

# Unveiling Performance Challenges of Large Language Models in Low-Resource Healthcare: A Demographic Fairness Perspective

Yue Zhou Barbara Di Eugenio Lu Cheng

University of Illinois Chicago

{yzhou232, bdieugen, lucheng}@uic.edu

## Abstract

This paper studies the performance of large language models (LLMs), particularly regarding demographic fairness, in solving real-world healthcare tasks. We evaluate state-of-the-art LLMs with three prevalent learning frameworks across six diverse healthcare tasks and find significant challenges in applying LLMs to real-world healthcare tasks and persistent fairness issues across demographic groups. We also find that explicitly providing demographic information yields mixed results, while LLM's ability to infer such details raises concerns about biased health predictions. Utilizing LLMs as autonomous agents with access to up-to-date guidelines does not guarantee performance improvement. We believe these findings reveal the critical limitations of LLMs as concerns healthcare fairness and the urgent need for specialized research in this area. **WARNING:** This paper contains model outputs that may be considered offensive in nature.

## 1 Introduction

The application of Artificial Intelligence (AI) in healthcare is almost as old as AI itself<sup>1</sup>. Over the years, the penetration of AI techniques in healthcare has increased, from early expert systems like MYCIN (Shortliffe, 1976) to NLP techniques applied to clinical notes (Friedman et al., 1999) to the current proliferation of applications of Large Language Models (LLMs) (He et al., 2024). The assumption is that LLMs will be equally successful in healthcare as they have been in other domains (Srivastava et al., 2023; Rae et al., 2022; Liang et al., 2023), especially given emerging learning frameworks such as chain-of-thought, parameter-efficient fine-tuning, and LLM as autonomous agents to address in-context reasoning, data scarcity and factual

knowledge (Wei et al., 2022; Kojima et al., 2022; Zhou et al., 2024c; Yao et al., 2023; Shinn et al., 2023; Wang et al., 2024c).

Healthcare applications present unique challenges due to the complexity of knowledge involved, limited data resources, and inherent ethical considerations, including how to mitigate health disparities and achieve health equity (Pereira, 1993; LaVeist, 2005; Waters, 2000; Braveman, 2006; Lane et al., 2017; Ndugga and Artiga, 2021). While recent studies have begun exploring LLMs in medical QA, bio-medicine understanding, and disease diagnosis (Singhal et al., 2023; Tian et al., 2023; Zhou et al., 2024a; He et al., 2024; Wang et al., 2024b), there remains a significant gap in comprehensive evaluations of LLM performance on real-world healthcare tasks, particularly as concerns their potential to reinforce health disparities, a crucial consideration given LLM known biases and their potential impact on patient care (Schick et al., 2021; Weidinger et al., 2021; Huang et al., 2024; Gallegos et al., 2024).

To address this gap and provide insights into best practices for utilizing LLMs on low-resource healthcare tasks, we present a comprehensive study examining the performance of LLMs across diverse healthcare benchmarks. Concretely, we formulate six benchmarks, including mortality, readmission, health coaching outcome prediction, and mental health diagnosis. We evaluate three state-of-the-art LLMs, GPT-4 (OpenAI, 2023), Claude-3 (Anthropic, 2024), and LLaMA-3 (AI@Meta, 2024), with prevalent frameworks: in-context learning with chain-of-thought reasoning (Wei et al., 2022; Kojima et al., 2022), parameter-efficient fine-tuning (Hu et al., 2022; Dettmers et al., 2023), and LLM-as-agent leveraging external factual knowledge. We employ two standard fairness metrics, Demographic Parity Difference (DPD) and Equal Opportunity Difference (EOD) (Zemel et al., 2013; Wang et al., 2024a; Liu et al., 2023) to quantify dis-

<sup>1</sup>The first issue of the journal *AI in medicine* is from 1989; <https://www.sciencedirect.com/journal/artificial-intelligence-in-medicine/>

parities across racial and gender groups, offering insights into the models’ potential biases.

Our findings show significant challenges in applying LLMs to real-world healthcare tasks. Contrary to their success in other domains, LLMs struggle to achieve high accuracy across our benchmarks, with some models barely surpassing random guessing. We also observe persistent fairness issues, with considerable disparities in performance across demographic groups, particularly regarding ethnicity. Notably, explicitly prompting LLMs with demographic information yields mixed results and does not consistently improve either prediction performance or fairness. We also explore LLM ability to infer demographic information from conversations and find that LLMs can deduce demographic details with serious biases, raising concerns about their potential influence on health predictions. Finally, we show that access to up-to-date guidelines and factual information does not guarantee accurate predictions in healthcare scenarios.

## 2 Related Work

**LLMs in Healthcare.** This domain has recently seen a surge in the application of LLMs: surveys such as Zhou et al. (2024a); Clusmann et al. (2023); Tian et al. (2023) discuss the current applications and future landscape of LLMs in medicine; He et al. (2024) offers a review of healthcare data and applications with LLMs. Google proposed PalmMed2 (Singhal et al., 2023), an LLM in the medical domain. Wang et al. (2024b) explore utilizing LLMs for rare case diagnosis. Hu et al. (2024); Monajatiipoor et al. (2024) study named entity recognition with LLMs in clinical and biomedicine settings. However, limited work exists on demographic fairness in LLMs across multiple healthcare applications.

**LLMs and Fairness.** Large language models have demonstrated considerable in-context learning abilities (Wei et al., 2022; Kojima et al., 2022) and parameter-efficient fine-tuning adaptability such as Low-Rank Adaptation (Hu et al., 2022; Dettmers et al., 2023). Recently, the use of LLMs as autonomous agents equipped with tool usage capabilities shows promising results (Yao et al., 2023; Shinn et al., 2023; Wang et al., 2024d). Nonetheless, LLMs can exhibit limitations on generating unbiased and faithful output, with performance deterioration among underrepresented groups (Rae et al., 2022; Srivastava et al., 2023; Liang et al.,

2023; Schick et al., 2021; Weidinger et al., 2021; Huang et al., 2024; Gallegos et al., 2024).

## 3 Datasets and Task Formulation

To facilitate evaluating LLM performance in healthcare, especially as concerns demographic fairness, we formulate six tasks based on four healthcare datasets containing demographic information, such as age, gender, and ethnicity. We point out that public healthcare datasets with demographic information are scarce due to potential ethical and privacy concerns. We employ MEDQA (publicly available), MIMIC, which is available upon request, and two others that are either partially available (a subset of the *Health Coaching Dataset* is available, but not the demographic information) or not available (*Bipolar Disorder and Schizophrenia Interviews*). Our results on the publicly available ones should attest to the generalizability of our approach.

**MIMIC-IV** (Johnson et al., 2020) is a large, publicly available healthcare dataset containing hospital health records. Taking the patients’ clinical notes as input, we formulate two tasks: **one-year mortality** prediction and **90-day readmission** prediction, which emulate patient outcome prediction in real-world settings. The input note and label pairs are created by joining three tables in the MIMIC database: ‘patients,’ ‘admissions,’ and ‘discharge.’ The note mainly contains sections including chief complaint, history of present illness, past medical history, and lab results. We removed the discharge instruction from the note in the input to our models since they can reveal direct information on mortality/readmission (e.g., Hospice - hospice care is provided to patients who are approaching the end of life and have stopped treatment).

**Health Coaching Datasets** Dialogues are inherently unique as concerns fairness evaluation since they often contain implicit demographic cues. This raises questions about how LLMs handle these subtleties and whether their responses might exhibit unfairness. The health coaching datasets (Gupta et al., 2020; Zhou et al., 2024b) comprise SMS conversations between patients and certified health coaches over several weeks, focusing on creating and accomplishing S.M.A.R.T. goals to promote health behavior changes (Doran, 1981). Each week, the conversation starts with a *goal setting stage*, where the coach and the patient discuss and cre-

ate a concrete and measurable goal for physical activities. Then, the coach follows up on the patient’s progress and maintains engagement, which is called the *goal implementation stage*. We formulate another patient **outcome prediction** task, which predicts whether the patient will accomplish the next goal based on the dialogue history over the past two weeks and the current goal-setting stage. <sup>2</sup>

**Bipolar Disorder and Schizophrenia Interviews** (Aich et al., 2022) contains transcribed interactions between a trained clinician and outpatients with schizophrenia, bipolar disorder, and healthy controls. These interactions pertain to two situations: (1) **Meeting New Neighbor**: the participant is asked to imagine an affiliative scene<sup>3</sup> and converse with the interviewee (role-playing as the new neighbor) as they have just moved into the neighborhood. (2) **Complaining to a Landlord**: A confrontational scene where the participant role-plays the tenant and complains to the landlord (role-played by the interviewee) about issues such as pipe leakage. These contrasting scenarios aim to assess mental status in both friendly and stressful situations. We utilize both scenes and ask the model to predict the outpatients’ cohort based on the conversations. We aim to test the LLM diagnosis ability in real-world settings and whether they can potentially identify linguistic cues to mental illness in the conversations. This task, again, aims to investigate how LLMs process and respond to implicit demographic cues within dialogues and introduce potential unfairness in their outputs.

**Medical Question Answering Dataset** contains multiple-choice medical problems collected from medical board exams (Jin et al., 2021). We use the English subsection adapted from the USMLE (United States Medical Licensing Examination) and only select the test problems targeting specific ethnic groups. We use this dataset to study the performance inequality in LLMs in typical-case diagnoses for specific demographic subgroups. While commonly used to evaluate LLM performance in healthcare, this dataset can understate the challenges of real-world medical data. We hope it can serve as a comparison, highlighting the gap

<sup>2</sup>Conversations and patient demographics are available at: [github.com/uic-nlp-lab/virtualcoachdata](https://github.com/uic-nlp-lab/virtualcoachdata)

<sup>3</sup>Affiliative means “relating to the formation of social and emotional bonds with others or to the desire to create such bonds”, the Merriam-Webster dictionary.

Task	SFT	Test
MIMIC-Mortality	5000	500
MIMIC-Readmission	5000	500
Health Coaching	120	60
MedQA	5000	175
SCZ and Bipolar: Neighbor	190	261
SCZ and Bipolar: Landlord	188	261

Table 1: Data splits for fine-tuning (SFT) and testing of the tasks.

between controlled benchmarks and the complexity of real-world healthcare applications.

Figure 1 provides data samples from the six healthcare benchmarks. Statistics on data split are shown in Table 1. The small size of the Health Coaching and Schizophrenia/Bipolar datasets reflect the inherent scarcity of authentic data prevalent in patient-facing healthcare applications. For tasks with access to larger databases (MIMIC and MedQA), we intentionally capped the training data at 5000 examples to mimic similar low-resource conditions. We sample each dataset such that the classes  $C$ , the demographic attributes  $Z$ , and the distribution  $P(C = c|Z = z)$  are roughly balanced.

We find it challenging to gather publicly available healthcare corpora where demographic information is also available, due to ethical concerns and privacy issues in the healthcare domain. We posit that there is a need for more open-source, well-structured healthcare data to facilitate research on AI fairness.

## 4 Baselines

Our study aims to evaluate the performance of trendy frameworks utilizing large language models in real-world, low-resource healthcare settings. We seek to provide insights into best practices for leveraging LLMs when building applications under these constraints. To this end, we evaluate the performance of LLMs using three representative frameworks:

- **In-Context Learning (ICL) with Chain-of-Thought** enhances LLM inherent reasoning capabilities by prompting the LLM to provide a step-by-step reasoning chain (Wei et al., 2022; Kojima et al., 2022; Zhou et al., 2024c). We implement two schemes: (1) **Zero-Shot Chain-of-Thought (CoT)**, which appends “*Let’s think step by step.*” to the question text; and (2) **N-Shot CoT**, where we append four to eight-shot in-context examples

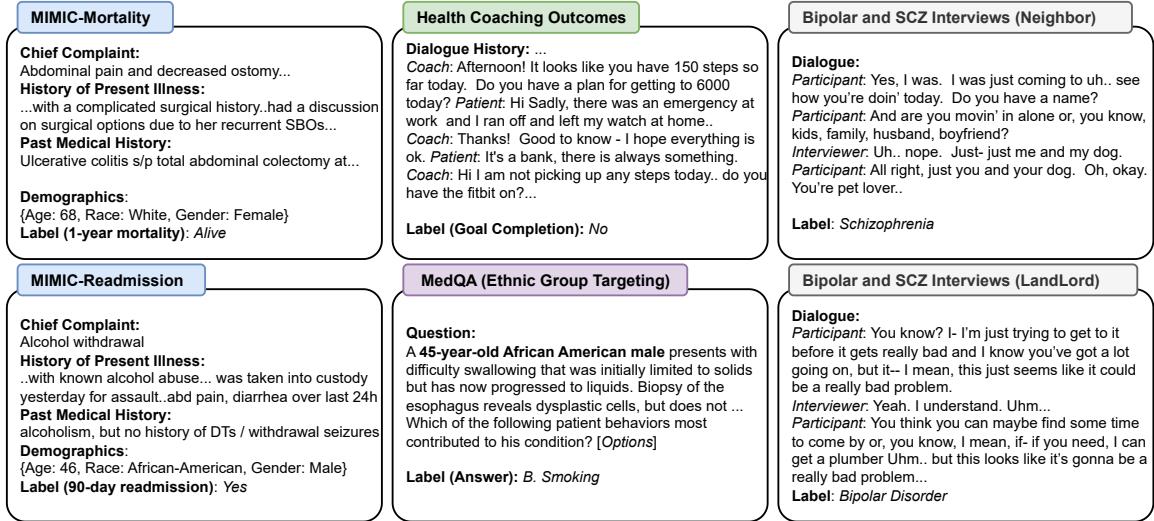


Figure 1: Overview of the six health benchmarks, with illustrative examples.

with CoT to the LLM when solving the problems. The examples are demographically balanced following (Wang et al., 2024a). We report the best performance between Zero-Shot and N-Shot in this setting. Since previous work found weak evidence on prompting demographic information to improve fairness (Wang et al., 2024a), we provide baselines with and without explicit demographic information in the input in this framework.

• **Parameter Efficient Fine-Tuning (PEFT)** is a set of techniques that adapt pre-trained language models to downstream tasks by updating only a small portion of the parameters, reducing computational costs and storage requirements while maintaining performance. Given the resource constraints, we specifically fine-tune LLaMA-3 with **Low-Rank Adaptation (LoRA)** (Hu et al., 2022), where the pre-trained model weights are untouched, yet small-scale trainable rank decomposition matrices are injected.

• **LLM as Agents** enhances LLMs with specialized modules for planning and tool usage, enabling them to solve complex tasks beyond pre-trained knowledge. In the paper, we propose a simplified pipeline based on **ReAct** (Yao et al., 2023) and **Reflexion** (Shinn et al., 2023). Concretely, for each task, we prompt the LLM to web search for the latest guidelines for analyzing the underlying [example] on the [task]. Then, we prompt the LLM to generate a concise guide based on the retrieved top 10 most relevant Google search results. Finally, the LLM generates predictions based on the question, the example, and the generated guide. We additionally allow a maximum of two

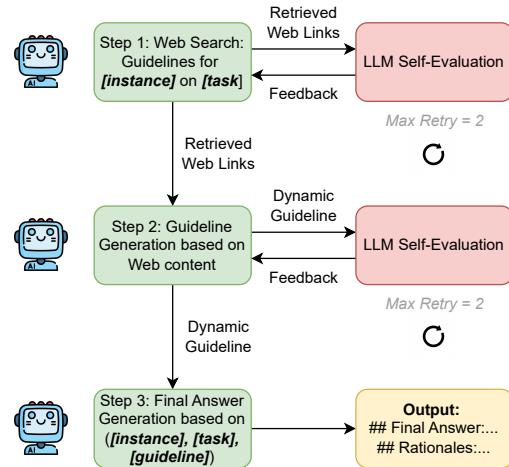


Figure 2: An overview of the LLM as Agent framework, including search and generating guidelines for the underlying data instance and task, and generating the final answer referencing the guidelines.

re-attempts for the first two steps based on LLM’s self-evaluation. Figure 2 shows an overview of the framework.

## 5 Metrics for Fairness

Existing literature predominantly adopts two metrics to evaluate the demographic fairness of the model prediction (Zemel et al., 2013; Wang et al., 2024a; Liu et al., 2023). The first metric is called *Statistical Parity* or *Demographic Parity*. Statistical parity is achieved when favorable decision outcomes are unrelated to the protected attributes.

The rationale is to test whether the model treats various subgroups similarly. Take fraud detection as an example; the model should output “good credit” with a similar chance for both males and females. Note that it does not consider the ground truth label. Consider the sensitive/demographic attribute  $Z$  and the predicted outcome  $\hat{Y}$ , the (one-vs-all) **Demographic Parity Difference (DPD)** for subgroup  $z_i$  can be defined as:

$$PDP = P(\hat{Y} = 1|Z = z_i) - P(\hat{Y} = 1|Z \neq z_i)$$

This metric may pose challenges when assessing model performance in healthcare applications, as the attribute  $Z$  could be a prior factor influencing model predictions. For instance, when predicting a patient’s one-year mortality, age may significantly influence risk, with individuals above the age of 90 facing greater risk compared to those below. Consequently, an LLM which obtained such knowledge during pre-training may be more likely to predict mortality for patients above 90 years old. Nonetheless, we include this metric as it provides valuable insights into the model’s prediction tendencies across different demographic groups in healthcare contexts and is crucial for understanding potential biases.

The second metric, *Equality of Opportunity*, evaluates model fairness based on the ground truth labels. It indicates that different subgroups should have an equal likelihood of being accurately classified by the model. One way to formulate the metric is to measure the true positive rates of class  $Y$  across various subgroups. We report the **Equal Opportunity Difference (EOD)** as:

$$EOD = P(\hat{Y} = 1|Y = 1, Z = z_i) - P(\hat{Y} = 1|Y = 1, Z \neq z_i)$$

Note the definition of favorable attributes in healthcare is more nuanced than in other domains like fraud detection or tweet classification. While “good credit” or “non-toxic” are straightforward favorable attributes in those fields, healthcare scenarios often have context-dependent favorable classes. However, for clarity and consistency in our main experiments, we define favorable attributes by *any positive health indicators* across different tasks. These include, for example, Low Mortality Risk out of {Low Mortality Risk, High Mortality Risk} in mortality prediction and Healthy Control group out of

{SCZ, Bipolar, Healthy Control} in SCZ and Bipolar Interviews. We report global accuracy and accuracy per demographic group instead of PDP and EOD for the MedQA task, which involves open-ended question answering.

## 6 Experiments

In this section, we describe our experiment results evaluating the effectiveness of LLMs in solving real-world healthcare tasks with various frameworks and settings, as well as additional discussions on demographic awareness and qualitative examples.

### 6.1 Experimental Settings

**Language Models** We utilize three state-of-the-art large language models for evaluation, including two closed-source models, OpenAI GPT-4 (OpenAI, 2023) and Claude-3 (Sonnet) (Anthropic, 2024), and one open-source model, LLaMA-3 (8b) (AI@Meta, 2024). In compliance with the responsible use guidelines for MIMIC data with online services, we utilize the Azure OpenAI service for GPT-4 and opt out of human data review.<sup>4</sup> We have also ensured that our usage of Claude-3 adheres to the stipulated agreement.<sup>5</sup> The LLaMA-3 model is run locally on our machines. The Schizophrenia and Bipolar dataset is the only dataset in our study that is not publicly available.<sup>6</sup>

**Implementation Details** For fine-tuning, we employed LoRA with a rank of 8 across all trainable layers. We use a dropout rate of 0.1, a learning rate  $1e^{-5}$ , and a batch size of 8 for all experiments. Our implementation adheres to the recommendations outlined in QLoRA (Dettmers et al., 2023), except for the LoRA scaling factor (Alpha), which is set equal to the LoRA rank. We choose the temperature  $T = 0.3$  for all three language models for inference. The full implementation details and prompt templates used in the experiments are available in Appendix A.

### 6.2 Main Results

Table 2 shows the accuracy results for six healthcare tasks using different LLM frameworks. The numbers outside parentheses represent accuracy without explicit demographic information as input,

<sup>4</sup>Responsible use of MIMIC data

<sup>5</sup>Anthropic’s data usage policy

<sup>6</sup>Access is regulated by the Institutional Review Board (IRB) of the University of California San Diego.

while those inside parentheses show results when demographic information is explicitly prompted to the LLMs. There are several key observations:

1. Despite their impressive performance in various domains, LLMs struggle with real-world healthcare tasks across all prevalent frameworks. Many implementations in Readmission, Neighbor, and Landlord barely surpass random guess baselines. The claims that LLMs can easily solve classification tasks with few examples are unsupported and inconsistent with our findings regarding real-world healthcare applications.
2. While closed-source, large-scale LLMs generally outperform open-source, smaller models in in-context learning, the most effective framework varies by task. For instance, the "Schizophrenia and bipolar" diagnosis tasks achieve the best results with fine-tuning despite many fewer training examples compared to MIMIC-based tasks. In contrast, for the MIMIC and health coaching tasks, in-context learning achieves the best performance.
3. The LLM-as-Agent approach shows mixed results across tasks. It excels in MedQA, presumably due to its ability to search online for open-book guidelines for USMLE questions. However, it underperforms in real-world healthcare applications despite generating seemingly convincing thought processes. The following subsection will present qualitative examples to illustrate these findings.
4. Explicitly prompting LLMs with demographic information does not necessarily improve performance. The impact varies depending on both the specific task and the LLM used.

We next present further demographic fairness results. Table 3 shows Demographic Parity Difference (DPD) and Equal Opportunity Difference (EOD) across six tasks, with results inside parentheses indicating explicit demographic prompts. For brevity, we mainly focus on White vs. African American and Female vs. Male comparisons. The PDP/EOD metrics are calculated as (African American - White) for race and (Female - Male) for gender, and the +/- indicates the sign of the difference. To interpret the results, consider the Demographic Parity Difference (DPD) and Equal Opportunity Difference (EOD) metrics. A DPD value of -8.2

for African Americans in the mortality task indicates that the model is 8.2% less likely to predict a favorable outcome (e.g., low mortality risk) for this group compared to White patients, regardless of the ground truth. Similarly, an EOD value of -3.5 for African Americans signifies that the model's true positive rate in predicting favorable outcomes is 3.5% lower for this group, highlighting a performance disparity. There are several key observations:

1. Unfairness exists across all tasks, frameworks, and demographics, with racial disparities more prominent than gender disparities.
2. LLMs consistently predict less favorable outcomes for African American patients, while a lower Equality of Opportunity for African Americans is observed in most tasks, except health coaching.
3. Explicitly prompting demographic information yields mixed results on fairness. DPD mostly improves for GPT-4 but not for other models. EOD is less influenced by demographic prompts compared to DPD.
4. Fine-tuning's impact on fairness varies by task, improving for some (readmission, neighbor, landlord) while worsening for others (health coaching, mortality). The agent approach can mitigate unfairness in certain cases. Note that LLaMA's all-zero results for landlord and neighbor tasks stem from blindly predicting all participants as having schizophrenia.

Finally, all LLMs exhibit a discrepancy in in-context learning performance in MedQA targeting different racial groups. See Table 4.

We additionally use the mortality dataset as an example to showcase fairness results across diverse demographic subgroups in gender, age, and race. We can observe that LLMs predict high mortality risks for the geriatric age group and African Americans and lower prediction performance for these groups. See Appendix B.

### 6.3 Additional Studies

**Demographic Awareness and Biases** An intriguing question arises regarding the risk that LLMs infer irrelevant or wrong demographic information from conversational data and hence, that they introduce bias into health outcome predictions. To

Baselines	Backbones	Mortality	Readmission	MedQA	HealthCo	Neighbor	Landlord
Random Guess	-	50.0	50.0	20.0	50.0	33.3	33.3
ICL/Chain-of-Thought	GPT-4	77.0 ( <b>79.0</b> )	55.3 (56.8)	- (68.6)	76.7 (76.7)	41.0 (39.1)	38.3 (36.8)
ICL/Chain-of-Thought	Claude-3	72.6 (50.8)	52.9 ( <b>57.6</b> )	- (65.7)	<b>80.0</b> ( <b>80.0</b> )	37.5 (38.3)	37.2 (34.1)
ICL/Chain-of-Thought	LLaMA-3	73.0 (72.2)	53.3 (55.3)	- (59.4)	70.0 (70.0)	33.0 (33.3)	34.1 (33.7)
Supervised Fine-Tuning	LLaMA-3	68.1	48.2	71.4	70.0	<b>49.4</b>	<b>41.4</b>
LLM as Agent	GPT-4	66.7	53.3	<b>86.9</b>	70.0	36.8	36.8

Table 2: Global accuracy across the six tasks with various LLM frameworks. ICL Results with and without explicit demographic prompts are inside and outside parentheses, respectively.

Task	Setting	DPD		EOD	
		African-American - White	Female - Male	African-American - White	Female - Male
Mortality	COT-GPT-4	-8.2 (-7.4)	-5.1 (-2.7)	-3.5 (-4.4)	-4.6 (-5.3)
	COT-Claude-3	-11.5 (-6.4)	-8.7 (+1.5)	-3.8 (-3.8)	-4.5 (-7.2)
	COT-LLaMA-3	-8.3 (-17.6)	-2.8 (-3.0)	-4.2 (-8.0)	-5.9 (-5.2)
	SFT (LoRA)	-14.3	-12.7	-9.4	-5.3
	LLM as Agent	+2.1	-2.1	+2.4	+2.8
Readmission	COT-GPT-4	-10.5 (-7.1)	-0.5 (-0.6)	-9.8 (-7.1)	-1.9 (-1.1)
	COT-Claude-3	-7.9 (-12.2)	-2.8 (+6.4)	-3.1 (-2.6)	-6.1 (-1.6)
	COT-LLaMA-3	-9.3 (-26.7)	-1.3 (-0.9)	-0.5 (-7.3)	-2.3 (-3.1)
	SFT (LoRA)	+0.5	+0.6	+4.9	+1.1
	LLM as Agent	-12.1	+4.2	-6.8	+4.6
Neighbor	COT-GPT-4	-7.5 (-5.7)	+3.8 (+3.6)	-1.5 (-1.5)	-9.0 (-9.0)
	COT-Claude-3	-7.9 (-7.3)	+4.1 (-3.2)	-0.1 (-0.9)	-7.2 (-9.7)
	COT-LLaMA-3	0.0 (0.0)*	0.0 (0.0)*	0.0 (0.0)*	0.0 (0.0)*
	SFT (LoRA)	-4.8	+7.1	-4.5	+6.5
	Agent (LLaMA)	-6.7	-3.6	-2.5	-8.6
Landlord	COT-GPT-4	-10.3 (-4.7)	+8.5 (+6.1)	-0.7 (0.0)	-9.9 (-8.9)
	COT-Claude-3	-4.1 (-20.2)	+0.8 (+0.9)	+2.4 (-4.9)	-11.4 (-10.4)
	COT-LLaMA-3	0.0 (0.0)*	0.0 (0.0)*	0.0 (0.0)*	0.0 (0.0)*
	SFT (LoRA)	-0.8	-0.1	+2.5	-4.7
	Agent (LLaMA)	-5.2	+3.7	+2.2	-8.7
HealthCoaching	African-American - Hispanic		Female - Male	African-American - Hispanic	
	COT-GPT-4	-4.2 (-4.2)	-4.3 (-4.3)	+15.0 (+15.0)	-16.1 (-16.1)
	COT-Claude-3	-8.3 (-8.3)	-8.7 (-8.7)	+15.0 (+15.0)	-16.1 (-16.1)
	COT-LLaMA-3	-12.5 (-12.5)	+5.6 (+5.6)	-6.7 (-6.7)	-6.2 (-6.2)
	SFT (LoRA)	-13.3	-1.8	+18.3	-19.2
	LLM as Agent	-4.2	+14.3	+10.8	-1.8

Table 3: Demographic Parity Difference (DPD) and Equal Opportunity Difference (EOD) across the six tasks. Results with and without explicit demographic prompts are inside and outside parentheses, respectively. The PDP/EOD metrics are calculated as (African American - White) for race and (Female - Male) for gender, and the +/- indicates the sign of the difference.

Baselines	Accuracy			
	African-American	Asian	Hispanic	White
COT-GPT-4	62	60	68	84
COT-Claude-3	64	62	60	74
COT-LLaMA-3	60	56	52	66
SFT (LoRA)	86	62	52	76
Agent (GPT-4)	86	86	88	88

Table 4: Accuracy across different race-targeting MedQA problems.

explore this, we ask three LLMs to identify the patient’s race, given the conversation. Table 6 shows that GPT-4 achieved the highest accuracy in predicting race information, followed by LLaMA-3, outperforming random guessing. Notably, Claude-3 refused to make race predictions, stating, “*I would prefer not to speculate about the participant’s race*

*or ethnicity.*” GPT-4 exhibited similar reluctance for the health coaching dataset. However, refusal to answer does not necessarily indicate an inability to do so or a lack of potential biases. Table 5 provides a qualitative example of LLaMA’s predictions with rationales, which heavily contain stereotypical and biased assumptions (even those based on AAV (African American Vernacular) characteristics are incorrect, according to a sociolinguistics expert on AAV we consulted<sup>7</sup>). While the model appears to predict schizophrenia (SCZ) based on unbiased criteria independent of race, it is unclear whether racial biases subtly influence these diagnoses. In fact, we observed a high Demographic

<sup>7</sup>Details of this consultation are reported in Appendix C.

Correct	Output	Task	Model
Healthy	<p>1. Linguistic features: The participant's language use, such as the use of "ain't" and "gonna" is more commonly associated with African American Vernacular English (AAVE).</p> <p>2. Tone and demeanor: The participant's tone and demeanor throughout the conversation, such as <b>expressing frustration and anger</b> when discussing the leak in their apartment, are consistent with the cultural norms of African American individuals.</p> <p>3. Language patterns: ..such as using short sentences and a more informal tone, are also AAVE.</p> <p>4. Lack of formal language use: ..does not use formal language or attempt to articulate their concerns in a more sophisticated manner, is consistent with the cultural norms of some African American individuals.</p>	Race Inference	LLaMA-3
Healthy	<p><b>Answer:</b> Schizophrenia</p> <p><b>Rationale:</b> 1. Disorganized and illogical speech patterns...2. Lack of coherence and clarity...3. Irritability and agitation... 4. Disregard for social norms and conventions...</p>	Diagnosis	LLaMA-3
C. 46XX	<p><b>Answer:</b> C. 46XX</p> <p><b>Rationale:</b> The clinical presentation and ultrasound findings are indicative of a molar pregnancy. The markedly elevated HCG levels and the "snowstorm" appearance on ultrasound strongly suggest a complete molar pregnancy. Complete molar pregnancies most commonly have a 46XX karyotype, resulting from the fertilization of an empty ovum by ...</p> <p><b>Retrieved Links:</b> <a href="https://www.ncbi.nlm.nih.gov/books/NBK459155/">https://www.ncbi.nlm.nih.gov/books/NBK459155/</a>  <a href="https://step2.medbullets.com/gynecology/121751/hydatidiform-mole">https://step2.medbullets.com/gynecology/121751/hydatidiform-mole</a></p>	QA	LLM as Agent
Bipolar Disorder	<p><i>Reasoning through retrieved guidelines for SCZ (Bold Text)</i></p> <ul style="list-style-type: none"> <li>- <b>Disorganized Speech:</b> The participant repeats phrases like "it's getting worse"</li> <li>- <b>Reduced Complexity:</b> straightforward and repetitive sentence structures, such as "I'm- I'm- I'm gonna take"</li> <li>- <b>Limited Vocabulary:</b> repeating the same words like "worse," "leaking," and "important."</li> <li>- <b>Poverty of Speech:</b> The participant's responses are often brief and lack depth, such as "Huh?"</li> </ul> <p><b>Answer:</b> Schizophrenia</p>	Diagnosis	LLM as Agent

Table 5: Qualitative examples of model outputs on health outcome prediction and race inference.

	HealthCoaching	Neighbor	Landlord
Random	33.3	50.0	50.0
GPT-4	Refusal	75.2	78.8
Claude3	Refusal	Refusal	Refusal
LLaMA-3	40.0	59.6	61.7

Table 6: Accuracy of predicting patient's race from conversations.

Parity Difference in schizophrenia predictions for African-American individuals. This raises serious concerns about the underlying biases that LLMs may possess and how these biases could influence healthcare-related predictions and decision-making processes.

**Agent and Factual Knowledge Retrieval in Healthcare** One potential advantage of using LLMs as agents equipped with tool usage capabilities is to retrieve external facts and knowledge to guide predictions rather than relying solely on potentially hallucination-prone pre-trained knowledge. This approach yields impressive performance in solving MedQA questions. We hypothesize that the performance stems from the LLM web search for direct guidelines, particularly for questions that require memorization rather than complex reasoning. Table 5 provides an example where the LLM agent directly located guideline links.

However, our findings suggest that access to up-to-date guidelines and factual information does not necessarily guarantee accurate final predictions. Table 5 illustrates a factual guideline retrieved from the latest research yet incorrect reasoning from linguistic cues. The LLM erroneously overemphasized fragments and other speech patterns and thus predicted the patient as having schizophrenia, failing to account for the fact that these were spoken dialogue transcripts (despite this being explicitly stated in the prompt). This example highlights the challenges in applying retrieved knowledge appropriately and the potential for misinterpretation even when given access to current and factual information in healthcare.

## 7 Conclusions and Future Work

We explored LLM performance and demographic fairness across diverse healthcare tasks. Our experimental results highlight LLM difficulties in solving real-world healthcare tasks and significant disparities across demographic groups. We showed that explicitly providing demographic information to LLMs does not guarantee improved performance or fairness. While LLM-as-agent can retrieve factual knowledge, the reasoning processes may still lead to inaccurate conclusions. We also observed potential biases in race inference, which could influence

health outcome predictions. These findings highlight the urgent need for future research addressing LLM fairness and reliability in healthcare.

## Limitations

While our study provides valuable insights into the performance disparities of LLMs across demographic groups in healthcare tasks, we still need to identify a systematic strategy to mitigate biases and improve fairness across different healthcare scenarios. Additionally, investigating the relationship between LLM demographic awareness and the fairness of health outcome predictions could yield noteworthy insights, particularly given our observations of biases in demographic inference and demographic parity differences.

## Ethics Statement

This study examined LLM performance across diverse healthcare tasks, revealing disparities in accuracy and fairness among demographic groups. Our findings highlight the potential risks of propagating LLM unfairness and the need for rigorous fairness assessments and bias mitigation in healthcare. Future work should focus on developing LLMs for equitable outcomes in healthcare.

## Acknowledgement

The work is supported by the National Science Foundation under Grants IIS-1838770 and IIS-2312862.

## References

- Ankit Aich, Avery Quynh, Varsha Badal, Amy Pinkham, Philip Harvey, Colin Depp, and Natalie Parde. 2022. [Towards intelligent clinically-informed language analyses of people with bipolar disorder and schizophrenia](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 2871–2887, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- AI@Meta. 2024. [Llama 3 model card](#).
- Anthropic. 2024. [The claude 3 model family: Opus, sonnet, haiku](#).
- Paula Braveman. 2006. Health disparities and health equity: concepts and measurement. *Annual Review of Public Health*, 27(1):167–194.
- Jan Clusmann, Fiona R Kolbinger, Hannah Sophie Muti, Zunamys I Carrero, Jan-Niklas Eckardt, Narmin Ghaffari Laleh, Chiara Maria Lavinia Löffler, Sophie-Caroline Schwarzkopf, Michaela Unger, Gregory P Veldhuizen, et al. 2023. The future landscape of large language models in medicine. *Communications medicine*, 3(1):141.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. [Qlora: Efficient finetuning of quantized llms](#). In *Advances in Neural Information Processing Systems (NeurIPS)*.
- George T Doran. 1981. There’s a SMART way to write management’s goals and objectives. *Management review*, 70(11):35–36.
- Carol Friedman, George Hripcsak, et al. 1999. Natural language processing and its future in medicine. *Acad Med*, 74(8):890–5.
- Isabel O. Gallegos, Ryan A. Rossi, Joe Barrow, Md Mehrab Tanjim, Sungchul Kim, Franck Dernoncourt, Tong Yu, Ruiyi Zhang, and Nesreen K. Ahmed. 2024. [Bias and fairness in large language models: A survey](#). *Computational Linguistics*, 50(3):1097–1179.
- Itika Gupta, Barbara Di Eugenio, Brian Ziebart, Aiswarya Baiju, Bing Liu, Ben Gerber, Lisa Sharp, Nadia Nabulsi, and Mary Smart. 2020. [Human-human health coaching via text messages: Corpus, annotation, and analysis](#). In *Proceedings of the 21th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 246–256, 1st virtual meeting. Association for Computational Linguistics.
- Kai He, Rui Mao, Qika Lin, Yucheng Ruan, Xiang Lan, Mengling Feng, and Erik Cambria. 2024. [A survey of large language models for healthcare: from data, technology, and applications to accountability and ethics](#). *Preprint*, arXiv:2310.05694.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [Lora: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations (ICLR)*.
- Yan Hu, Qingyu Chen, Jingcheng Du, Xueqing Peng, Vipina Kuttichi Keloth, Xu Zuo, Yujia Zhou, Zehan Li, Xiaoqian Jiang, Zhiyong Lu, et al. 2024. Improving large language models for clinical named entity recognition via prompt engineering. *Journal of the American Medical Informatics Association*, page ocad259.
- Yue Huang, Lichao Sun, Haoran Wang, Siyuan Wu, Qihui Zhang, Yuan Li, Chujie Gao, Yixin Huang, Wenhan Lyu, Yixuan Zhang, Xiner Li, Hanchi Sun, Zhengliang Liu, Yixin Liu, Yijue Wang, Zhikun Zhang, Bertie Vidgen, Bhavya Kailkhura, Caiming Xiong, Chaowei Xiao, Chunyuan Li, Eric P. Xing, Furong Huang, Hao Liu, Heng Ji, Hongyi Wang, Huan Zhang, Huaxiu Yao, Manolis Kellis, Marinka Zitnik, Meng Jiang, Mohit Bansal, James Zou, Jian Pei, Jian Liu, Jianfeng Gao, Jiawei Han, Jieyu Zhao, Jiliang Tang, Jindong Wang, Joaquin

- Vanschoren, John Mitchell, Kai Shu, Kaidi Xu, Kai-Wei Chang, Lifang He, Lifu Huang, Michael Backes, Neil Zhenqiang Gong, Philip S. Yu, Pin-Yu Chen, Quanquan Gu, Ran Xu, Rex Ying, Shuiwang Ji, Suman Jana, Tianlong Chen, Tianming Liu, Tianyi Zhou, William Yang Wang, Xiang Li, Xiangliang Zhang, Xiao Wang, Xing Xie, Xun Chen, Xuyu Wang, Yan Liu, Yanfang Ye, Yinzhi Cao, Yong Chen, and Yue Zhao. 2024. *Position: TrustLLM: Trustworthiness in large language models*. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 20166–20270. PMLR.
- Di Jin, Eileen Pan, Nassim Oufattoule, Wei-Hung Weng, Hanyi Fang, and Peter Szolovits. 2021. What disease does this patient have? a large-scale open domain question answering dataset from medical exams. *Applied Sciences*, 11(14):6421.
- Alistair Johnson, Lucas Bulgarelli, Tom Pollard, Steven Horng, Leo Anthony Celi, and Roger Mark. 2020. Mimic-iv. *PhysioNet*. Available online at: <https://physionet.org/content/mimiciv/1.0/> (accessed August 23, 2021).
- Takeshi Kojima, Shixiang (Shane) Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. In *Advances in Neural Information Processing Systems*, volume 35, pages 5547–5560.
- William Labov. 1969. A study of non-standard english.
- Haylee Lane, Mitchell Sarkies, Jennifer Martin, and Terry Haines. 2017. Equity in healthcare resource allocation decision making: a systematic review. *Social science & medicine*, 175:11–27.
- Thomas A. LaVeist. 2005. *Minority populations and health: An introduction to health disparities in the United States*, volume 4. John Wiley & Sons.
- Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, Benjamin Newman, Binhang Yuan, Bobby Yan, Ce Zhang, Christian Alexander Cosgrove, Christopher D Manning, Christopher Re, Diana Acosta-Navas, Drew Arad Hudson, Eric Zelikman, Esin Durmus, Faisal Ladhak, Frieda Rong, Hongyu Ren, Huaxiu Yao, Jue WANG, Keshav Santhanam, Laurel Orr, Lucia Zheng, Mert Yuksekgonul, Mirac Suzgun, Nathan Kim, Neel Guha, Niladri S. Chatterji, Omar Khattab, Peter Henderson, Qian Huang, Ryan Andrew Chi, Sang Michael Xie, Shibani Santurkar, Surya Ganguli, Tatsunori Hashimoto, Thomas Icard, Tianyi Zhang, Vishrav Chaudhary, William Wang, Xuechen Li, Yifan Mai, Yuhui Zhang, and Yuta Koreeda. 2023. *Holistic evaluation of language models*. *Transactions on Machine Learning Research*. Featured Certification, Expert Certification.
- Yanchen Liu, Srishti Gautam, Jiaqi Ma, and Himabindu Lakkaraju. 2023. *Investigating the fairness of large language models for predictions on tabular data*. In *Socially Responsible Language Modelling Research*.
- Masoud Monajatipoor, Jiaxin Yang, Joel Stremmel, Melika Emami, Fazlolah Mohaghegh, Mozhdeh Rouhsedaghat, and Kai-Wei Chang. 2024. *LLMs in biomedicine: A study on clinical named entity recognition*. *Preprint*, arXiv:2404.07376.
- Nambi Ndugga and Samantha Artiga. 2021. Disparities in health and health care: 5 key questions and answers. *Kaiser Family Foundation*, 11.
- OpenAI. 2023. *Gpt-4 technical report*. *Preprint*, arXiv:2303.08774.
- Joao Pereira. 1993. What does equity in health mean? *Journal of Social Policy*, 22(1):19–48.
- Jack W. Rae, Sebastian Borgeaud, Trevor Cai, Katie Millican, Jordan Hoffmann, Francis Song, et al. 2022. *Scaling language models: Methods, analysis and insights from training gopher*. *Preprint*, arXiv:2112.11446.
- Timo Schick, Sahana Udupa, and Hinrich Schütze. 2021. *Self-diagnosis and self-debiasing: A proposal for reducing corpus-based bias in nlp*. *Transactions of the Association for Computational Linguistics*, 9:1408–1424.
- Noah Shinn, Federico Cassano, Edward Berman, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. Reflexion: Language agents with verbal reinforcement learning. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Edward Shortliffe. 1976. *Computer-based medical consultations: MYCIN*. Elsevier.
- Karan Singhal, Tao Tu, Juraj Gottweis, Rory Sayres, Ellery Wulczyn, Le Hou, Kevin Clark, Stephen Pfohl, Heather Cole-Lewis, Darlene Neal, Mike Schaekermann, Amy Wang, Mohamed Amin, Sami Lachgar, Philip Mansfield, Sushant Prakash, Bradley Green, Ewa Dominowska, Blaise Aguera 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. 2023. *Towards expert-level medical question answering with large language models*. *Preprint*, arXiv:2305.09617.
- Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R. Brown, et al. 2023. *Beyond the imitation game: Quantifying and extrapolating the capabilities of language models*. *Transactions on Machine Learning Research*.
- Shubo Tian, Qiao Jin, Lana Yeganova, Po-Ting Lai, Qingqing Zhu, Xiuying Chen, Yifan Yang, Qingyu Chen, Won Kim, Donald C Comeau, Rezarta Islamaj, Audit Kapoor, Xin Gao, and Zhiyong Lu. 2023. *Opportunities and challenges for chatgpt and large*

- language models in biomedicine and health. *Briefings in Bioinformatics*, 25(1).
- Boxin Wang, Weixin Chen, Hengzhi Pei, Chulin Xie, Mintong Kang, Chenhui Zhang, Chejian Xu, Zidi Xiong, Ritik Dutta, Rylan Schaeffer, et al. 2024a. Decodingtrust: A comprehensive assessment of trustworthiness in gpt models. *Advances in Neural Information Processing Systems*, 36.
- Guanchu Wang, Junhao Ran, Ruixiang Tang, Chia-Yuan Chang, Chia-Yuan Chang, Yu-Neng Chuang, Zirui Liu, Vladimir Braverman, Zhandong Liu, and Xia Hu. 2024b. Assessing and enhancing large language models in rare disease question-answering. *Preprint*, arXiv:2408.08422.
- Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jiakai Tang, Xu Chen, Yankai Lin, Wayne Xin Zhao, Zhewei Wei, and Jirong Wen. 2024c. A survey on large language model based autonomous agents. *Frontiers of Computer Science*, 18(6).
- Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jiakai Tang, Xu Chen, Yankai Lin, et al. 2024d. A survey on large language model based autonomous agents. *Frontiers of Computer Science*, 18(6):186345.
- Hugh R Waters. 2000. Measuring equity in access to health care. *Social science & medicine*, 51(4):599–612.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837.
- Laura Weidinger, John F. J. Mellor, M. Rauh, C. Griffin, J. Uesato, Po-Sen Huang, M. Cheng, Mia Glaese, B. Balle, A. Kasirzadeh, Z. Kenton, S. Brown, W. Hawkins, T. Stepleton, C. Biles, A. Birhane, Julia Haas, Laura Rimell, Lisa Anne Hendricks, William S. Isaac, Sean Legassick, Geoffrey Irving, and Iason Gabriel. 2021. Ethical and social risks of harm from language models. *ArXiv*, abs/2112.04359.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2023. React: Synergizing reasoning and acting in language models. In *The Eleventh International Conference on Learning Representations*.
- Rich Zemel, Yu Wu, Kevin Swersky, Toni Pitassi, and Cynthia Dwork. 2013. Learning fair representations. In *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pages 325–333, Atlanta, Georgia, USA. PMLR.
- Hongjian Zhou, Fenglin Liu, Boyang Gu, Xinyu Zou, Jinfa Huang, Jing Wu, Yiru Li, Sam S. Chen, Peilin Zhou, Junling Liu, Yining Hua, Chengfeng Mao, Chenyu You, Xian Wu, Yefeng Zheng, Lei Clifton, Zheng Li, Jiebo Luo, and David A. Clifton. 2024a. A survey of large language models in medicine: Progress, application, and challenge. *Preprint*, arXiv:2311.05112.
- Yue Zhou, Barbara Di Eugenio, Brian Ziebart, Lisa Sharp, Bing Liu, and Nikolaos Agadakos. 2024b. Modeling low-resource health coaching dialogues via neuro-symbolic goal summarization and text-units-text generation. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 11498–11509, Torino, Italia. ELRA and ICCL.
- Yue Zhou, Yada Zhu, Diego Antognini, Yoon Kim, and Yang Zhang. 2024c. Paraphrase and solve: Exploring and exploiting the impact of surface form on mathematical reasoning in large language models. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 2793–2804, Mexico City, Mexico. Association for Computational Linguistics.

## A Prompt Templates

We list the prompt templates used in the paper below.

### Question Formulation Prompt

**Mortality / Readmission:** Assess the 1-year mortality risk (low or high) / 90-day readmission chance for the patient based on the given de-identified clinical note. Using medical evidence only. If uncertain, provide your best evidence-based estimate.

[Note]

**Health Coaching:** Given the following conversation between the patient and the health coach, predict the patient's outcome regarding goal achievement in the following week. Will the patient be most likely to complete the goal in the following week based on the conversation history?

[Conversation]

**MedQA:** Answer the following USMLE question with medical evidence only. No assumptions. If unsure, give your best evidence-based guess. [question] [options]

**Neighbor Scene:** The following is a transcribed conversation from an audio recording between a participant and an interviewer, who is a trained psychologist. [scene]: *In this scenario, the participant imagines they have just moved into a new neighborhood and must introduce themselves to a new neighbor.* The purpose of this exercise is to gather information about the participant's mental status through their linguistic cues. The participant belongs to one of three groups: individuals with schizophrenia, those with bipolar disorder, or healthy controls. Based on the dialogue, which group is the participant most likely to belong to?

[Conversation]

**Landlord Scene:** Same as above except for [scene]: *In this confrontational scenario, the participant imagines having a leaky pipe in their apartment that has not been fixed for a while, and they need to complain to their landlord and get it fixed.*

### Output Formatting and Other

#### Output format:

### Answer: [Task-Dependent Choices]

Then, provide your rationale.

### Rationales: ...

Each task prompt contains optional demographic information between the instruction and the data example: [Optional Demographic Information]

### Agent Prompt

To facilitate implementation, we utilized the third-party package Crewai <https://github.com/crewAIInc/crewAI> for autonomous agent prompting.

**Web Search:** *Description:* Search for the latest research, guidelines, or expert recommendations on analyzing [task] based on [example]. *Expected Output:* Provide a concise summary (within 200 words) of key points to aid in analysis. List each point with its rationale in bullet form. [Tools = Web Search]

**Analysis:** Using the provided guidelines, analyze [example] and predict [task] (varies by [Question Formulation Prompt] and [Output Formatting])

## B Additional Results

We show more fairness discrepancy across demographic subgroups in mortality prediction in Figure 3. The Demographic Parity Difference (DPD) and Equal Opportunity Difference (EOD) are calculated using a one-vs-all approach. Both GPT-4 and LLaMA-3 exhibit similar bias patterns: they are more likely to predict high mortality risks for the geriatric age group and African Americans. Additionally, these models demonstrate lower prediction performance (True Positive Rate) for these groups. These findings highlight the persistent challenges in achieving LLM fairness across different demographic subgroups in healthcare settings.

## C Sociolinguistic Consultation

We consulted with a sociolinguist regarding language model outputs that attempt to infer demo-

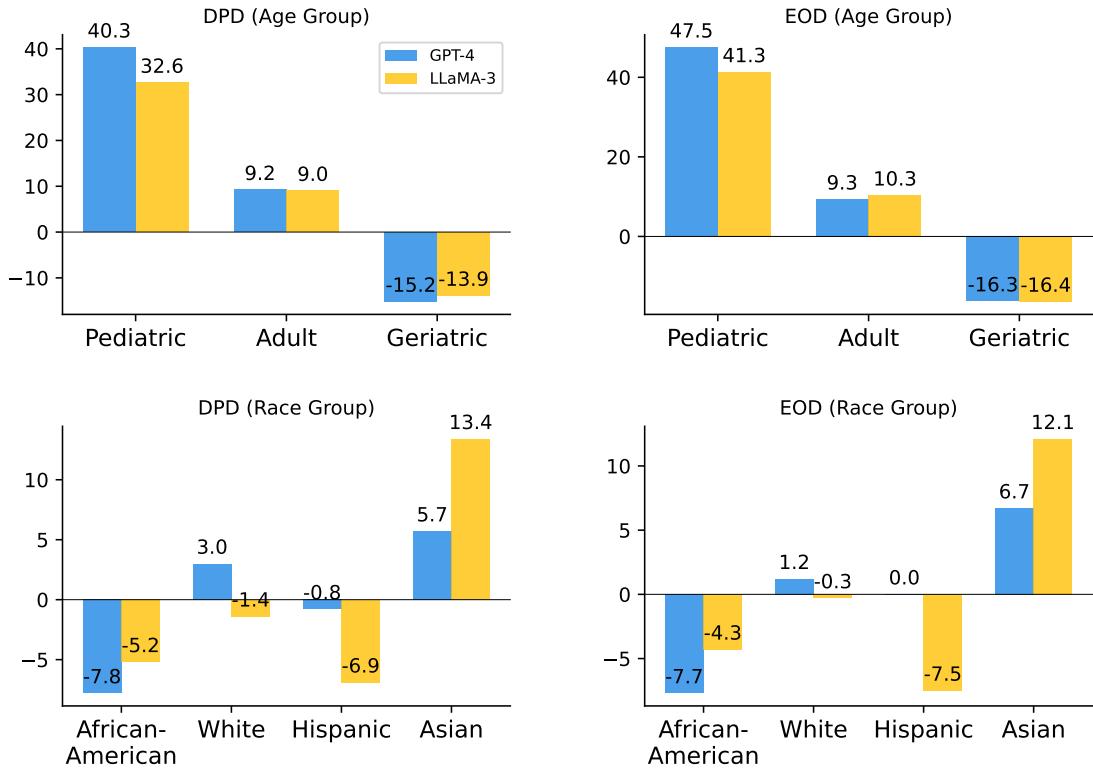


Figure 3: PDP and EOD results for more demographic subgroups (one-vs-all) for the mortality prediction task.

graphic characteristics from conversational patterns for diagnosis. The consultation revealed significant concerns about linguistic stereotyping in current LLMs.

**Misattribution of Common Linguistic Features** Features like “ain’t” and “gonna,” which LLMs often flag as African American Vernacular English (AAVE), are prevalent across multiple dialects. The expert notes that while “ain’t” showed some demographic correlation in specific contexts (e.g., Oak Park school study, Chicago area), it is not unique to AAVE. Similarly, “gonna” is a common informal contraction across all English dialects.

**Problematic Behavioral Assumptions** The models demonstrate concerning biases in attributing emotional expressions (e.g., frustration, anger) to cultural norms of specific demographic groups. The expert emphasized that such reactions are universal human responses to situations like unresolved maintenance issues, not characteristics of any particular group.

**Misinterpretation of Speech Patterns** The models incorrectly classify common features of verbal communication (e.g., short sentences, informal tone) as dialect-specific markers. However,

these are typical characteristics of spoken language across all demographics.

**Unfounded Assumptions About Language Sophistication** The models exhibit bias in equating informal language with a lack of sophistication, particularly problematic when associating this with specific demographic groups. As referenced by the expert, this misconception has been thoroughly addressed in seminal sociolinguistic works (Labov, 1969).