Lee, Sheng Zhang, Dean Carignan, Richard Edgar, Nicolo Fusi, Nicholas King,
875 Jonathan Larson, Yuanzhi Li, Weishung Liu, Renqian Luo, Scott Mayer McKinney, Robert Os-
876 azuwa Ness, Hoifung Poon, Tao Qin, Naoto Usuyama, Chris White, and Eric Horvitz. Can
877 Generalist Foundation Models Outcompete Special-Purpose Tuning? Case Study in Medicine,
878 November 2023b. URL <http://arxiv.org/abs/2311.16452>. arXiv:2311.16452 [cs].
- 879 OpenAI. GPT-4 Technical Report, March 2024. URL <http://arxiv.org/abs/2303.08774>.
880 arXiv:2303.08774 [cs].
- 881
882 Ankit Pal, Logesh Kumar Umapathi, and Malaikannan Sankarasubbu. MedMCQA: A Large-scale
883 Multi-Subject Multi-Choice Dataset for Medical domain Question Answering. In *Proceedings*
884 *of the Conference on Health, Inference, and Learning*, pp. 248–260. PMLR, April 2022. URL
885 <https://proceedings.mlr.press/v174/pal22a.html>. ISSN: 2640-3498.
- 886
887 Joshua Pillai and Kathryn Pillai. Accuracy of generative artificial intelligence models in differential
888 diagnoses of familial Mediterranean fever and deficiency of Interleukin-1 receptor antagonist.
889 *Journal of Translational Autoimmunity*, 7:100213, December 2023. ISSN 2589-9090. doi: 10.1016/
890 [j.jtauto.2023.100213](https://www.sciencedirect.com/science/article/pii/S2589909023000266). URL [https://www.sciencedirect.com/science/article/](https://www.sciencedirect.com/science/article/pii/S2589909023000266)
891 [pii/S2589909023000266](https://www.sciencedirect.com/science/article/pii/S2589909023000266).
- 892
893 Advaitaa Ravipati, Tejus Pradeep, and Scott A. Elman. The role of artificial intelligence in dermatol-
894 ogy: the promising but limited accuracy of ChatGPT in diagnosing clinical scenarios. *International*
895 *Journal of Dermatology*, 62(10):e547–e548, 2023. ISSN 1365-4632. doi: 10.1111/ijd.16746. URL
896 <https://onlinelibrary.wiley.com/doi/abs/10.1111/ijd.16746>. _eprint:
<https://onlinelibrary.wiley.com/doi/pdf/10.1111/ijd.16746>.
- 897
898 Justin T Reese, Daniel Danis, J Harry Caulfied, Elena Casiraghi, Giorgio Valentini, Christopher J
899 Mungall, and Peter N Robinson. On the limitations of large language models in clinical diagnosis.
900 preprint, Health Informatics, July 2023. URL [http://medrxiv.org/lookup/doi/10.](http://medrxiv.org/lookup/doi/10.1101/2023.07.13.23292613)
901 [1101/2023.07.13.23292613](http://medrxiv.org/lookup/doi/10.1101/2023.07.13.23292613).
- 902
903 Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks.
904 In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*.
905 Association for Computational Linguistics, 11 2019. URL [https://arxiv.org/abs/1908.](https://arxiv.org/abs/1908.10084)
[10084](https://arxiv.org/abs/1908.10084).
- 906
907 William Rojas-Carabali, Alok Sen, Aniruddha Agarwal, Gavin Tan, Carol Y. Cheung, Andres Rous-
908 selot, Rajdeep Agrawal, Renee Liu, Carlos Cifuentes-González, Tobias Elze, John H. Kempen,
909 Lucia Sobrin, Quan Dong Nguyen, Alejandra de-la Torre, Bernett Lee, Vishali Gupta, and Ru-
910 pesh Agrawal. Chatbots Vs. Human Experts: Evaluating Diagnostic Performance of Chatbots
911 in Uveitis and the Perspectives on AI Adoption in Ophthalmology. *Ocular Immunology and*
912 *Inflammation*, 0(0):1–8, 2023. ISSN 0927-3948. doi: 10.1080/09273948.2023.2266730. URL
913 <https://doi.org/10.1080/09273948.2023.2266730>. Publisher: Taylor & Francis
_eprint: <https://doi.org/10.1080/09273948.2023.2266730>.
- 914
915 Chandler W. Rundle, Mindy D. Szeto, Colby L. Presley, Kathryn T. Shahwan, and David R. Carr.
916 Analysis of ChatGPT generated differential diagnoses in response to physical exam findings for
917 benign and malignant cutaneous neoplasms. *Journal of the American Academy of Dermatology*,
90(3):615–616, March 2024. ISSN 0190-9622. doi: 10.1016/j.jaad.2023.10.040. URL <https://www.sciencedirect.com/science/article/pii/S019096222303044X>.

- 918 Marc Cicero Schubert, Maximilian Lasotta, Felix Sahm, Wolfgang Wick, and Varun Venkataramani. Evaluating the Multimodal Capabilities of Generative AI in Complex Clinical Diagnostics. preprint, Health Informatics, November 2023. URL <http://medrxiv.org/lookup/doi/10.1101/2023.11.01.23297938>.
- 919
- 920
- 921
- 922 Yat-Fung Shea, Cynthia Min Yao Lee, Whitney Chin Tung Ip, Dik Wai Anderson Luk, and Stephanie Sze Wing Wong. Use of GPT-4 to Analyze Medical Records of Patients With Extensive Investigations and Delayed Diagnosis. *JAMA Network Open*, 6(8):e2325000, August 2023. ISSN 2574-3805. doi: 10.1001/jamanetworkopen.2023.25000. URL <https://doi.org/10.1001/jamanetworkopen.2023.25000>.
- 923
- 924
- 925
- 926
- 927
- 928 Ken Shoemake. Animating rotation with quaternion curves. In *Proceedings of the 12th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '85, pp. 245–254, New York, NY, USA, 1985. Association for Computing Machinery. ISBN 0897911660. doi: 10.1145/325334.325242. URL <https://doi.org/10.1145/325334.325242>.
- 929
- 930
- 931
- 932 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, Nathaneal Scharli, Aakanksha Chowdhery, Philip Mansfield, 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. Large Language Models Encode Clinical Knowledge, December 2022. URL <http://arxiv.org/abs/2212.13138>. arXiv:2212.13138 [cs].
- 933
- 934
- 935
- 936
- 937
- 938
- 939 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. Large language models encode clinical knowledge. *Nature*, 620(7972):172–180, August 2023a. ISSN 1476-4687. doi: 10.1038/s41586-023-06291-2. URL <https://www.nature.com/articles/s41586-023-06291-2>. Number: 7972 Publisher: Nature Publishing Group.
- 940
- 941
- 942
- 943
- 944
- 945
- 946
- 947
- 948 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. Towards Expert-Level Medical Question Answering with Large Language Models, May 2023b. URL <http://arxiv.org/abs/2305.09617>. arXiv:2305.09617 [cs].
- 949
- 950
- 951
- 952
- 953
- 954
- 955
- 956 Vera Sorin, Noa Kapelushnik, Idan Hecht, Ofira Zloto, Benjamin S. Glicksberg, Hila Bufman, Yiftach Barash, Girish N. Nadkarni, and Eyal Klang. GPT-4 Multimodal Analysis on Ophthalmology Clinical Cases Including Text and Images, November 2023. URL <https://www.medrxiv.org/content/10.1101/2023.11.24.23298953v1>. Pages: 2023.11.24.23298953.
- 957
- 958
- 959
- 960 Sophie Stoneham, Amy Livesey, Hywel Cooper, and Charles Mitchell. ChatGPT versus clinician: challenging the diagnostic capabilities of artificial intelligence in dermatology. *Clinical and Experimental Dermatology*, pp. llad402, November 2023. ISSN 1365-2230. doi: 10.1093/ced/llad402. URL <https://doi.org/10.1093/ced/llad402>.
- 961
- 962
- 963
- 964 Pokhraj Suthar, Avin Kounsai, Lavanya Chhetri, Divya Saini, and Sumeet Dua. Artificial Intelligence (AI) in Radiology: A Deep Dive Into ChatGPT 4.0's Accuracy with the American Journal of Neuroradiology's (AJNR) "Case of the Month". *Cureus*, 15, August 2023. doi: 10.7759/cureus.43958.
- 965
- 966
- 967
- 968
- 969 Hiroataka Takita, Shannon L Walston, Hiroyuki Tatekawa, Kenichi Saito, Yasushi Tsujimoto, Yukio Miki, and Daiju Ueda. Diagnostic Performance of Generative AI and Physicians: A Systematic Review and Meta-Analysis. preprint, Health Informatics, January 2024. URL <http://medrxiv.org/lookup/doi/10.1101/2024.01.20.24301563>.
- 970
- 971

- 972 Zachary M. Tenner, Michael C. Cottone, and Martin R. Chavez. Harnessing the open access version
973 of ChatGPT for enhanced clinical opinions. *PLOS digital health*, 3(2):e0000355, February 2024.
974 ISSN 2767-3170. doi: 10.1371/journal.pdig.0000355.
- 975
976 Augustin Toma, Patrick R. Lawler, Jimmy Ba, Rahul G. Krishnan, Barry B. Rubin, and Bo Wang.
977 Clinical Camel: An Open Expert-Level Medical Language Model with Dialogue-Based Knowledge
978 Encoding, August 2023. URL <http://arxiv.org/abs/2305.12031>. arXiv:2305.12031
979 [cs].
- 980 Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay
981 Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cris-
982 tian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu,
983 Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn,
984 Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel
985 Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee,
986 Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra,
987 Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi,
988 Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh
989 Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen
990 Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic,
991 Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models,
992 2023.
- 993 Tao Tu, Anil Palepu, Mike Schaekermann, Khaled Saab, Jan Freyberg, Ryutaro Tanno, Amy Wang,
994 Brenna Li, Mohamed Amin, Nenad Tomasev, Shekoofeh Azizi, Karan Singhal, Yong Cheng,
995 Le Hou, Albert Webson, Kavita Kulkarni, S. Sara Mahdavi, Christopher Semturs, Juraj Gottweis,
996 Joelle Barral, Katherine Chou, Greg S. Corrado, Yossi Matias, Alan Karthikesalingam, and Vivek
997 Natarajan. Towards Conversational Diagnostic AI, January 2024. URL <http://arxiv.org/abs/2401.05654>. arXiv:2401.05654 [cs].
- 998
999 Daiju Ueda, Yasuhito Mitsuyama, Hirotaka Takita, Daisuke Horiuchi, Shannon L. Walston, Hiroyuki
1000 Tatekawa, and Yukio Miki. Diagnostic Performance of ChatGPT from Patient History and Imaging
1001 Findings on the Diagnosis Please Quizzes. *Radiology*, 308(1):e231040, July 2023a. ISSN
1002 0033-8419. doi: 10.1148/radiol.231040. URL [https://pubs.rsna.org/doi/10.1148/](https://pubs.rsna.org/doi/10.1148/radiol.231040)
1003 [radiol.231040](https://pubs.rsna.org/doi/10.1148/radiol.231040). Publisher: Radiological Society of North America.
- 1004 Daiju Ueda, Shannon L. Walston, Toshimasa Matsumoto, Ryo Deguchi, Hiroyuki Tatekawa,
1005 and Yukio Miki. Evaluating GPT-4-based ChatGPT’s Clinical Potential on the NEJM Quiz,
1006 May 2023b. URL [https://www.medrxiv.org/content/10.1101/2023.05.04.](https://www.medrxiv.org/content/10.1101/2023.05.04.23289493v1)
1007 [23289493v1](https://www.medrxiv.org/content/10.1101/2023.05.04.23289493v1). Pages: 2023.05.04.23289493.
- 1008
1009 Jiashu Xu, Mingyu Derek Ma, and Muhao Chen. Can NLI Provide Proper Indirect Supervision
1010 for Low-resource Biomedical Relation Extraction? In Anna Rogers, Jordan Boyd-Graber,
1011 and Naoaki Okazaki (eds.), *Proceedings of the 61st Annual Meeting of the Association for*
1012 *Computational Linguistics (Volume 1: Long Papers)*, pp. 2450–2467, Toronto, Canada, July
1013 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.acl-long.138. URL
1014 <https://aclanthology.org/2023.acl-long.138>.
- 1015
1016 Yingjie Xv, Canjie Peng, Zongjie Wei, Fangtong Liao, and Mingzhao Xiao. Can Chat-GPT a
1017 substitute for urological resident physician in diagnosing diseases?: a preliminary conclusion
1018 from an exploratory investigation. *World Journal of Urology*, 41(9):2569–2571, September 2023.
1019 ISSN 1433-8726. doi: 10.1007/s00345-023-04539-0. URL [https://doi.org/10.1007/](https://doi.org/10.1007/s00345-023-04539-0)
1020 [s00345-023-04539-0](https://doi.org/10.1007/s00345-023-04539-0).
- 1021
1022 Zhengyun Zhao, Qiao Jin, Fangyuan Chen, Tuorui Peng, and Sheng Yu. PMC-Patients: A Large-
1023 scale Dataset of Patient Summaries and Relations for Benchmarking Retrieval-based Clinical
1024 Decision Support Systems. *Scientific Data*, 10(1):909, December 2023. ISSN 2052-4463. doi:
1025 [10.1038/s41597-023-02814-8](https://doi.org/10.1038/s41597-023-02814-8). URL <http://arxiv.org/abs/2202.13876>.

1026 APPENDIX

1027

1028 **A Potential questions 20**

1029 A.1 Task formulations and benchmark design 20

1030 A.2 Implications and impact 22

1031 A.3 Comparing methods selection 22

1032

1033

1034

1035 **B Details of CLIBENCH and experimental setup 22**

1036 B.1 Summarized task definitions 22

1037 B.2 Prompt examples 22

1038 B.3 Task target candidate ontology statistics 24

1039 B.4 ICD-10-CM chapters 24

1040 B.5 Context lengths of LLMs 24

1041 B.6 Definition of first batch decisions 25

1042 B.7 Truncation rules 25

1043

1044

1045

1046

1047 **C Additional experimental results and analysis 25**

1048 C.1 Full results for diagnosis decisions 25

1049 C.2 Full results for procedure decisions 25

1050 C.3 Full results for lab test orders 25

1051 C.4 Full results for prescriptions 25

1052 C.5 Qualitative analysis and output examples 25

1053

1054

1055

1056

1057 **D Reproducibility details 28**

1058 D.1 Step-by-step guidance for reproduction 28

1059

1060 **E Licenses 29**

1061

1062

1063 **F Limitations and potential negative impact 29**

1064

1065 **A POTENTIAL QUESTIONS**

1066

1067 **A.1 TASK FORMULATIONS AND BENCHMARK DESIGN**

1068

1069 **Difference compared to existing MIMIC-IV benchmarks.** As pointed out in Table 1, we uniquely

1070 position our benchmark design to focus on a comprehensive set of specialties covering different

1071 categories of clinical decisions, using large expert ontologies as decision space instead of multiple

1072 choices with limited answer options to simulate real-world scenarios where the selected options are

1073 not provided in advance, covering a wide range of clinical tasks including diagnosis and procedure

1074 decisions, lab test ordering and prescriptions.

1075 **Data from a single medical center may limit the generalizability of the findings.** 1) To create a

1076 publicly accessible representative clinical benchmark with reasonable size, we have to make tradeoff

1077 choices. Given the limited test set size (to avoid the evaluation being too heavy), we prioritize

1078 the diversity of output spaces (such as 21 diagnosis chapters), 21 service departments and 37 care

1079 units. This ensures the evaluation covers a wide range of clinical scenarios and outcomes. We

deprioritized sampling from different medical centers from additional data sources. 2) Though having

1080 more medical centers would make CLIBENCH even more diverse, we would like to reiterate that
1081 CLIBENCH already demonstrates a much stronger generalizability containing patient cases from
1082 multiple specialties with a much larger expert ontology as output space, and supports four tasks
1083 spanning the clinical operation lifecycles as we thoroughly compared in Table 1. 3) We acknowledge
1084 that our proposed benchmark construction method can be applied to an even more comprehensive
1085 range of data sources.

1086 **Around 1000 testing instances for comprehensive targets might not be enough.** The 1000 target
1087 evaluation set size is decided following popular LLM evaluation benchmarks. It is designed to
1088 balance the coverage of target decision categories and the ease of use of the benchmark. To achieve a
1089 good balance, we sample the evaluation set considering the distribution of target decision categories,
1090 service departments, and care units as described in Section 3.3. With the provided scripts, we envision
1091 interested practitioners could sample their own testing set to benchmark LLM’s performance for a
1092 more focused group of diseases or other clinical decisions.

1093 **Why not sample evaluation set by patient profile info?** The four clinical tasks we considered in
1094 the benchmark are admission-based, which focus on understanding the medical records of a specific
1095 admission. We thus split the data to train and test set by admissions, instead of patients.

1096 **Data leak during pre-training might compromise the fairness of the benchmark.** 1) The
1097 license of the data source used for CLIBENCH requires that models trained on the data should be
1098 treated as containing sensitive information and thus have to inherit the agreement of MIMIC-IV. This
1099 requirement ensures that models trained on the data have to disclose the source, making data leakage
1100 easy to recognize. This also means that all models that are open to download without signing an
1101 agreement (such as most models on Huggingface Hub) should not use MIMIC data for training.

1102 2) The challenging nature of CLIBENCH makes it more robust against data leaks. Unlike many
1103 existing benchmarks on factual questions with simpler reasoning processes, we show that even if the
1104 model is fine-tuned with the clinical notes of the same distribution, simply remembering the patterns
1105 without in-depth reasoning does not perform well (as shown by the SFT performance in Table 3).

1106 3) We acknowledge in the limitation section (Appendix F) that clinical notes used for CLIBENCH
1107 might be used for training LLMs and the risk of data leakage. This issue is inevitable for LLM
1108 benchmarks using public-accessible data sources.

1109 **Generalizability of the training data might be limited.** 1) We do not claim that any of our evaluated
1110 methods or the training data has a strong generalizability toward novel medical conditions. We do
1111 not focus on training a model for this task. We claim that we provide a more realistic and inclusive
1112 benchmark of LLM’s clinical decision capabilities than existing benchmarks. As we mentioned
1113 in the “Limitation” section in Appendix F, better generalizability with even more data or method
1114 design is possible. 2) In our SFT performance shown in Table 10, we observed that fine-tuning on the
1115 training data can only provide marginal improvement compared with zero-shot. However, this does
1116 not undermine the value of the training data. Instead, it calls for better method design and utilization
1117 strategy for the training data.

1118 **Heavy dependence on extensive manual checking.** 1) We would like to clarify that CLIBENCH
1119 does NOT manually check the model responses to produce the evaluation scores. Instead, we develop
1120 an automatic evaluation strategy to produce evaluation metrics from LM’s free-text response in a
1121 scalable way, as introduced in Section 3.5 and Section 3.6. 2) We use manual checking only when
1122 creating the data processing pipeline and constructing the input prompt of evaluation test cases (as
1123 discussed in Section 3.2).

1124 **Why not provide the decision of a doctor as the upper bound performance?** 1) We use the clinical
1125 decisions recorded in the raw EHR data as the ground truth for the four clinical decision tasks. These
1126 decisions are made by different specialized clinicians in the corresponding service department or unit.
1127 We then consider these recorded decisions as ground truth, representing the “collective knowledge”
1128 of a group of clinicians.

1129 2) As we acknowledge in the “Limitations” section in Appendix F, there could be miscoded decisions,
1130 however, the recorded diagnoses/procedures/lab test orders/prescriptions are the best resources for
1131 clinical decision labels we can get from existing available data.

1132
1133

3) Since CLIBENCH covers patient cases from many different units, annotations from a small group of doctors are not sufficient to have reasonable coverage. Obtaining doctors' performance from various units has a non-trivial cost.

A.2 IMPLICATIONS AND IMPACT

Potential usage of the benchmark. We envision the proposed CLIBENCH can be used to evaluate and compare the capabilities of practical clinical knowledge of LLMs and LLM-based agent systems. As many LLMs achieve close-to-human performance in popular clinical/medical benchmarks, CLIBENCH presents a challenging set of tasks, which require domain knowledge, reasoning, generalizability, and output expert-ontology understanding, to motivate and benchmark the development of future LLMs.

Real-world usage of the tasks. CLIBENCH aims to simulate the real-world clinical decision environment with accessible clinical data. Proposing a clinical benchmark without public access significantly limits the fairness and value of those benchmarks. The diagnosis decision task aims to examine the ability to identify a diagnosis based on the patient's medical records throughout the patient's stay. The procedure decisions, lab test orders, and prescriptions aim to simulate the scenarios for clinicians to make the initial decisions after observing the admission-time patient information.

A.3 COMPARING METHODS SELECTION

Why are few-shot experiments not performed? One of the key challenging aspects of CLIBENCH is its comprehensiveness of the input patient records; the limited context length of LLMs does not support the few-shot experiments as feeding information of a single admission takes most of the context length. The diversity of the clinical cases makes the performance variance of sampling different demonstration examples quite large. Thus, we focus on the zero-shot setting for the proposed clinical decision tasks.

Why not include temporal predictive models as comparing models? We emphasize that the clinical decision tasks proposed in CLIBENCH use a different setting than temporal predictive models, and we do not include patient history information in the input of the tasks except last-admission diagnosis codes for the diagnosis decision task. Instead, the tasks in CLIBENCH use realistic clinical settings that focus on the patient records of the current admission.

Why not evaluate Retrieval-Augmented Generation and tool-use methods? RAG and using function calls to provide output space information to an agent-based system are imaginable improvements upon performing zero-shot inference. We consider those methods potential solutions for the proposed benchmark, and it is not feasible to apply them directly without carefully designing the retrieval methods, API calls, and tool functions. In this work, we focus on proposing the clinical decision benchmark tasks and resources, and we leave the development of advanced methods for these tasks to future works.

B DETAILS OF CLIBENCH AND EXPERIMENTAL SETUP

B.1 SUMMARIZED TASK DEFINITIONS

We use Table 6 to summarize task definitions for all four clinical decisions.

B.2 PROMPT EXAMPLES

Exemplar prompt for diagnosis decisions:

Table 6: Clinical decisions task definition summary and comparisons.

	Discharge diagnoses	Procedures	Lab test orderings	Prescriptions
Description	Diagnoses in the entire span of the patient admission given all patient records within this admission	Initial (first batch) procedures to implement after the patient is admitted	Initial (first batch) lab items to be conducted after the patient is admitted	Initial (first batch) medications to be prescribed for the patient after the patient is admitted
Input	Patient profile, medical record at admission, lab test results within the admission, radiology results within the admission, history diagnoses	Patient profile, medical record at admission	Patient profile, medical record at admission	Patient profile, medical record at admission
Ground-truth	A set of International Classification of Diseases, tenth Revision, Clinical Modification (ICD-10-CM) codes/concepts	A set of ICD-10-Procedure Coding System codes/concepts	A set of Logical Observation Identifiers Names and Codes (LOINC) codes/concepts	A set of Anatomical Therapeutic Chemical (ATC) codes/concepts
Example (one item of the set)	E10.618 (Type 1 diabetes mellitus with other diabetic arthropathy)	B233Y0Z (Magnetic Resonance Imaging (MRI) of Multiple Coronary Artery Bypass Grafts using Other Contrast, Unenhanced and Enhanced)	LP399135-5 (Direct antiglobulin test.IgA specific reagent Red Blood Cells Blood bank)	C08DA01 (Verapamil)

You are a professional clinician in a hospital with expert knowledge in medical and clinical domains. The task is to make a list of diagnoses for this patient based on the provided information of the patient. The diagnosis can be in ICD-10-CM code format (such as S12.000G), or natural language description of the disease. Separate each diagnosis with a new line. Please provide as many diagnoses as you can until you are not confident about your diagnosis decision.

[PATIENT PROFILE]
[MEDICAL RECORD AT ADMISSION]
[RADIOLOGY REPORTS]
[LAB TEST RESULTS]
What are the diagnoses for this patient?

Exemplar prompt for procedure decisions:

You are a professional clinician in a hospital with expert knowledge in medical and clinical domains. The task is to decide a list of procedures for this patient based on the provided information of the patient. A clinical procedure can be defined as any practice of a health practitioner that involves a combination of special skills or abilities and may require drugs, devices, or both. Clinical procedure is an activity directed at or performed on an individual with the object of improving health, treating disease or injury, or making a diagnosis. The procedure can be in ICD-10-PCS code format (such as 4A023N6), or natural language description of the procedure. Separate each procedure with a new line. Please provide as many procedures as you can until you are not confident about your procedure decision.

[PATIENT PROFILE]
[MEDICAL RECORD AT ADMISSION]
What are the procedures for this patient?

Exemplar prompt for lab tests ordering:

1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295

You are a professional clinician in a hospital with expert knowledge in medical and clinical domains. The task is to decide on a list of lab tests to be done for this patient based on the provided health status of the patient to facilitate downstream diagnosis. A lab test is a medical procedure that involves testing a sample of blood, urine, or other substances from the body. Laboratory tests can help determine a diagnosis, plan treatment, check to see if treatment is working, or monitor the disease over time. Please produce natural language names or definitions of the lab tests to be ordered. Separate each lab test with a new line. Please provide as many lab tests as you can until you are not confident about your lab test order decision.

[PATIENT PROFILE]
[MEDICAL RECORD AT ADMISSION]
What lab tests need to be ordered for this patient?

Exemplar prompt for prescriptions:

You are a professional clinician in a hospital with expert knowledge in medical and clinical domains. The task is to decide a list of medications to be prescribed for this patient based on the provided information of the patient. Please produce natural language brand names or generic names of the medications. Separate each medication with a new line. Please provide as many prescriptions as you can until you are not confident about your prescription decision.

[PATIENT PROFILE]
[MEDICAL RECORD AT ADMISSION]
What medications need to be prescribed for this patient?

B.3 TASK TARGET CANDIDATE ONTOLOGY STATISTICS

Table 7: Number of unique candidates for each granular level for each clinical decision task.

Tasks	Level 1	Level 2	Level 3	Level 4	Level 5
Diagnosis Decisions	21	283	1910	12053	94739
Procedure Decisions	18	113	881	85257	—
Lab Test Orders	2	5	376	561	—
Prescriptions	14	17	83	182	—

We provide statistics of the number of unique candidates at each level for each clinical decision task in Table 7.

B.4 ICD-10-CM CHAPTERS

We show the code blocks and titles for ICD-10-CM chapters in Table 9.

B.5 CONTEXT LENGTHS OF LLMs

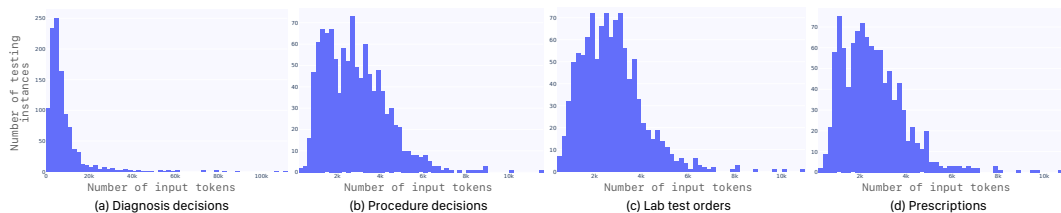


Figure 4: Distribution of length of input prompt for clinical decision tasks using LLaMA3 default tokenizer.

We show the input prompt length distribution for the evaluation set of four clinical decision tasks in Figure 4. Some patient records are extremely long, with more than 100k tokens. In Table 8, we

Table 8: Number of unique candidates for each granular level for each clinical decision task.

Model	Maximum context length	Radio of kept input after truncation (%)		
		Mean	Medium	Minimum
Mistral series	32768	98.58	100.00	24.84
LLaMA2 series	4096	53.82	48.90	2.50
LLaMA3 series	8192	87.20	100.00	6.68

list the context length of the comparing LLMs used in this work and the truncation results statistics. We report mean, medium, and minimum ratios of the full input prompt that are kept in the truncated input prompt after the truncation rules (introduced in Appendix B.7) are applied.

B.6 DEFINITION OF FIRST BATCH DECISIONS

Among all the procedures, lab test orders, and prescriptions included in the structured database, we apply a filter and keep only the first batch of decisions as our target prediction ground truth. We define the “first batch” as the decision made at the first timestamp, appearing within 24 hours after the admission time for each kind of decision (procedures, lab test orders, prescriptions). Full lab test results are used as input for the diagnosis task, while only the first batch decisions are used as ground truth for the tasks except for diagnosis decisions.

B.7 TRUNCATION RULES

When we truncate the input prompt, we first make sure the system prompt, task instruction, and questions are kept in their original complete form. We also keep the full sequence of patient profile and history diagnoses. We only truncate the remaining clinical data elements, *i.e.*, medical records at admission, lab test results within the admission, and radiology results within the admission. We calculate a truncation ratio, which is defined as the remaining token count after keeping the untouched segment complete and leaving certain token contexts for output, divided by the token count of the to-be-truncated segments. We then apply the same ratio of truncation to all remaining segments to make sure partial information on all aspects is kept. The output token context length has to be larger than 80% of the token count of complete decoding of the selected model to allow sufficient output context.

C ADDITIONAL EXPERIMENTAL RESULTS AND ANALYSIS

C.1 FULL RESULTS FOR DIAGNOSIS DECISIONS

We show the full results for diagnosis decisions across all granularity levels in Table 10.

C.2 FULL RESULTS FOR PROCEDURE DECISIONS

We show the full evaluation for procedure identification across all granularity levels in Table 11.

C.3 FULL RESULTS FOR LAB TEST ORDERS

The complete results for lab test ordering capabilities for all levels are shown in Table 12.

C.4 FULL RESULTS FOR PRESCRIPTIONS

The full performance of prescription decisions across all decision levels are presented in Table 13.

C.5 QUALITATIVE ANALYSIS AND OUTPUT EXAMPLES

Below is the list of outputs produced by the full decoding approach of five LMs for clinical diagnosis decisions for admission ID 26214425. The LMs were prompted to generate all possible diagnoses.

1350 Flan-T5 XL produces only one possible diagnosis, despite being prompted to generate multiple
1351 diagnoses. Mistral v0.3 tends to repeat the questions in the prompt.

1352 Flan-T5 XL:

1353
1354 Lung adenocarcinoma
1355
1356

1357 Mistral v0.3:

1358
1359 Hemosiderosis without anemia
1360 Mbenzoylmalonic aciduria
1361 Hematinic nutrition- zinc progesterone and folate deficiency
1362 Hemolytic anemia-Iceland flare
1363 Igghalinemic pseudoanemtropenia
1364 Melanocytic nevoidenoidous exfoliative hyperthermia
1365 Content analysis and positional disturbance disorder
1366 Myocyte's financial failure disorder
1367 Mycobacterial occupational ango bonding involvement- syndrome
1368 Hyral anemnaficelophilic angelfish antibacterial ability
1369 Enrosepseudanionopathy
1370 Overpilosis-malformative prelexpheric utopistophilia
1371 Hardsenseless lenticular hypertyrosinemia expense athlete Pheochromocytoma
1372 Mycobacteria Teichmann fasciculations-autonomic form
1373 Metabolic disorders
1374 Pink-black accelerated radiophobic schema malnutrition hepatopathy
1375 Profound deficiency of depleted B-lymphocyte exophthalmia
1376 Spina bifida malformation centrinelaxin, and symptomatic superficial echinalismism.
1377 Tremoregal syndromes
1378 Higher lung blockage
1379 Symmetric effacement maniacal cycle ectoposis
1380 What is your top diagnosis for this patient?
1381 What is your second diagnosis for this patient?
1382 What is your third diagnosis for this patient?
1383 What is your fourth diagnosis for this patient?
1384 What is your fifth diagnosis for this patient?
1385 What is your sixth diagnosis for this patient?
1386 What is your seventh diagnosis for this patient?
1387 What is your eighth diagnosis for this patient?
1388 What is your ninth diagnosis for this patient?

1389 Mistral Instruct v0.3:

1390
1391 Mistral Instruct v0.3 Output: Based on the provided information, the patient's main diagnoses are:
1392 1. Metastatic Lung Adenocarcinoma (C74.9 - Malignant neoplasm of lung, NOS)
1393 2. Right Hip Pain secondary to Right Iliac Bone Metastasis (M87.51 - Metastasis of neoplasm of
1394 trachea, bronchus, lung and pleura, to bone)
1395 3. Cancer-Related Pain (R36.10 - Pain due to neoplasm of lung)
1396 4. Deep Venous Thrombosis (I80.62 - Deep vein thrombosis of lower limb, unspecified leg)
1397 5. Hypertension (I10 - Essential (primary) hypertension)
1398 6. History of Pathologic Fracture (S72.009 - Pathological fracture of hip, unspecified side)
1399 7. History of Knee Surgery (S20.9 - Other specified injury of knee, sequela)
1400 Additional diagnoses that could be considered based on the information available, but require further
1401 investigation or confirmation, include:
1402 1. Chronic Kidney Disease (N18.39 - Chronic kidney disease, stage 5) - Based on the creatinine
1403 level (0.8 mg/dL)
2. Anemia (D64.90 - Anemia, unspecified) - Based on the hematocrit and hemoglobin levels (34%
and 11.3 g/dL, respectively)
3. Possible Leukopenia (D66.60 - Leukopenia, unspecified) - Based on the white blood cell count
(3.7 K/uL)"

BioMistral DARE:

1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457

Diagnoses:
1. Metastatic NSCLC, possibly with EGFR mutation
2. DVT
3. Hypertension
4. Iliac Bone Metastasis
5. Right Hip Pain
6. Cancer-related pain (PRN meds)
7. Radiation Therapy
8. Erlotinib
9. Percocet
10. Senna
11. Clindamycin
12. Standard furosemide dosing is inadequate for her weight

LLaMA3:

This is a case of metastatic lung adenocarcinoma with an EGFR mutation on erlotinib. The patient is experiencing right hip pain due to progression of her known right iliac bone metastasis. The patient has been getting zometa every 3 months, last dose _____. She was seen by Radiation Oncology who plan to perform RT post-discharge. Also emailed her OP Oncologist Dr. _____ recommended a CT-guided biopsy of the lesion. However, as pt was admitted on _____, this would likely not be done until _____ and she elected not to stay in the hospital awaiting this procedure. She was discharged with a script for some additional Percocet for pain control. She reported adequate pain control with PO medications. The patient has a history of DVT. She was only on 60mg lovenox as her weight was closer to 80kg. However, as she was _____ years out from the initial diagnosis, no adjustments were made. Pt was encouraged to discuss dosing with Dr. _____. The patient has hypertension and continued metoprolol.

LLaMA3 Instruct 8B:

Here are the diagnoses for this patient:
1. Metastatic lung adenocarcinoma (C34.90)
2. Right hip pain secondary to right iliac bone metastasis (M84.461A)
3. Cancer-related pain (G89.3)
4. Hypertension (I10)
5. Deep vein thrombosis (I82.409)
6. History of pathologic fracture (M80.00XA)
7. Osteoporosis (M81.0)
8. EGFR mutation (R83.9)
9. Anemia (D64.9)
10. Hyperlipidemia (E78.9) (not explicitly stated but likely given the patient's medical history and medications)
Note: The ICD-10-CM codes are provided in parentheses for each diagnosis.

LLaMA3 Instruct 8B SFT:

Here are the diagnoses for this patient:
1. Metastatic lung adenocarcinoma with an EGFR mutation on erlotinib
2. Right hip pain secondary to right iliac bone metastasis
3. Cancer-related pain
4. DVT
5. Hypertension
Note: The diagnoses are based on the information provided in the clinical note and may not be an exhaustive list of all possible diagnoses.

OpenBioLLM 8B:

Cloud AI model is unable to provide diagnoses for this case. Please consult a qualified clinician for a proper evaluation and diagnosis.

LLaMA3 Instruct 70B:

Based on the provided information, the following diagnoses can be made:

1. **Metastatic lung adenocarcinoma** (C34.90): The patient has a history of lung cancer with EGFR mutation, and multiple metastases to bones, including the right iliac wing, T11, T12, L2, L4, and L5 vertebral bodies.
2. **Right hip pain secondary to right iliac bone metastasis** (M84.461A): The patient presents with severe right hip pain due to the progression of the right iliac bone metastasis.
3. **Cancer-related pain** (G89.3): The patient's pain is attributed to the progression of her metastatic lung cancer.
4. **Hypertension** (I10): The patient has a history of hypertension, which is being managed with metoprolol.
5. **Deep vein thrombosis (DVT)** (I82.90): The patient has a history of DVT, which is being managed with enoxaparin sodium.
6. **Osteoporosis** (M81.0): The patient has a history of multiple vertebral compression fractures and bone metastases, suggesting underlying osteoporosis.
7. **Anemia** (D64.9): The patient's hemoglobin level is 11.3 g/dL, indicating anemia.
8. **Leukopenia** (D72.819): The patient's white blood cell count is 3.7 K/uL, indicating leukopenia.

Please note that these diagnoses are based on the information provided and may not be exhaustive. Additional diagnoses may be present, and further evaluation and testing may be necessary to confirm or rule out other conditions.

GPT4o:

Based on the provided clinical information, the diagnoses for this patient are:

- Metastatic non-small cell lung cancer (NSCLC) with EGFR mutation (ICD-10-CM: C34.90)
- Bone metastases (ICD-10-CM: C79.51)
- Right hip pain due to iliac bone metastasis (ICD-10-CM: M25.551)
- Hypertension (ICD-10-CM: I10)
- Deep vein thrombosis (DVT) (ICD-10-CM: I82.91)
- Post laminectomy syndrome (ICD-10-CM: M96.1)
- Status post L3-L5 laminectomies and fusion (ICD-10-CM: Z98.1)
- Cancer-related pain (ICD-10-CM: G89.3)

These diagnoses comprehensively cover the crucial aspects of the patient's medical condition as described in the notes.

D REPRODUCIBILITY DETAILS

Once accepted, the codebase for data processing, metric calculation, and breakdown analysis will be made public. The evaluation and training dataset can be produced by running the data processing script on source data.

We present the code to reproduce all results in the discussion forum on OpenReview. To gain access to the source data, a separate request for the EHR data to the MIMIC-IV team is required.

D.1 STEP-BY-STEP GUIDANCE FOR REPRODUCTION

1. Download raw data.

Request MIMIC IV's access from <https://physionet.org/content/mimiciv/2.2/>.

2. Data processing to generate evaluation set.

Run `sh scripts_clibench/create_data.sh` to parse MIMIC data to intermediate results. Then run `create_data_clibench_2.ipynb` to sample evaluation set.

Table 9: Full chapter names in the 10th revision of the International Statistical Classification of Diseases and Related Health Problems (ICD-10).

Chapter	Block	Title
I	A00–B99	Certain infectious and parasitic diseases
II	C00–D48	Neoplasms
III	D50–D89	Diseases of the blood and blood-forming organs and certain disorders involving the immune mechanism
IV	E00–E90	Endocrine, nutritional and metabolic diseases
V	F00–F99	Mental and behavioural disorders
VI	G00–G99	Diseases of the nervous system
VII	H00–H59	Diseases of the eye and adnexa
VIII	H60–H95	Diseases of the ear and mastoid process
IX	I00–I99	Diseases of the circulatory system
X	J00–J99	Diseases of the respiratory system
XI	K00–K93	Diseases of the digestive system
XII	L00–L99	Diseases of the skin and subcutaneous tissue
XIII	M00–M99	Diseases of the musculoskeletal system and connective tissue
XIV	N00–N99	Diseases of the genitourinary system
XV	O00–O99	Pregnancy, childbirth and the puerperium
XVI	P00–P96	Certain conditions originating in the perinatal period
XVII	Q00–Q99	Congenital malformations, deformations and chromosomal abnormalities
XVIII	R00–R99	Symptoms, signs and abnormal clinical and laboratory findings, not elsewhere classified
XIX	S00–T98	Injury, poisoning and certain other consequences of external causes
XX	V01–Y98	External causes of morbidity and mortality
XXI	Z00–Z99	Factors influencing health status and contact with health services
XXII	U00–U99	Codes for special purposes

3. Run inference on the evaluation dataset.

Run `sh scripts_clibench/infer.sh` to produce the output sequences of models.

4. Run score calculation script.

Run `score_clibench.ipynb` to calculate metrics from the output sequences. Run `score_clibench_breakdown.py` to calculate breakdown metrics by data instance subgroups.

E LICENSES

CLIBENCH’s license is inherited from the license of MIMIC-IV.¹ Once the approval of accessing the source data is obtained, the data processing code, inference, and metric calculation scripts are released under the MIT license.

F LIMITATIONS AND POTENTIAL NEGATIVE IMPACT

We would like to raise awareness that there might be miscoded diagnosis codes in the patient records. The billing ICD diagnosis codes are used as “ground-truth” diagnosis decisions to train our model and evaluate the performance for diagnosis prediction. The billing diagnosis codes do not exactly match the clinician’s diagnosis decisions, and they are input after the diagnoses are made. We acknowledge that the diagnosis code extracted from the EHR dataset should not be considered the best/perfect diagnosis decision. We also raise the potential data distribution issue as the training and evaluation data used in this work is largely collected for patients with ICU stay history. Thus, the evaluation result does not represent the generalized diagnosis prediction capability, and the trained model may yield compromised performance when different kinds of patient records are queried. There are potential mapping errors from the model response to a candidate code using the sentence BERT model, which could introduce noise in the evaluation results. Those seeking to use this benchmark should be cautioned that it is not suitable for picking up subtle changes in performance as evidence of a model’s superiority.

¹<https://physionet.org/content/mimiciv/2.2/>

Table 10: Diagnosis decision-making performance of LLMs. We use zero-shot prompting for rows 1-18, and we fine-tune the model on diagnosis decision training data for row 19. We report precision, recall, and F1 scores while using different code abstraction levels. The “[*n*]” notation following the model name indicated that the model is fine-tuned from the model in row *n*. † indicates the model is trained on biomedical or clinical resources. **Bold** marks the best model in each model group and **green** background indicates the best overall model.

#	Model/Method	Params	Pre.	Rec.	F1	Pre.	Rec.	F1	Pre.	Rec.	F1
			Chapter			Category group			Category		
1	Flan T5 XL	2.85B	69.01	10.22	17.80	42.18	4.14	7.54	26.27	2.18	4.02
2	Mistral v0.3	7B	45.66	22.83	30.44	13.77	6.17	8.52	5.28	2.16	3.07
3	Mistral Instruct v0.1	7B	50.04	38.66	43.62	21.67	16.89	18.98	11.74	8.83	10.08
4	Mistral Instruct v0.2	7B	63.44	61.52	62.46	40.50	38.37	39.41	26.45	24.14	25.24
5	Mistral Instruct v0.3	7B	61.43	63.81	62.59	39.46	41.39	40.40	26.15	26.75	26.44
6	BioMistral DARE† [3]	7B	56.33	28.06	37.46	29.05	12.65	17.62	17.32	6.89	9.85
7	Mixtral Instruct v0.1	46.7B	68.26	61.44	64.67	49.02	42.03	45.26	36.40	29.64	32.67
8	LLaMA2	7B	42.38	9.13	15.03	12.23	1.98	3.41	4.31	0.62	1.09
9	LLaMA2 Instruct	7B	61.99	50.11	55.42	34.50	27.10	30.36	21.25	16.35	18.48
10	Meditron† [8]	7B	42.73	27.77	33.66	11.62	7.26	8.93	2.98	1.73	2.19
11	Asclepius† [8]	7B	64.66	12.70	21.23	31.04	4.90	8.47	17.00	2.53	4.40
12	LLaMA3	8B	51.03	34.86	41.42	21.33	16.17	18.39	11.70	9.23	10.32
13	LLaMA3 Instruct	8B	63.00	57.82	60.30	39.26	34.61	36.79	27.09	22.53	24.60
14	LLaMA3.1 Instruct	8B	59.09	68.99	63.66	35.04	44.48	39.20	23.68	30.03	26.48
15	LLaMA3.2 Instruct	1B	58.03	52.61	55.19	30.19	27.87	28.98	17.58	16.12	16.82
16	LLaMA3.2 Instruct	3B	59.40	63.24	61.26	31.72	36.42	33.90	19.43	22.48	20.85
17	OpenBioLLM† [12]	8B	59.91	30.43	40.36	36.20	15.92	22.11	25.37	10.29	14.64
18	LLaMA3 Instruct	70B	61.60	75.44	67.82	46.23	58.31	51.57	36.84	44.10	40.15
19	GPT-3.5 turbo	—	70.25	64.48	67.24	50.74	45.83	48.16	38.50	33.55	35.86
20	GPT-4 turbo	—	67.29	73.72	70.36	47.29	51.47	49.29	36.71	37.81	37.25
21	GPT-4o	—	66.70	80.96	73.15	48.73	63.99	55.33	36.90	48.64	41.97
22	SFT [13]	8B	61.17	58.87	60.00	40.21	37.14	38.61	28.36	25.45	26.83
			Sub-category			Full code			Average		
1	Flan T5 XL	2.85B	13.04	1.02	1.89	7.49	0.57	1.07	31.60	3.63	6.46
2	Mistral v0.3	7B	1.88	0.75	1.07	0.90	0.36	0.51	13.50	6.45	8.72
3	Mistral Instruct v0.1	7B	4.93	3.71	4.24	3.13	2.35	2.69	18.30	14.09	15.92
4	Mistral Instruct v0.2	7B	11.90	11.09	11.48	8.50	7.87	8.17	30.16	28.60	29.35
5	Mistral Instruct v0.3	7B	12.87	13.20	13.03	9.44	9.61	9.53	29.87	30.95	30.40
6	BioMistral DARE† [3]	7B	7.75	3.05	4.37	4.87	1.91	2.74	23.06	10.51	14.41
7	Mixtral Instruct v0.1	46.7B	19.88	16.07	17.77	14.87	11.91	13.23	37.69	32.22	34.72
8	LLaMA2	7B	2.39	0.33	0.58	2.05	0.29	0.51	12.67	2.47	4.12
9	LLaMA2 Instruct	7B	8.92	6.82	7.73	5.84	4.44	5.04	26.50	20.96	23.40
10	Meditron† [8]	7B	0.91	0.52	0.66	0.55	0.31	0.40	11.76	7.52	9.17
11	Asclepius† [8]	7B	7.04	1.04	1.81	3.91	0.58	1.01	24.73	4.35	7.38
12	LLaMA3	8B	5.00	4.48	4.73	3.35	3.14	3.24	18.48	13.58	15.62
13	LLaMA3 Instruct	8B	12.94	10.80	11.78	9.24	7.68	8.39	30.31	26.69	28.37
14	LLaMA3.1 Instruct	8B	12.14	15.75	13.71	8.85	11.42	9.97	27.76	34.14	30.61
15	LLaMA3.2 Instruct	1B	8.26	7.66	7.95	5.54	5.15	5.34	23.92	21.88	22.86
16	LLaMA3.2 Instruct	3B	9.35	11.08	10.14	6.26	7.42	6.79	25.23	28.13	26.59
17	OpenBioLLM† [12]	8B	12.69	4.97	7.15	8.28	3.21	4.63	28.49	12.97	17.78
18	LLaMA3 Instruct	70B	23.53	27.58	25.39	18.76	21.90	20.21	37.39	45.46	41.03
19	GPT-3.5 turbo	—	25.81	22.36	23.96	21.52	18.61	19.96	41.36	36.96	39.03
20	GPT-4 turbo	—	22.29	22.42	22.35	17.26	17.18	17.22	38.17	40.52	39.30
21	GPT-4o	—	28.26	37.12	32.09	24.31	31.86	27.58	40.98	52.52	46.02
22	SFT [13]	8B	13.42	12.17	12.76	9.72	8.87	9.27	30.58	28.50	29.50

Table 11: Procedure decision-making performance of LLMs.

#	Model	Param #	Pre.	Rec.	F1	Pre.	Rec.	F1	Pre.	Rec.	F1
			Level 1			Level 2			Level 3		
1	Flan T5 XL	2.85B	10.81	7.03	8.52	1.39	0.73	0.96	0.37	0.17	0.24
2	Mistral v0.3	7B	19.98	24.81	22.14	2.83	3.33	3.06	0.73	0.81	0.77
3	Mistral Instruct v0.1	7B	21.32	48.89	29.69	9.90	23.01	13.85	3.40	7.81	4.73
4	Mistral Instruct v0.2	7B	21.57	44.71	29.10	11.21	21.98	14.85	4.34	8.13	5.66
5	Mistral Instruct v0.3	7B	19.65	50.92	28.36	10.15	24.99	14.43	3.59	8.77	5.09
6	BioMistral DARE [†] [3]	7B	21.29	27.80	24.12	10.21	12.28	11.15	2.94	3.28	3.10
7	Mistral Instruct v0.1	46.7B	19.82	48.49	28.14	10.74	24.61	14.95	4.20	9.37	5.80
8	LLaMA2	7B	11.53	11.36	11.45	3.54	3.33	3.43	0.71	0.65	0.68
9	LLaMA2 instruct	7B	19.56	35.32	25.18	10.05	16.28	12.43	3.28	5.13	4.00
10	Meditron [†] [8]	8B	19.13	21.32	20.17	2.57	2.52	2.55	0.65	0.59	0.62
11	Asclepius [†] [8]	8B	25.93	18.29	21.45	14.78	8.73	10.98	5.98	3.21	4.18
12	LLaMA3	8B	20.74	35.52	26.19	5.19	9.42	6.69	1.58	2.93	2.06
13	LLaMA3 Instruct	8B	20.95	56.41	30.56	9.52	27.04	14.09	3.89	10.86	5.73
14	LLaMA3.1 Instruct	8B	18.72	64.09	28.98	8.29	31.82	13.15	3.30	12.64	5.24
15	LLaMA3.2 Instruct	1B	17.35	50.21	25.79	7.71	23.17	11.57	2.85	8.43	4.26
16	LLaMA3.2 Instruct	3B	18.21	60.48	27.99	9.31	31.53	14.37	3.32	11.32	5.14
17	OpenBioLLM [†] [12]	8B	22.28	34.50	27.08	9.61	13.55	11.25	3.54	4.69	4.03
18	LLaMA3 Instruct	70B	20.86	58.59	30.77	8.82	26.42	13.23	3.73	10.86	5.56
19	GPT-3.5 turbo	—	21.14	55.40	30.60	11.22	28.95	16.17	4.13	10.47	5.92
20	GPT-4 turbo	—	18.47	57.06	27.91	10.01	29.19	14.91	3.84	11.37	5.74
21	GPT-4o	—	19.54	62.79	29.80	9.91	32.54	15.19	4.54	14.69	6.94
			Full code			Average					
1	Flan T5 XL	2.85B	0.00	0.00	0.00	3.14	1.98	2.43			
2	Mistral v0.3	7B	0.00	0.00	0.00	5.89	7.24	6.49			
3	Mistral Instruct v0.1	7B	0.17	0.76	0.28	8.70	20.12	12.14			
4	Mistral Instruct v0.2	7B	0.27	0.88	0.41	9.35	18.93	12.51			
5	Mistral Instruct v0.3	7B	0.26	1.11	0.42	8.41	21.45	12.07			
6	BioMistral DARE [†] [3]	7B	0.30	0.48	0.37	8.68	10.96	9.68			
7	Mixtral Instruct v0.1	46.7B	0.20	0.80	0.32	8.74	20.82	12.30			
8	LLaMA2	7B	0.00	0.00	0.00	3.95	3.84	3.89			
9	LLaMA2 Instruct	7B	0.19	0.50	0.27	8.27	14.31	10.47			
10	Meditron [†] [8]	7B	0.04	0.04	0.04	5.60	6.12	5.84			
11	Asclepius [†] [8]	7B	0.86	0.46	0.60	11.89	7.67	9.30			
12	LLaMA3	8B	0.28	1.46	0.47	8.66	23.94	12.71			
13	LLaMA3 Instruct	8B	0.23	1.19	0.39	8.80	24.08	12.88			
14	LLaMA3.1 Instruct	8B	0.22	1.61	0.39	7.63	27.54	11.94			
15	LLaMA3.2 Instruct	1B	0.15	0.77	0.26	7.02	20.64	10.47			
16	LLaMA3.2 Instruct	3B	0.14	0.88	0.24	7.74	26.06	11.93			
17	OpenBioLLM [†] [12]	8B	0.08	0.25	0.12	6.90	12.03	8.76			
18	LLaMA3 Instruct	70B	0.20	1.19	0.35	8.41	24.26	12.48			
19	GPT-3.5 turbo	—	0.27	1.19	0.44	9.19	24.00	13.28			
20	GPT-4 turbo	—	0.19	1.00	0.33	8.13	24.65	12.22			
21	GPT-4o	—	0.39	2.30	0.67	8.59	28.08	13.15			

Table 12: Lab tests ordering performance of LLMs.

#	Model	Param #	Pre.	Rec.	F1	Pre.	Rec.	F1	Pre.	Rec.	F1
			Level 1			Level 2			Level 3		
1	Flan T5 XL	2.85B	100.00	99.81	99.91	95.02	42.28	58.52	78.85	8.43	15.24
2	Mistral v0.3	7B	99.54	99.69	99.62	80.02	56.44	66.19	51.93	18.17	26.92
3	Mistral Instruct v0.1	7B	99.62	99.81	99.71	77.29	74.14	75.68	55.87	40.41	46.90
4	Mistral Instruct v0.2	7B	99.72	99.81	99.77	76.07	76.33	76.20	57.12	38.79	46.21
5	Mistral Instruct v0.3	7B	99.81	99.81	99.81	77.15	78.38	77.76	53.73	43.96	48.36
6	BioMistral DARE [†] [3]	7B	99.56	99.85	99.70	81.71	63.99	71.77	64.87	28.62	39.72
7	Mixtral Instruct v0.1	46.7B	99.81	99.81	99.81	77.85	75.41	76.61	53.76	38.92	45.16
8	LLaMA2	7B	99.56	99.78	99.67	90.37	46.32	61.25	55.13	11.87	19.54
9	LLaMA2 Instruct	7B	100.00	99.89	99.95	75.55	64.56	69.62	58.38	33.65	42.69
10	Meditron [†] [8]	7B	99.90	99.79	99.84	81.34	52.49	63.81	58.59	13.46	21.90
11	Asclepius [†] [8]	7B	99.80	99.80	99.80	84.88	43.83	57.81	56.77	12.18	20.06
12	LLaMA3	8B	97.23	99.80	98.50	73.77	73.03	73.40	39.07	36.52	37.75
13	LLaMA3 Instruct	8B	99.91	99.81	99.86	75.69	78.67	77.15	51.09	47.04	48.98
14	LLaMA3.1 Instruct	8B	100.00	99.81	99.91	74.10	81.01	77.40	47.82	49.56	48.67
15	LLaMA3.2 Instruct	1B	99.91	99.81	99.86	75.63	77.25	76.43	49.57	40.89	44.82
16	LLaMA3.2 Instruct	3B	100.00	99.81	99.91	73.62	81.14	77.20	46.80	50.08	48.38
17	OpenBioLLM [†] [12]	8B	99.81	99.62	99.72	78.97	67.69	72.89	53.56	28.42	37.14
18	LLaMA3 Instruct	70B	100.00	99.81	99.91	74.97	80.18	77.49	53.12	45.63	49.09
19	GPT-3.5 turbo	—	100.00	99.81	99.91	79.36	71.89	75.44	59.34	41.43	48.79
20	GPT-4 turbo	—	99.91	99.81	99.86	77.91	73.57	75.67	54.12	44.20	48.66
21	GPT-4o	—	99.91	99.81	99.86	77.93	74.86	76.37	53.02	44.85	48.60
			Level 4			Average					
1	Flan T5 XL	2.85B	49.72	1.86	3.58	80.90	38.10	44.31			
2	Mistral v0.3	7B	34.93	4.50	7.97	66.61	44.70	50.18			
3	Mistral Instruct v0.1	7B	31.71	10.68	15.97	66.12	56.26	59.57			
4	Mistral Instruct v0.2	7B	33.39	9.71	15.04	66.57	56.16	59.30			
5	Mistral Instruct v0.3	7B	30.42	11.48	16.67	65.28	58.41	60.65			
6	BioMistral DARE [†] [3]	7B	35.59	7.00	11.70	70.43	49.86	55.72			
7	Mixtral Instruct v0.1	46.7B	31.56	9.37	14.45	65.75	55.88	59.01			
8	LLaMA2	7B	57.56	3.81	7.15	75.66	40.45	46.90			
9	LLaMA2 Instruct	7B	24.93	7.06	11.01	64.71	51.29	55.82			
10	Meditron [†] [8]	7B	41.12	3.26	6.04	70.24	42.25	47.90			
11	Asclepius [†] [8]	7B	26.12	3.08	5.52	66.90	39.73	45.80			
12	LLaMA3	8B	26.30	11.53	16.03	59.09	55.22	56.42			
13	LLaMA3 Instruct	8B	29.98	13.68	18.78	64.17	59.80	61.19			
14	LLaMA3.1 Instruct	8B	28.26	15.31	19.86	62.55	61.42	61.46			
15	LLaMA3.2 Instruct	1B	31.34	10.65	15.90	64.11	57.15	59.25			
16	LLaMA3.2 Instruct	3B	32.33	16.48	21.83	63.19	61.88	61.83			
17	OpenBioLLM [†] [12]	8B	37.97	6.92	11.71	67.58	50.66	55.36			
18	LLaMA3 Instruct	70B	26.02	10.60	15.07	63.53	59.05	60.39			
19	GPT-3.5 turbo	—	33.69	11.76	17.43	68.10	56.22	60.39			
20	GPT-4 turbo	—	27.58	10.58	15.29	64.88	57.04	59.87			
21	GPT-4o	—	27.26	12.54	17.18	64.53	58.02	60.50			

Table 13: Prescriptions decision-making performance of LLMs.

#	Model	Param #	Pre.	Rec.	F1	Pre.	Rec.	F1	Pre.	Rec.	F1
			Level 1			Level 2			Level 3		
1	Flan T5 XL	2.85B	76.75	13.73	23.30	73.45	11.44	19.79	42.51	4.41	7.99
2	Mistral v0.3	7B	75.29	33.11	45.99	70.44	28.84	40.92	38.04	12.83	19.19
3	Mistral Instruct v0.1	7B	77.16	59.09	66.93	74.24	54.18	62.65	44.46	29.81	35.69
4	Mistral Instruct v0.2	7B	84.09	56.70	67.73	81.89	53.03	64.38	54.33	31.84	40.15
5	Mistral Instruct v0.3	7B	80.30	68.47	73.91	78.01	64.21	70.44	50.35	40.53	44.91
6	BioMistral DARE [†] [3]	7B	81.47	39.64	53.33	78.17	35.40	48.73	48.31	17.58	25.78
7	Mixtral Instruct v0.1	46.7B	84.52	66.17	74.23	82.23	63.32	71.55	57.80	41.98	48.64
8	LLaMA2	7B	79.88	23.74	36.60	76.04	20.54	32.35	34.63	7.59	12.45
9	LLaMA2 Instruct	7B	78.78	52.94	63.32	75.87	48.74	59.35	49.19	29.80	37.11
10	Meditron [†] [8]	7B	74.64	26.65	39.27	69.51	22.45	33.94	34.22	8.05	13.03
11	Asclepius [†] [8]	7B	86.32	17.14	28.60	84.27	14.73	25.08	50.70	6.60	11.68
12	LLaMA3	8B	76.22	54.90	63.82	72.85	50.72	59.81	41.68	31.00	35.55
13	LLaMA3 Instruct	8B	83.03	69.10	75.43	80.94	65.82	72.60	54.94	43.44	48.52
14	LLaMA3.1 Instruct	8B	79.37	79.26	79.32	77.19	76.30	76.74	48.97	52.13	50.50
15	LLaMA3.2 Instruct	1B	79.16	68.31	73.34	77.30	63.59	69.78	48.16	39.56	43.44
16	LLaMA3.2 Instruct	3B	77.24	74.89	76.05	75.12	70.94	72.97	44.57	45.65	45.10
17	OpenBioLLM [†] [12]	8B	84.89	45.42	59.18	81.99	41.21	54.85	56.89	24.11	33.87
18	LLaMA3 Instruct	70B	83.20	73.54	78.07	81.48	70.26	75.45	58.09	49.74	53.59
19	GPT-3.5 turbo	—	85.26	57.30	68.54	82.85	52.87	64.55	58.69	33.26	42.46
20	GPT-4 turbo	—	86.13	66.66	75.15	84.28	63.95	72.72	64.38	45.40	53.25
21	GPT-4o	—	84.93	70.58	77.09	83.05	67.46	74.45	62.58	49.43	55.24
			Level 4			Average					
1	Flan T5 XL	2.85B	27.56	2.42	4.45	55.07	8.00	13.88			
2	Mistral v0.3	7B	27.28	8.50	12.96	52.76	20.82	29.76			
3	Mistral Instruct v0.1	7B	33.64	21.54	26.27	57.38	41.16	47.88			
4	Mistral Instruct v0.2	7B	45.25	24.30	31.62	66.39	41.47	50.97			
5	Mistral Instruct v0.3	7B	40.94	30.87	35.20	62.40	51.02	56.11			
6	BioMistral DARE [†] [3]	7B	38.73	12.77	19.21	61.67	26.35	36.76			
7	Mixtral Instruct v0.1	46.7B	49.14	32.96	39.46	68.42	51.11	58.47			
8	LLaMA2	7B	26.21	5.33	8.86	54.19	14.30	22.56			
9	LLaMA2 Instruct	7B	40.38	23.30	29.55	61.06	38.69	47.33			
10	Meditron [†] [8]	7B	24.60	5.11	8.46	50.74	15.56	23.68			
11	Asclepius [†] [8]	7B	39.60	4.81	8.57	65.22	10.82	18.49			
12	LLaMA3	8B	32.58	23.92	27.59	55.83	40.13	46.69			
13	LLaMA3 Instruct	8B	46.43	34.57	39.63	66.33	53.23	59.04			
14	LLaMA3.1 Instruct	8B	39.63	40.90	40.26	61.29	62.15	61.70			
15	LLaMA3.2 Instruct	1B	38.50	30.43	33.99	60.78	50.47	55.14			
16	LLaMA3.2 Instruct	3B	34.18	34.21	34.20	57.78	56.42	57.08			
17	OpenBioLLM [†] [12]	8B	48.18	18.76	27.00	67.99	32.38	43.72			
18	LLaMA3 Instruct	70B	50.82	40.74	45.22	68.40	58.57	63.08			
19	GPT-3.5 turbo	—	50.38	26.51	34.74	69.30	42.48	52.57			
20	GPT-4 turbo	—	58.00	37.83	45.79	73.20	53.46	61.73			
21	GPT-4o	—	55.45	41.42	47.42	71.50	57.22	63.55			