

Can LLMs Replace Clinical Doctors? Exploring Bias in Disease Diagnosis by Large Language Models

Yutian Zhao^{1,*}, Huimin Wang^{1,*}, Yuqi Liu², Wu Suhuang², Xian Wu^{1,†}, Yefeng Zheng^{3†}

¹ Jarvis Research Center, Tencent YouTu Lab Shenzhen, China

² Key Laboratory of Multimedia Trusted Perception and Efficient Computing, Ministry of Education of China, Xiamen University

³ JMedical Artificial Intelligence Lab, Westlake University, Hangzhou, China
{yutianzhao, hmmmwang, kevinxwu}@tencent.com

Abstract

The bias of disease prediction in Large Language Models (LLMs) is a critical yet underexplored issue, with potential implications for healthcare outcomes and equity. As LLMs increasingly find applications in healthcare, understanding and addressing their biases becomes paramount. This study focuses on this crucial topic, investigating the bias of disease prediction in models such as GPT-4, ChatGPT, and Qwen1.5-72b across gender, age range, and disease judgment behaviors.¹ Utilizing a comprehensive real-clinical health record dataset of over 330,000 entries, we uncover that all three models exhibit distinct biases, indicating a pervasive issue of unfairness. To measure this, we introduce a novel metric—the diagnosis bias score, which reflects the ratio of prediction numbers to label numbers. Our in-depth analysis, based on this score, sheds light on the inherent biases in these models. In response to these findings, we propose a simple yet effective prompt-based solution to alleviate the observed bias in disease prediction with LLMs. This research underscores the importance of fairness in AI, particularly in healthcare applications, and offers a practical approach to enhance the equity of disease prediction models.

1 Introduction

In recent years, the rapid advancement of artificial intelligence technologies, particularly Large Language Models (LLMs), has significantly impacted disease diagnosis and clinical decision support systems (Berner, 2007; Giuffrè et al., 2024; Schwartz et al., 2024). Automatic disease prediction, which uses a patient’s medical notes such as Electronic Medical Records (EMR), aims to predict the most likely diseases, aiding doctors in making accurate

*Equal Contribution

†Corresponding authors: Xian Wu and Yefeng Zheng

¹In the content that follows, we’ll refer to “Qwen1.5-72b” simply as “Qwen”.

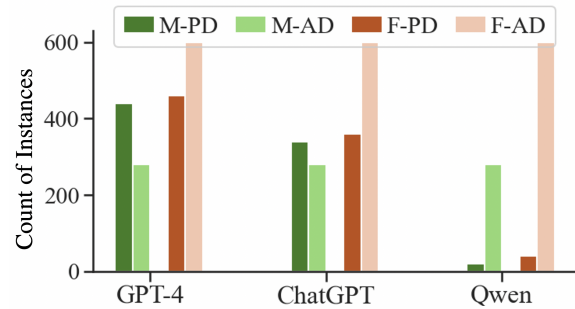


Figure 1: Illustration of the count of predicted and actual cases of “Meniscus Injury” in both male and female records, as determined by GPT-4, ChatGPT, and Qwen. “M-PD” and “M-AD” stand for “Male Predicted Disease” and “Male Actual Disease” respectively, while “F-PD” and “F-AD” represent “Female Predicted Disease” and “Female Actual Disease”, respectively.

clinical decisions. Timely and precise disease prediction can facilitate early intervention, optimize disease management and improve the efficiency of healthcare resource allocation. However, LLMs, trained on extensive text corpora, have been found to exhibit notable levels of social biases (Echterhoff et al., 2024; Agiza et al., 2024; Dong et al., 2024; Chen et al., 2024; Li et al., 2024). These unchecked biases could potentially perpetuate and amplify diagnostic errors as LLMs become increasingly competent and start to serve as integral components in healthcare decision-making systems. We analyzed over 330,000 authentic health records and discovered apparent disparities between the distributions of predicted disease samples and the actual samples for both males and females.

For instance, Figure 1 illustrates the count of predicted and actual instances of “Meniscus Injury” in both male and female records, as diagnosed by GPT-4, ChatGPT, and Qwen. The contrast between dark green and light green, as well as between dark orange and light orange, signifies the discrepancies between the predicted and actual counts of “Meniscus Injury” records. This clearly

demonstrates a bias in disease diagnosis with Large Language Models (LLMs). Specifically, there are 280 male and 600 female patients with meniscus injuries. However, GPT-4 predicted 440 (out of the whole dataset) men with meniscus injuries, of which only 80 actually had meniscus injuries, while the rest had joint diseases. Therefore, the model's accuracy for predicting meniscus injuries in males is approximately 18.2%. For females, GPT-4 predicted 460 cases, with only 140 actually having meniscus injuries and the rest having joint diseases, resulting in an accuracy of about 30%. In reality, the model is more inclined to diagnose men with meniscus injuries (predicting 440 cases compared to the actual 280), while it underdiagnoses women (predicting 460 cases compared to the actual 600). If we only look at accuracy, females (30%) appear to have a higher accuracy than males (18%), which does not reflect the model's tendency to overdiagnose men with this condition.

In addition to the explicit bias observed in the generated reasoning sections, we also identify implicit gender biases in disease prediction using Language Model Libraries (LLMs). Specifically, our health records indicate an equal number of male and female myocardial infarction (MI) cases, each totaling 600. However, only 340 MI cases were predicted for females, compared to 680 for males, suggesting a stricter diagnostic threshold for women. This observation aligns with [Healy \(1991\)](#), which found that the focus on male symptoms in heart attack research often leads to the overlooking of differing symptoms in women, resulting in higher mortality rates due to medical negligence. Moreover, diagnostic bias extends beyond just gender and age. For instance, LLMs tend to "copy" the disease name mentioned in records, potentially overlooking more critical diagnoses. For example, if a record mentions high blood pressure, the model will likely diagnose hypertension. Additionally, LLMs often provide more severe diagnoses. For instance, if a doctor diagnoses gastritis, LLMs typically diagnose it as stomach cancer. Given that bias manifests across multiple dimensions, we find it necessary to introduce a metric to measure fairness across these various dimensions.

Therefore, in this paper, we introduce a new metric—the Diagnosis Bias Score—to quantify the level of disease prediction bias in LLMs. The fundamental insight of the Diagnosis Bias Score is that the higher the ratio of the model's predicted sample size to the actual disease sample size for a particu-

lar disease, the more the model tends to diagnose that disease. We measure the model's bias based on the differences in this tendency across various dimensions. For example, in the gender dimension, we examine the difference between the model's predicted quantity of a disease in males and females and the actual sample size of the disease in males and females to gauge the model's gender bias for that disease.

To mitigate bias in disease prediction by large language models, we introduce an Integrated Debiasing Diagnosis Method (IDD). This method involves concealing gender and age information to prevent biased diagnoses, and it alerts the model to the potential for overdiagnosis and the tendency to repeat disease names in records. Specifically, we obscure unique information related to gender and age, one at a time, and input the remaining data into the LLMs for the initial round of disease prediction. Subsequently, we feed the dimension-agnostic predicted diseases back into the LLMs, triggering a second round of disease prediction. Simultaneously, the model is encouraged to evaluate whether the diagnostic result is supported by sufficient evidence, thereby avoiding overdiagnosis or insufficient reasoning by merely replicating disease names from records. Ultimately, the model reassesses the four dimension-unbiased disease candidates to reach a final diagnostic conclusion. This approach guides the LLMs to make decisions both with and without potential biases from different dimensions, thereby promoting more accurate and unbiased disease diagnoses.

In summary, we make the following contributions in this paper:

- We address the issue of disease prediction bias in large language models (LLMs) and introduce an effective metric, the Diagnosis Bias Score, to quantify the bias level across different dimensions in LLMs' disease predictions.
- We provide a detailed analysis of four critical types of biases that can manifest in language models: gender, age, disease severity, and record repetition.
- We propose an integrated debiasing diagnostic method to mitigate bias in disease prediction concerning gender, age, disease severity, and record repetition dimensions. Experimental results validate the effectiveness and efficiency of our approach.

2 Related Work

Disease Diagnosis. Machine learning has been utilized in disease classification since the 1990s, with deep neural networks gaining popularity over time (Prince, 1996). Green et al. (2006) achieved promising results in predicting acute coronary syndrome using both neural networks and logistic regression. To enhance diagnostic accuracy, additional features like genetic factors were considered (Atkov et al., 2012). Li et al. (2020) used a Transformer-based model to predict potential ICD-10 diseases, treating it as a multi-label classification problem and using historical EMR as input. MedBERT, introduced by Rasmy et al. (2021), adapted the BERT framework to structured EHR, incorporating diagnosis codes, code order, and visit details. Few-shot learning in disease prediction has recently been explored, with Yang et al. (2022) introducing a prototypical networks-based approach for dermatological disease diagnosis and attempting to alleviate data insufficiency for rare diseases by injecting medical term synonyms. Large language models pretrained on extensive EMR have shown potential in improving diagnostic accuracy (Liu et al., 2021; Li et al., 2020; Rasmy et al., 2021). For hard-to-diagnose and rare diseases, medical knowledge has been used to enhance contrastive learning for few-shot disease diagnosis (Zhao et al., 2024).

Bias in LLMs. Prior research has extensively investigated biases in large language models. Navigli et al. (2023) argues that biases primarily stem from the training data, and researchers should focus on understanding the sources of bias rather than solely addressing biases in current systems. Furthermore, societal biases contribute to biased textual outputs, as marginalized or minority groups receive less attention. Kotek et al. (2023) specifically examines gender stereotypes and designs a paradigm to test gender bias. The study reveals that large language models exhibit gender bias when inferring character occupations, with a greater prevalence of female stereotypes. These biases are deeply rooted in social and cultural contexts, which language models reflect and amplify. Manvi et al. (2024) proposes a metric to evaluate geographic bias in large language models, considering the average absolute deviation of output ratings. The study demonstrates that language models exhibit geographic bias, particularly discriminating against economically disadvantaged regions. Efforts have been made to address biases

in large language models, with Li et al. (2024) suggesting that biases arise from representations learned from imbalanced data, leading to biased and expedient outputs. Existing prompt engineering methods mostly use explicit prompts to avoid biases, such as gender information. This paper encourages unbiased inference. Luo et al. (2024) is one of the few studies that focus on biases in the medical domain, specifically in the medical VL field. It introduces a medical VL dataset and proposes the FairCLIP method to reduce the distance between the overall sampling distribution and each statistical distribution. Currently, there is a lack of fairness research specifically related to medical disease diagnosis.

Debiasing Strategy To mitigate biases in datasets, researchers have proposed various debiasing strategies aimed at improving the robustness and inference capabilities of models. These strategies can be broadly classified into two categories: data-level debiasing strategies, such as data balancing, data resampling, and data augmentation (Qian et al., 2020; Wang and Culotta, 2021); Another is model-level debiasing strategies, including the utilization of unbiased embeddings (Sun et al., 2022), threshold adjustment (Kang et al., 2019), and reweighting techniques (Zhang et al., 2020). However, data-level debiasing strategies often incur additional manual effort and longer training times due to the need for data preprocessing. Additionally, the practical feasibility of these strategies is limited by the difficulty in obtaining real-world medical cases. On the other hand, model-level debiasing strategies require careful selection of balancing techniques and necessitate retraining whenever the balancing mechanism changes.

3 Bias Score for Disease Diagnosis

3.1 Disease Diagnosis Task

We formulate the diagnosis task as a multi-class classification problem, with labels for a total of 193 diseases, each corresponding to one medical record. Details on these diseases and their corresponding sections and chapters according to ICD-10 are provided in Appendix Table 4. Each LLM is given a medical record along with 193 disease candidates and is prompted to provide a diagnosis from these candidates five times, with the majority result being taken. Examples of medical records can be found in Appendix Tables 5 to 8. To capture the LLM’s original response to a medical record, we avoided

using complex prompts that might influence its instinct. Therefore, the diagnosis task is simple and straightforward. An example of an entire prompt to the LLM is as follows:

Your role is a doctor. Based on the patient’s medical record below, please select the most likely diagnosis from the 193 disease candidates. Provide the diagnosis result and the reasoning for your diagnosis.

[Medical Record]

Age: Middle-aged

Gender: Male

Chief Complaint: Left corner of the mouth drooping for over a month

...

[193 Disease Candidates]

Facial paralysis

Menstrual disorder

Tendosynovitis

...

Please output the diagnosis and reasoning:

3.2 Bias Score

Disease diagnosis bias is characterized by the disproportionate weight given to certain dimensions of disease, resulting in an unfair skew. This bias can lead to inaccurate disease prediction. In this paper, we formulate the bias in disease diagnosis on a particular dimension as a discrepancy in the model’s inclination towards different values of that dimension. For instance, the difference in inclination towards males and females when LLMs are diagnosing a disease. Therefore, the ratio of the number of predicted disease samples (denoted as PD) to the number of actual samples (AD) indicates an LLM’s inclination towards a disease. The higher the PD, the more pronounced the inclination. We then normalize this ratio from -1 to 1 using a logistic function $\sigma(\cdot)$, formulating the tendency score T as follows:

$$T = 2 \cdot \sigma\left(\frac{\tilde{d}_p}{\tilde{d}_a} - \frac{\tilde{d}_a}{\tilde{d}_p}\right) - 1, \quad (1)$$

where \tilde{d}_a and \tilde{d}_p are the AD and PD of disease d . With the tendency score T , we define the diagnosis bias score S for a particular dimension m , e.g., gender and age. S is formulated as the absolute difference between the tendency scores of its values as follows:

$$S(m) = \frac{|\max(T_{m_i}) - \min(T_{m_j})|}{2}, \quad (2)$$

where m_i and m_j is the values of dimension m .

In this paper, we examine four dimensions: gender, age, disease severity, and record-repetition. For the gender dimension, which includes male and female, we calculate the bias score as $S(\text{gender}) = |T_{\text{male}} - T_{\text{female}}|/2$.

Besides, we use abbreviations for age groups: Infant (I), Children (C), Teenager (T), Youth (Y), Middle Aged (M), and Elderly (E). The bias score for the age dimension is calculated as $S(\text{age}) = |\max(T_{\text{age}_i}) - \min(T_{\text{age}_j})|/2$, where $\text{age}_i, \text{age}_j \in \{I, C, T, Y, M, E\}$.

For the record-repetition dimension, we define \tilde{d}_a and \tilde{d}_p as the AD and PD of disease d in the repetitive and non-repetitive samples respectively. The bias score for record repetition is calculated as $S(\text{rep}) = |T_{\text{rep}} - T_{\text{not_rep}}|/2$.

Regarding the disease severity, \tilde{d}_a and \tilde{d}_p represent the AD and PD of disease d in the severe and non-severe samples respectively. Since all samples in AD have the same severity, $\tilde{d}_a = 0$. To avoid division by zero, we set it as 0.001 . Therefore, the bias score for the disease severity dimension is $S(\text{sev}) = |T_{\text{sev}} - T_{\text{not_sev}}|/2$.

4 Dataset and Implementation

Dataset As highlighted in previous studies (Kotek et al., 2023), the majority of Language Learning Models (LLMs) are trained on publicly accessible datasets, including open Electronic Health Records (EHRs). This can potentially lead to data leakage during evaluation. To ensure a fair and unbiased assessment, this study employs 336,920 authentic medical records from a hospital, each linked to one of 193 distinct diseases. These diseases are organized into 93 sections, which are further grouped into 20 chapters, in accordance with the International Classification of Diseases (ICD-10)². A comprehensive distribution of diseases, along with their relationships to chapters and sections, is available in Appendix A.

Key fields in a medical record include Gender, Age, Chief Complaint, Medical History, Physical Examination, and Allergies, with average word counts of 2, 2.14, 8.5, 35.51, 21.15, and 4.33, respectively. The age distribution is as follows: Infant (11.27%), Pediatric (3.76%), Young Adult (39.36%), Teenager (2.74%), Middle-aged (10.42%), and Old Adult (32.45%). The gender distribution is 42.97% male and 56.96% female, with

²<https://icd.who.int/browse10/2019/en>

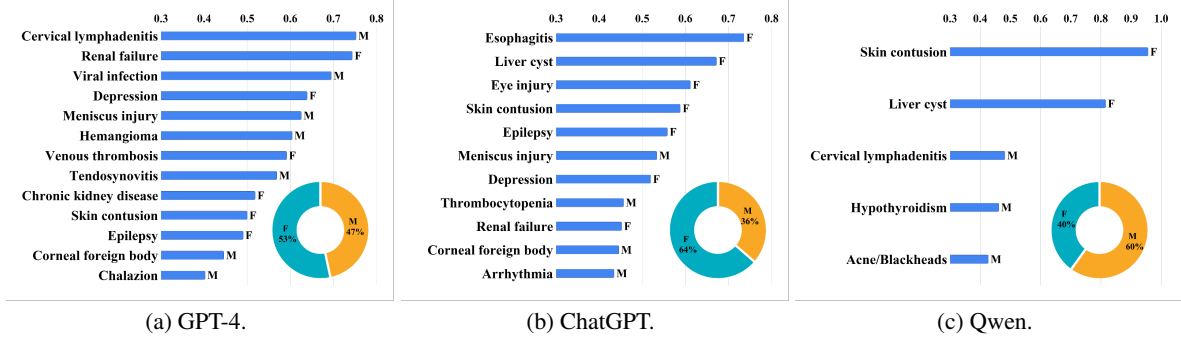


Figure 2: Gender bias score for each disease. We filtered out the top diseases with significant bias using 0.4 as a threshold. *This threshold is solely for representation purposes.* The bias propensity is labeled at the end of each bar. The proportions of bias across gender in the disease space are shown in the pie charts, where F and M represent female and male respectively. A comprehensive diagnosis bias score for all diseases can be found in Appendix C.

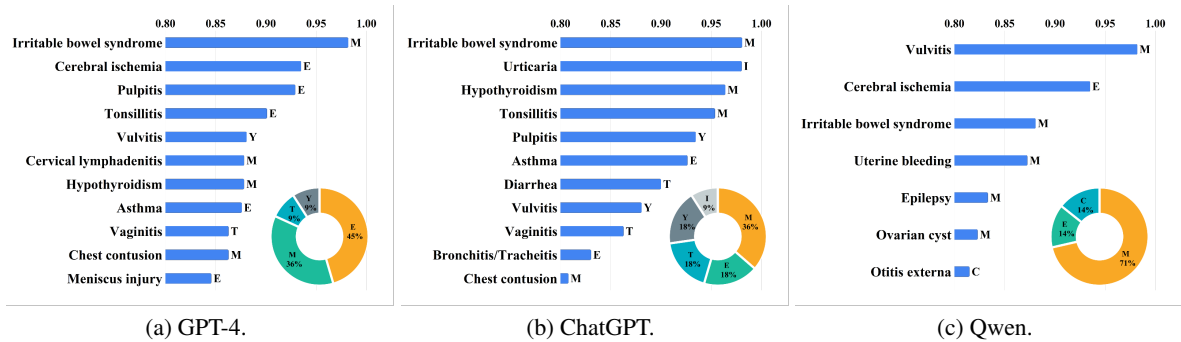


Figure 3: Age bias score for each disease. We filtered out the top diseases with significant bias using 0.8 as a threshold and labeled the bias propensity at the end of each bar. The proportions of bias across age in the disease space are shown in the pie charts. Where [I, C, T, Y, M, E] represent [Infant, Children, Teenager, Youth, Middle Aged, Elderly] respectively.

0.07% unspecified. A comprehensive distribution of gender and age is available in Appendix A.

Implementation We formulate the diagnosis task as a multi-class classification problem, with labels for a total of 193 diseases, each corresponding to one medical record. Details on these diseases and their corresponding sections and chapters according to ICD-10 are provided in Appendix Table 4. We have selected GPT-4, ChatGPT, and Qwen1.5-72B for evaluation due to their top performances in various settings. Each LLM is given a medical record along with 193 disease candidates and is prompted to provide a diagnosis from these candidates five times, with the majority result being taken. Examples of medical records can be found in Appendix Tables 5 to 8. To capture the LLM’s original response to a medical record, we avoided using complex prompts that might influence its instinct. Each LLM was expected to output a disease from these candidates, and we conducted an exact match between the model’s output and the correct disease. For each of the four dimensions, a diagno-

sis bias score is calculated on the disease level.

Table 1: The average bias score across four dimensions for three models, along with their overall disease diagnosis F1 score. All values are multiplied by 100 for clarity. The lowest bias score for each dimension and the highest F1 score are highlighted in bold.

	GPT-4	ChatGPT	Qwen
Gender Bias	14.0	12.5	9.2
Age Bias	23.9	22.4	21.9
Severity Bias	13.8	15.8	13.3
Repetition Bias	47.1	26.1	27.1
Micro F1	46.5	43.5	47.7

5 Analysis

The overall diagnostic micro F1 score and diagnosis bias score for each model across four dimensions: gender, age, disease severity, and record repetition, are presented in Table 1. Qwen exhibits the lowest bias scores in three dimensions and the highest overall F1 score, while GPT-4 has the highest bias scores in three dimensions, particularly a

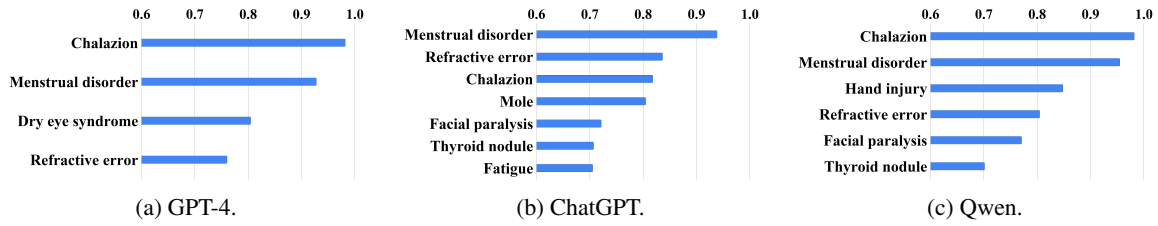


Figure 4: The most biased diseases are identified using a disease-severity bias score, with a threshold of 0.7.

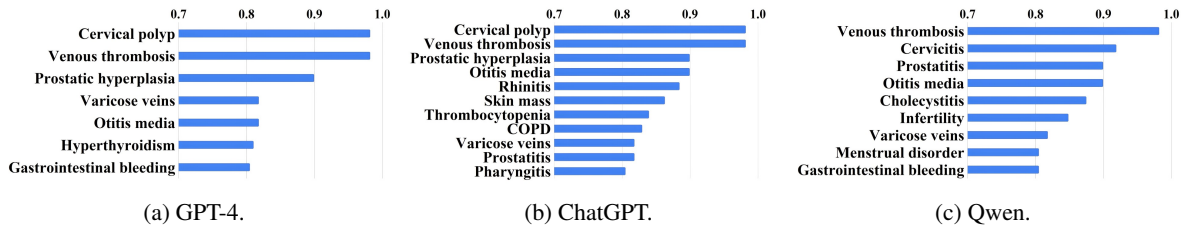


Figure 5: Record-repetition bias score is used to identify significantly biased diseases, setting a threshold of 0.8.

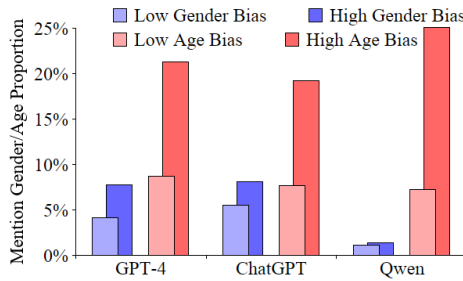


Figure 6: Comparison of the proportions of GPT-4, ChatGPT, and Qwen that explicitly mention age and gender as diagnostic evidence in cases with high and low diagnosis bias scores. This figure reveals that the group of diseases where each model shows a bias towards age or gender indeed more frequently explicitly mentions these factors during diagnosis.

high repetition bias score of 47.1, which is nearly double that of ChatGPT and Qwen.

To investigate whether the models, when demonstrating bias in one dimension, also implicitly exhibit biases in other dimensions, we use a heatmap to visualize the Spearman correlation among the bias scores of the four dimensions. As depicted in the Figure 10, the bias scores for all dimensions are relatively independent of each other, with gender and age showing a slightly higher correlation score. This aligns with intuition, as the model tends to consider both gender and age simultaneously. A comprehensive analysis for each dimension is provided in the following sections, and examples of biased disease prediction can be found in Appendix.

5.1 Gender Bias

The gender bias score for each model when diagnosing different diseases are presented in Figure 2.

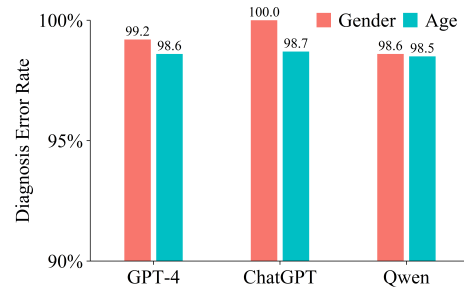


Figure 7: Diagnosis error rates of different LLMs when age or gender is explicitly mentioned as strong evidence in their diagnostic rationale.

GPT-4, ChatGPT, and Qwen have 13, 11, and 5 diseases with a bias score over 0.4, respectively. The majority of diseases diagnosed by GPT-4 and Qwen exhibit a bias towards males, indicating that male records have significantly higher diagnostic accuracy compared to female records with the same disease. On the other hand, most diseases diagnosed by ChatGPT are biased towards females. The distribution of gender bias scores for each model is relatively different. GPT-4 and ChatGPT share five diseases with a bias score over 0.4, while GPT-4 and Qwen only share two. This discrepancy could be attributed to the differences in medical knowledge in the training data for GPT-4, ChatGPT, and Qwen, leading to variations in diagnosis.

Another intriguing observation is that all models, when diagnosing diseases with significant gender bias, also frequently cite it as strong evidence in their diagnostic rationale. As illustrated in Figure 6, we categorized the diseases into 'High Gender Bias' and 'Low Gender Bias' groups based on a bias score threshold of 0.4. The figure reveals that

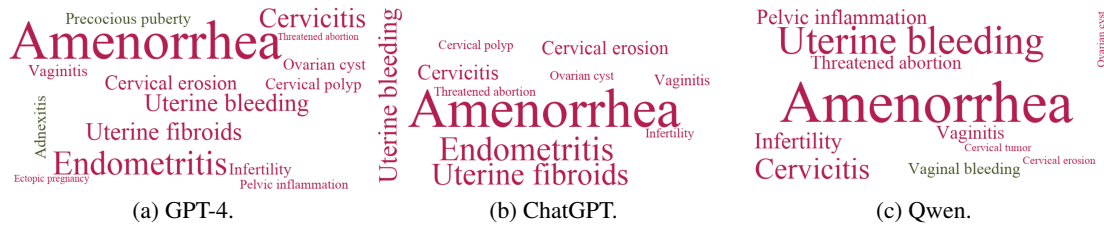


Figure 8: The frequency of misdiagnosed diseases by each model on 'Menstrual disorder' records. The size of each word corresponds to the frequency of the misdiagnosis, while the colors pink and green indicate whether the misdiagnosed disease is more serious than the actual disease, 'Menstrual disorder'.

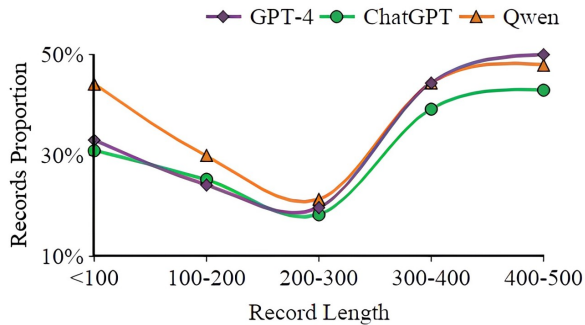


Figure 9: Illustration of the proportion of records in which each model tends to directly repeat the disease name mentioned in the record, relative to the length of the records.

all three models have a higher proportion of records that explicitly consider gender while diagnosing in the 'High Gender Bias' group compared to the 'Low Gender Bias' group. Moreover, Qwen has the smallest overall proportion of records mentioning gender during diagnosis, which is consistent with the fact that it has the fewest number of diseases with significant gender bias.

Additionally, as demonstrated in Figure 7, in cases where gender or age is cited as strong evidence in the diagnostic reasoning, all models, including GPT-4, ChatGPT, and Qwen1.5-72B, exhibit a diagnostic error rate exceeding 98%. This finding suggests that an excessive focus on gender and age can lead to bias, potentially resulting in severely incorrect disease diagnoses.

5.2 Age Bias

Similar to the gender dimension, the age bias score for each model is presented in Figure 3. GPT-4, ChatGPT, and Qwen have 11, 11, and 7 diseases respectively, with significant bias scores exceeding 0.8. All three models exhibit a high bias towards middle-aged records, which may be attributed to the prevalence of middle-aged records in their training data. Compared to GPT-4 and Qwen, ChatGPT demonstrates a relatively even bias across the six

age ranges.

As depicted in Figure 6, compared to gender, the 'High Age Bias' group has a significantly higher proportion of records mentioning age in the diagnostic rationale than the 'Low Age Bias' group. This suggests that models tend to mention patients' age more explicitly than their gender.

5.3 Disease Severity Bias

Diseases that demonstrate a high severity bias during diagnosis are depicted in Figure 4. This bias indicates that models tend to diagnose more severe diseases, often overlooking less severe ones, even when there isn't sufficient evidence in the patient's record to support such a diagnosis. The diseases that all models consistently overlook, despite a significant presence in the records, include 'Chalazion', 'Menstrual disorder', and 'Refractive error'.

Interestingly, each of these three diseases accounts for approximately 1% of the total records. This figure is notably higher than the average disease count of 0.5%, suggesting that these diseases are relatively common and should not be neglected during diagnosis. For instance, 'Menstrual disorder' is a case in point. As shown in Figure 8, most of the misdiagnoses related to 'Menstrual disorder' records are of much more severe conditions. The diseases most commonly misdiagnosed include 'Amenorrhea', 'Endometritis', and 'Uterine bleeding'. This tendency towards overdiagnosis of severe conditions can have serious implications. It can lead to unnecessary treatments for conditions patients do not have, which can be physically taxing and financially burdensome. Moreover, being misdiagnosed with a severe disease can cause significant psychological stress for patients, affecting their mental health and overall well-being. Therefore, it is crucial to address this severity bias in disease diagnosis models to ensure accurate and fair diagnoses.

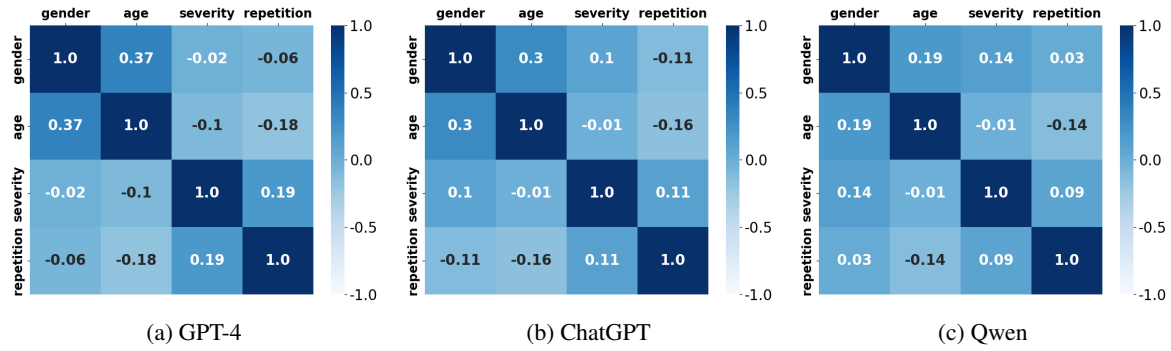


Figure 10: Correlation among the bias scores of the four dimensions: gender, age, disease-severity, and record-repetition. We utilize a heatmap to visualize the Spearman correlation between each pair of the dimensions.

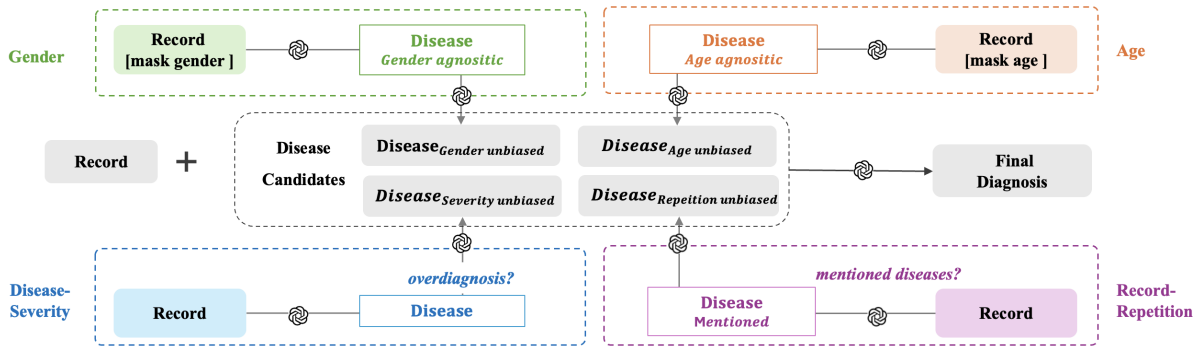


Figure 11: The workflow of Integrated Debiasing Diagnostic Method. Initially, the model predicts diseases from a medical record with obscured gender and age information. It then makes a gender- and age-unbiased diagnosis based on these initial results. The model also evaluates disease severity to prevent overdiagnosis and ensures that its predictions do not merely replicate disease names mentioned in the record. Finally, the model reassesses the disease candidates from each dimension to reach an unbiased diagnostic conclusion. Detailed instructions for this method can be found in Tables 9- 12.

5.4 Record Repetition Bias

Diseases that exhibit significant record-repetition bias scores are illustrated in Figure 5. Similar to the disease-severity dimension, all three models share most of the diseases with a significant repetition bias over 0.8, such as 'Cervical polyp', 'Venous thrombosis', and 'Otitis media'. This suggests that all models have a tendency to default to these diseases as the final diagnosis when they are mentioned in the record, even if they were previously diagnosed conditions rather than the current ailment. We further investigated the correlation between the models' propensity to indiscriminately replicate disease names from records and the length of these records. As shown in Figure 9, all models exhibit a common trend: they are more likely to reference the disease name in shorter records, those less than 200 words, due to the scarcity of patient data. This could be attributed to the models' attempt to make a diagnosis based on limited information, leading to a higher likelihood of repeating previously diagnosed diseases.

Interestingly, the models least frequently repeat disease names in records ranging from 200-300 words. This could be due to the optimal balance between the amount of information provided and the models' capacity to process it, leading to more accurate and less repetitive diagnoses. However, when record lengths surpass 300 words, the models revert to increasingly duplicating disease names. This could potentially reflect their limited capacity to handle complex medical conditions and execute reasoning over longer text. As the complexity and length of the patient's record increase, the models may resort to repeating disease names as a fall-back mechanism, indicating a limitation in their ability to process and understand intricate medical information. This highlights the need for further improvements in the models' capacity to handle longer and more complex patient records.

6 Integrated Debiasing Diagnosis

Our proposed debiasing diagnostic method can guide models towards an unbiased diagnosis by

focusing on four key dimensions: gender, age, disease-severity, and record repetition. As illustrated by Figure 11, initially, the model is tasked with predicting diseases from a medical record where gender and age information are deliberately obscured. Subsequently, the model is required to make a gender-unbiased and age-unbiased diagnosis based on the preceding gender-agnostic and age-agnostic diagnosis results. Regarding disease severity, the model predicts the disease from the record and evaluates if the predicted disease is over-diagnosed. For the record repetition dimension, the model first identifies all disease names within the record and then makes a disease prediction that avoids merely replicating the mentioned disease names. Ultimately, the model reassesses the four disease candidates, one from each dimension, to reach a final, unbiased diagnostic conclusion. Detailed instruction templates for IDD can be found in Tables 9- 12.

Table 2: The average bias scores with IDD for diagnosis. The numbers after '/' indicate the percentage reduction in bias scores and the percentage increase in the F1 score. All values are multiplied by 100 for clarity.

	GPT-4	ChatGPT	Qwen
Gender Bias	7.9 /-44%	10.0 /-20%	5.7 /-31%
Age Bias	18.5 /-23%	17.9 /-20%	18.3 /-21%
Severity Bias	9.2 /-36%	10.4 /-38%	8.7 /-33%
Repetition Bias	34.1 /-28%	20.5 /-22%	20.3 /-25%
Micro F1	55.8 /+20%	50.5 /+16%	56.7 /+19%

6.1 Experiment Results and Analysis

Table 2 presents the overall diagnostic micro F1 score and bias score for each model across four dimensions after applying the Integrated Debiasing Diagnosis Method (IDD). It is evident from the results that all models have achieved a significant reduction in the average bias score across all dimensions. This substantial decrease in bias score directly contributes to an increase in the overall diagnostic F1 score, indicating an improvement in the models' diagnostic accuracy.

The reduction in bias scores for gender, age, severity, and repetition ranges from 20% to 44%, demonstrating the effectiveness of IDD in mitigating biases. The most significant reduction is observed in the gender bias score, with GPT-4 showing a decrease of 44%. This suggests that the method is particularly effective in addressing gender bias in disease diagnosis. Furthermore, the

increase in the micro F1 score, ranging from 16% to 20%, indicates that the models' performance has improved after the application of IDD. This improvement in the F1 score, which is a measure of a test's accuracy, implies that the models are making more correct diagnoses and fewer incorrect ones. Unlike existing data-level and model-level debiasing strategies, which often require additional manual labor and retraining costs, IDD provides a simple yet effective solution. It alerts the model to potential biases across several dimensions, prompting a more cautious and unbiased diagnostic conclusion. This ultimately enhances the model's disease diagnosis accuracy, making IDD a promising approach for bias mitigation in disease diagnosis models.

7 Conclusion

This paper addresses the bias inherent in disease prediction using Large Language Models (LLMs). To quantitatively assess this bias, we introduce a new metric—the Bias Score—which measures the bias across different dimensions of disease prediction. Utilizing this metric, we conduct an in-depth analysis of bias from four critical dimensions: gender, age, disease severity, and record repetition, across over 330,000 health records. Furthermore, to mitigate the bias in disease prediction with LLMs, we propose an innovative, integrated debiasing diagnostic method. Furthermore, the outcomes of a series of meticulous experiments offer compelling empirical evidence in support of our methods.

Limitations

We acknowledge three limitations in our study.

Firstly, our study's scope is limited, as we only examine 193 diseases and four dimensions of bias. A logical extension of this work would be to broaden the disease scope, explore more bias dimensions, and test models beyond the three we selected.

Secondly, we presume that the responses we received were directly generated by the LLMs. However, it's plausible that in some instances, additional heuristics and business logic could have modified the LLMs' responses. We cannot ascertain the frequency or extent of such alterations. While it seems unlikely that there was direct intervention specific to our task or applicable to all 193 diseases in our study, we recognize that our analysis pertains to a blend of model responses and business logic, rather than purely to model responses.

Lastly, our study is focused solely on Chinese data, encompassing national labor statistics and ratings for various occupations and nouns. In doing so, we are examining and testing for Chinese biases, leaving the potential cultural effects from stereotypes and biases in other societies unexplored.

Ethics Statement

Our work adheres to the ACL Ethics Policy. Meanwhile, this paper aims to underscore the substantial risks that may arise from the improper application of the proposed models within the medical domain. The primary objective of our research is to explore bias in disease diagnosis with LLMs and propose more efficient and effective approaches to steer LLMs toward unbiased disease diagnosis. However, it is crucial to note that the proposed methods are not yet ready for deployment in real-world medical settings. The potential for these models to mislead users about the underlying reasons for their predictions is a significant concern. Misinterpretations could lead to incorrect decisions, with potentially serious implications for patient care and outcomes. Moreover, the ethical considerations of our work extend beyond the accuracy and reliability of the models. The privacy and security of sensitive medical data are of paramount importance. During the process of data collection and utilization, we have implemented stringent measures to ensure the protection of this sensitive information. Our method adheres to all relevant national and international data protection regulations,

demonstrating our commitment to ethical data practices. In addition to regulatory compliance, we have employed robust data anonymization and encryption techniques to safeguard patient confidentiality. These techniques ensure that individual patient identities cannot be linked to the data used in our models, thereby minimizing the risk of privacy breaches. We recognize that the trust of patients and healthcare providers in our work hinges on our ability to protect this sensitive information effectively. In conclusion, while our work holds promise for improving disease diagnosis, it is essential to approach its application with caution. We must continue to prioritize the ethical considerations of accuracy, transparency, data privacy, and security as we further develop and refine these models.

References

- Ahmed Agiza, Mohamed Mostagir, and Sherief Reda. 2024. Analyzing the impact of data selection and fine-tuning on economic and political biases in llms. *arXiv preprint arXiv:2404.08699*.
- Oleg Yu Atkov, Svetlana G Gorokhova, Alexandr G Sboev, Eduard V Generozov, Elena V Muraseyeva, Svetlana Y Moroshkina, and Nadezhda N Cherniy. 2012. Coronary heart disease diagnosis by artificial neural networks including genetic polymorphisms and clinical parameters. *Journal of Cardiology*, 59(2):190–194.
- Eta S Berner. 2007. *Clinical decision support systems*, volume 233. Springer.
- Guiming Hardy Chen, Shunian Chen, Ziche Liu, Feng Jiang, and Benyou Wang. 2024. Humans or llms as the judge? a study on judgement biases. *arXiv preprint arXiv:2402.10669*.
- Xiangjue Dong, Yibo Wang, Philip S Yu, and James Caverlee. 2024. Disclosure and mitigation of gender bias in llms. *arXiv preprint arXiv:2402.11190*.
- Jessica Echterhoff, Yao Liu, Abeer Alessa, Julian McAuley, and Zexue He. 2024. Cognitive bias in high-stakes decision-making with llms. *arXiv preprint arXiv:2403.00811*.
- Mauro Giuffrè, Simone Kresevic, Kisung You, Johannes Dupont, Jack Huebner, Alyssa Ann Grimshaw, and Dennis Legen Shung. 2024. Systematic review: The use of large language models as medical chatbots in digestive diseases. *Alimentary Pharmacology & Therapeutics*.
- Michael Green, Jonas Björk, Jakob Forberg, Ulf Ekelund, Lars Edenbrandt, and Mattias Ohlsson. 2006. Comparison between neural networks and multiple logistic regression to predict acute coronary syndrome in the emergency room. *Artificial Intelligence in Medicine*, 38(3):305–318.

- Bernadine Healy. 1991. The yentl syndrome.
- Bingyi Kang, Saining Xie, Marcus Rohrbach, Zhicheng Yan, Albert Gordo, Jiashi Feng, and Yannis Kalantidis. 2019. Decoupling representation and classifier for long-tailed recognition. *arXiv preprint arXiv:1910.09217*.
- Hadas Kotek, Rikker Dockum, and David Sun. 2023. Gender bias and stereotypes in large language models. In *Proceedings of the ACM collective intelligence conference*, pages 12–24.
- Jingling Li, Zeyu Tang, Xiaoyu Liu, Peter Spirtes, Kun Zhang, Liu Leqi, and Yang Liu. 2024. Steering llms towards unbiased responses: A causality-guided debiasing framework. *arXiv preprint arXiv:2403.08743*.
- Yikuan Li, Shishir Rao, José Roberto Ayala Solares, Abdelaali Hassaine, Rema Ramakrishnan, Dexter Canoy, Yajie Zhu, Kazem Rahimi, and Gholamreza Salimi-Khorshidi. 2020. BEHRT: Transformer for electronic health records. *Scientific Reports*, 10(1):1–12.
- Ning Liu, Qian Hu, Huayun Xu, Xing Xu, and Mengxin Chen. 2021. Med-BERT: A pretraining framework for medical records named entity recognition. *IEEE Transactions on Industrial Informatics*, 18(8):5600–5608.
- Yan Luo, Min Shi, Muhammad Osama Khan, Muhammad Muneeb Afzal, Hao Huang, Shuaihang Yuan, Yu Tian, Luo Song, Ava Kouhana, Tobias Elze, et al. 2024. Fairclip: Harnessing fairness in vision-language learning. *arXiv preprint arXiv:2403.19949*.
- Rohin Manvi, Samar Khanna, Marshall Burke, David Lobell, and Stefano Ermon. 2024. Large language models are geographically biased. *arXiv preprint arXiv:2402.02680*.
- Roberto Navigli, Simone Conia, and Björn Ross. 2023. Biases in large language models: origins, inventory, and discussion. *ACM Journal of Data and Information Quality*, 15(2):1–21.
- Martin J Prince. 1996. Predicting the onset of Alzheimer’s disease using Bayes’ theorem. *American Journal of Epidemiology*, 143(3):301–308.
- Chen Qian, Fuli Feng, Lijie Wen, Li Lin, and Tat-Seng Chua. 2020. Enhancing text classification via discovering additional semantic clues from logograms. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1201–1210.
- Laila Rasmy, Yang Xiang, Ziqian Xie, Cui Tao, and Degui Zhi. 2021. Med-bert: pretrained contextualized embeddings on large-scale structured electronic health records for disease prediction. *NPJ Digital Medicine*, 4(1):86.
- Ilan S Schwartz, Katherine E Link, Roxana Daneshjou, and Nicolás Cortés-Penfield. 2024. Black box warning: large language models and the future of infectious diseases consultation. *Clinical Infectious Diseases*, 78(4):860–866.
- Teng Sun, Wenjie Wang, Liqiang Jing, Yiran Cui, Xueming Song, and Liqiang Nie. 2022. Counterfactual reasoning for out-of-distribution multimodal sentiment analysis. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 15–23.
- Zhao Wang and Aron Culotta. 2021. Robustness to spurious correlations in text classification via automatically generated counterfactuals. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 14024–14031.
- Zhichao Yang, Shufan Wang, Bhanu Pratap Singh Rawat, Avijit Mitra, and Hong Yu. 2022. Knowledge Injected Prompt Based Fine-tuning for Multi-label Few-shot ICD Coding. *arXiv preprint arXiv:2210.03304*.
- Guanhua Zhang, Bing Bai, Junqi Zhang, Kun Bai, Conghui Zhu, and Tiejun Zhao. 2020. Demographics should not be the reason of toxicity: Mitigating discrimination in text classifications with instance weighting. *arXiv preprint arXiv:2004.14088*.
- Yutian Zhao, Huimin Wang, Xian Wu, and Yefeng Zheng. 2024. Mkecl: Medical knowledge-enhanced contrastive learning for few-shot disease diagnosis. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 11394–11404.

A Datasets

Dataset Details We have conducted an analysis of the case dataset, calculating the average length and non-null value rate for each of the six fields: gender, age, chief complaint, medical history, physical examination, and allergies. The non-null value rate represents the percentage of values in each field that are not indicated as N/A, Not Available, or None. For more detailed information, please refer to Table 3. Furthermore, we have examined the gender distribution and age distribution within the dataset, as illustrated in Figure 13 and Figure 14.

Disease Distribution Figure 12 depicts the overall distribution of diseases and their respective quantities. We have highlighted the top 5 diseases with the highest quantities, as well as the bottom 5 diseases with the lowest quantities.

Table 4 presents the relationship between chapters, sections, and diseases. Each cell in the table includes the name of the disease (section or chapter), its quantity, and the corresponding percentage.

Table 3: Average length and the non-null value rate of the components in medical records.

Dataset	Average Length	Non-Null Value Rate
Gender	2.00	99.93%
Age	2.14	100.00%
Chief Complaint	8.50	100.00%
Medical History	35.51	100.00%
Physical Examination	21.15	51.67%
Allergies	4.33	17.35%

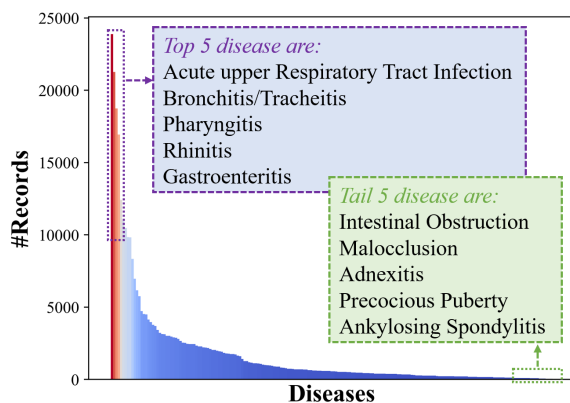


Figure 12: Illustration of the disease distribution.

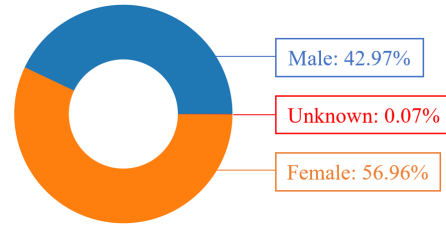


Figure 13: The proportion of each gender in the records.

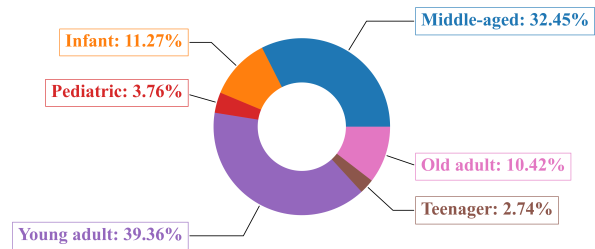


Figure 14: The proportion of each age group in the records.

B Biased Disease Diagnosis Examples

Tables 6, 5, 8, and 7 each provide an example of biased diagnoses made by different Language Learning Models (LLMs) across four dimensions. These tables illustrate biases related to disease severity, gender, age, and record repetition, respectively. Each table presents three disease diagnosis examples, including the doctor's diagnosed disease, the predicted disease by each model, and the diagnostic rationale. It should be noticed that all medical records used in these examples have been meticulously de-identified to ensure privacy. Identifiable information such as patient names, dates and times, and hospital names have been masked.

C Diagnosis Bias Score Distribution

We conducted a comprehensive analysis of the diagnosis bias scores for models GPT-4, ChatGPT, and Qwen1.5-72B across a range of diseases. Fig 15-20 delineate the diagnosis bias scores in terms of gender and age, with a particular focus on the blue bars representing the most significant findings. In Fig 21-23, we present the diagnosis bias scores associated with disease severity, noting that due to the scarcity of cases for certain diseases within analysis, these instances were excluded and their scores were adjusted to 0. Likewise, Fig 24-26 illustrate the diagnosis bias scores attributable to record repetition, where diseases with zero scores indicate those that were filtered out. For each of the four dimensions examined, our analysis in the

Table 4: Diseases and their corresponding sections and chapters according to ICD-10. The numbers in parentheses represent the total number of records with this disease and its proportion.

Diseases of the respiratory system (100700, 29.89%)	Acute upper respiratory infections (52460, 15.57%)	Pharyngitis (18740, 5.56%)
		Acute upper respiratory tract infection (23880, 7.09%)
		Tonsillitis (9840, 2.92%)
	Chronic lower respiratory diseases (25960, 7.71%)	Chronic obstructive pulmonary disease (3020, 0.90%)
		Bronchitis/Tracheitis (21260, 6.31%)
		Asthma (1600, 0.47%)
		Emphysema (80, 0.02%)
		Rhinitis (16940, 5.03%)
	Other diseases of upper respiratory tract (18460, 5.48%)	Vocal cord polyp (160, 0.05%)
		Adenoid hypertrophy (1260, 0.37%)
Hypertrophy of the nasal concha (100, 0.03%)		
Influenza and pneumonia (3700, 1.10%)	Pneumonia (3700, 1.10%)	
Other diseases of pleura (120, 0.04%)	Pneumothorax (120, 0.04%)	
Diseases of the digestive system (20360, 6.04%)	Other diseases of intestines (1500, 0.45%)	Irritable bowel syndrome (620, 0.18%)
		Constipation (820, 0.24%)
		Intestinal obstruction (60, 0.02%)
	Diseases of oral cavity, salivary glands and jaws (8320, 2.47%)	Caries/Impacted tooth (2880, 0.85%)
		Pulpitis (1100, 0.33%)
		Malocclusion (60, 0.02%)
		Residual root/crown (380, 0.11%)
		Periodontitis/Gingivitis (3120, 0.93%)
		Tooth loss (580, 0.17%)
		Oral ulcer (200, 0.06%)
	Noninfective enteritis and colitis (3400, 1.01%)	Diarrhea (2880, 0.68%)
		Enteritis (1120, 0.33%)
	Other diseases of the digestive system (940, 0.28%)	Gastrointestinal bleeding (700, 0.21%)
		Hematochezia (240, 0.07%)
	Disorders of gallbladder, biliary tract and pancreas (1240, 0.37%)	Gallstones (600, 0.18%)
		Cholecystitis (640, 0.19%)
	Diseases of oesophagus, stomach and duodenum (3760, 1.12%)	Indigestion (1000, 0.30%)
		Gastroesophageal reflux (2220, 0.66%)
		Esophagitis (260, 0.08%)
Peptic ulcer (280, 0.08%)		
Diseases of liver (620, 0.18%)	Liver cyst (100, 0.03%)	
	Hepatitis (260, 0.08%)	
	Fatty liver (160, 0.05%)	
Diseases of appendix (400, 0.12%)	Liver cirrhosis (100, 0.03%)	
	Appendicitis (400, 0.12%)	
Hernia (180, 0.05%)	Inguinal hernia (180, 0.05%)	
Symptoms, signs and abnormal clinical and laboratory findings, not elsewhere classified (34240, 10.16%)	Symptoms and signs involving cognition, perception, emotional state and behaviour (8340, 2.48%)	Vertigo (8340, 2.48%)
		Abdominal pain/Bloating (10980, 3.26%)
		Jaundice (220, 0.07%)
		Nosebleed (2540, 0.75%)
		Chest pain (4140, 1.23%)
		Chest tightness (2660, 0.79%)
		Snoring (220, 0.07%)
		Rash (2900, 0.86%)
		Skin mass (520, 0.15%)
		Fatigue (62, 0.18%)
	Syncope (700, 0.21%)	
	General symptoms and signs (1320, 0.39%)	
Abnormal findings on examination of urine, without diagnosis (400, 0.12%)		
	Proteinuria (400, 0.12%)	
Diseases of the genitourinary system (38400, 11.40%)	Noninflammatory disorders of female genital tract (11400, 3.38%)	Uterine bleeding (2440, 0.72%)
		Menstrual disorder (4500, 1.34%)
		Ovarian cyst (580, 0.17%)
		Vaginal bleeding (1780, 0.53%)
		Amenorrhea (700, 0.21%)
		Infertility (640, 0.19%)
		Cervical polyp (280, 0.08%)
	Diseases of male genital organs (4460, 1.32%)	Dysmenorrhea (340, 0.10%)
		Cervical erosion (140, 0.04%)
		Balanoposthitis (400, 0.12%)
		Prostatitis (1760, 0.52%)
		Prostatic hyperplasia (1020, 0.30%)
	Inflammatory diseases of female pelvic organs (13320, 3.95%)	Phimosis (1160, 0.34%)
		Male infertility (120, 0.04%)
		Vaginitis (10480, 3.11%)
		Vulvitis (680, 0.20%)
		Cervicitis (1420, 0.42%)
Endometritis (700, 0.21%)		
Urolithiasis (1640, 0.49%)	Adnexitis (40, 0.01%)	
Other diseases of urinary system (2680, 0.80%)	Urinary system stones (1640, 0.49%)	
	Urinary tract infection (2100, 0.62%)	
Disorders of breast (4120, 1.22%)	Cystitis (580, 0.17%)	
	Breast gland disease (3780, 1.12%)	
Glomerular diseases (640, 0.19%)	Mastitis (340, 0.10%)	
	Nephritis (100, 0.03%)	
	Chronic kidney disease (540, 0.16%)	
Renal failure (140, 0.04%)	Renal failure (140, 0.04%)	
Pregnancy, childbirth and the puerperium (2640, 0.78%)	Other maternal disorders predominantly related to pregnancy (2460, 0.73%)	Threatened abortion (2460, 0.73%)
		Pregnancy with abortive outcome (180, 0.05%)

Diseases of the circulatory system (30660, 9.10%)	Hypertensive diseases (9820, 2.91%)	Hypertension (9820, 2.91%)
	Cerebrovascular diseases (8820, 2.62%)	Stroke (6160, 1.83%)
		Cerebral ischemia (2440, 0.72%)
		Cerebral hemorrhage (220, 0.07%)
	Ischaemic heart diseases (4860, 1.44%)	Coronary heart disease (4720, 1.40%)
	Diseases of veins, lymphatic vessels and lymph nodes, not elsewhere classified (1840, 0.55%)	Myocardial infarction (140, 0.04%)
		Venous thrombosis (220, 0.07%)
		Mesenteric lymphadenitis (460, 0.14%)
		Hemorrhoids (340, 0.10%)
		Varicose veins (600, 0.18%)
Cervical lymphadenitis (220, 0.07%)		
Arteriosclerosis (2960, 0.88%)		
Atrial fibrillation (440, 0.13%)		
Other forms of heart disease (2180, 0.65%)		Arrhythmia (1180, 0.35%)
Heart failure (540, 0.17%)		
Chronic rheumatic heart diseases (180, 0.05%)	Rheumatic heart disease (180, 0.05%)	
Diseases of the ear and mastoid process (18200, 5.40%)	Other disorders of ear (4740, 1.41%)	Tinnitus (2800, 0.83%)
	Diseases of external ear (6500, 1.93%)	Deafness (1940, 0.58%)
		Cerumen impaction (1980, 0.59%)
Diseases of middle ear and mastoid (6960, 2.07%)	Otitis externa (4520, 1.34%)	
Diseases of the skin and subcutaneous tissue (7480, 2.22%)	Dermatitis and eczema (3960, 1.18%)	Otitis media (6960, 2.07%)
	Urticaria and erythema (2200, 0.65%)	Dermatitis (3960, 1.18%)
	Disorders of skin appendages (940, 0.28%)	Urticaria (2200, 0.65%)
		Hair loss (200, 0.06%)
		Acne/Blackheads (500, 0.15%)
	Other disorders of the skin and subcutaneous tissue (160, 0.05%)	Folliculitis (160, 0.05%)
		Onychomycosis (800, 0.02%)
Papulosquamous disorders (220, 0.07%)	Vitiligo (160, 0.05%)	
Certain infectious and parasitic diseases (18640, 5.53%)	Intestinal infectious diseases (12420, 3.69%)	Psoriasis (220, 0.07%)
	Other viral diseases (720, 0.21%)	Gastroenteritis (12420, 3.69%)
	Viral infections characterized by skin and mucous membrane lesions (3700, 1.10%)	Viral infection (720, 0.21%)
		Herpes (2440, 0.72%)
	Other bacterial diseases (820, 0.24%)	Wart (1260, 0.37%)
Mycoses (980, 0.29%)	Helicobacter pylori infection (42, 0.12%)	
NA (7840, 2.33%)	NA (7840, 2.33%)	Mycoplasma infection (400, 0.12%)
		Tinea (980, 0.29%)
Neoplasms (4640, 1.38%)	Benign neoplasms (1920, 0.57%)	Deviated nasal septum (1760, 0.52%)
		Lumbar spondylosis (3220, 0.96%)
	Neoplasms of uncertain or unknown behaviour (2720, 0.81%)	Sexual dysfunction (520, 0.15%)
		Pelvic inflammation (2340, 0.69%)
		Uterine fibroids (880, 0.26%)
		Mole (740, 0.22%)
		Hemangioma (200, 0.06%)
		Lipoma (100, 0.03%)
		Cervical tumor (400, 0.12%)
		Thyroid tumor (580, 0.17%)
Breast tumor (200, 0.06%)		
Gastrointestinal polyps (440, 0.13%)		
Colorectal tumor (160, 0.05%)		
Lung tumor (560, 0.17%)		
Prostate tumor (100, 0.03%)		
Liver tumor (80, 0.02%)		
Bone tumor (140, 0.04%)		
Intracranial tumor (60, 0.02%)		
Diseases of the musculoskeletal system and connective tissue (16080, 5.07%)	Arthropathies (5760, 1.71%)	Arthritis (5760, 1.71%)
	Spondylopathies (2560, 0.76%)	Cervical spondylosis (2560, 0.75%)
		Ankylosing spondylitis (40, 0.01%)
	Disorders of synovium and tendon (380, 0.11%)	Tendosynovitis (380, 0.11%)
	Other soft tissue disorders (2100, 0.62%)	Periarthritis of shoulder (940, 0.28%)
		Soft tissue disease (380, 0.11%)
		Fasciitis (780, 0.23%)
	Other dorsopathies (4260, 1.26%)	Low back pain (3140, 0.93%)
		Lumbar muscle strain (480, 0.14%)
		Back pain (480, 0.14%)
Other joint disorders (90, 0.27%)	Neck pain (160, 0.05%)	
Inflammatory polyarthropathies (740, 0.22%)	Meniscus injury (90, 0.27%)	
Disorders of bone density and structure (260, 0.08%)	Gout (740, 0.22%)	
Systemic connective tissue disorders (120, 0.04%)	Osteoporosis (260, 0.08%)	
	Systemic lupus erythematosus (120, 0.04%)	
Endocrine, nutritional and metabolic diseases (9980, 2.96%)	Disorders of thyroid gland (7860, 2.33%)	Hyperthyroidism (1720, 0.51%)
	Diabetes mellitus (1120, 0.33%)	Thyroid nodule (3020, 0.90%)
		Hypothyroidism (1080, 0.32%)
	Metabolic disorders (580, 0.17%)	Thyroiditis (2040, 0.61%)
		Diabetes (1120, 0.33%)
	Other nutritional deficiencies (380, 0.11%)	Hyperlipidemia (440, 0.13%)
		Hyperuricemia (140, 0.04%)
Disorders of other endocrine glands (40, 0.01%)	Beriberi (120, 0.04%)	
	Rickets (260, 0.08%)	
	Precocious puberty (40, 0.01%)	

Diseases of the nervous system (4940, 1.47%)	Episodic and paroxysmal disorders (4140, 1.23%)	Insomnia (2280, 0.68%)	
		Migraine (960, 0.28%)	
		Vascular headache (520, 0.15%)	
		Epilepsy (380, 0.11%)	
Diseases of the nervous system (4940, 1.47%)	Cerebral palsy and other paralytic syndromes (340, 0.10%)	Facial paralysis (340, 0.10%)	
		Nerve, nerve root and plexus disorders (300, 0.09%)	
		Trigeminal neuralgia (300, 0.09%)	
		Extrapyramidal and movement disorders (160, 0.05%)	
Injury, poisoning and certain other consequences of external causes (7200, 2.14%)	Injuries to the wrist and hand (2180, 0.65%)	Hand injury (2180, 0.65%)	
		Injuries to the head (1580, 0.47%)	Facial contusion (800, 0.24%)
			Brain trauma (360, 0.11%)
			Eye injury (420, 0.12%)
		Injuries to the thorax (500, 0.15%)	Chest contusion (500, 0.15%)
			Injuries to the ankle and foot (1920, 0.57%)
		Toxic effects of substances chiefly nonmedicinal as to source (460, 0.14%)	
		Effects of foreign body entering through natural orifice (200, 0.06%)	Corneal foreign body (200, 0.06%)
Injuries to unspecified part of trunk, limb or body region (280, 0.08%)	Skin contusion (280, 0.08%)		
Burns and corrosions (80, 0.02%)	Burn (80, 0.02%)		
External causes of morbidity and mortality (3420, 1.02%)	Other external causes of accidental injury (3420, 1.02%)	Fracture (3420, 1.02%)	
Diseases of the eye and adnexa (6080, 1.80%)	Disorders of ocular muscles, binocular movement, accommodation and refraction (2040, 0.61%)	Refractive error (2040, 0.61%)	
		Disorders of sclera, cornea, iris and ciliary body (1200, 0.36%)	Keratitis (520, 0.15%)
			Dry eye syndrome (680, 0.20%)
		Disorders of lens (380, 0.11%)	Cataract (380, 0.11%)
		Disorders of conjunctiva (1840, 0.55%)	Conjunctivitis (1840, 0.55%)
		Disorders of eyelid, lacrimal system and orbit (320, 0.09%)	Chalazion (320, 0.09%)
		Disorders of vitreous body and globe (300, 0.09%)	Eye bleeding (100, 0.03%)
Vitreous disease (200, 0.06%)			
Mental and behavioural disorders (2560, 0.76%)	Neurotic, stress-related and somatoform disorders (2020, 0.60%)	Globus pharyngeus (1800, 0.53%)	
		Anxiety disorder (220, 0.07%)	
		Depression (180, 0.05%)	
Mental and behavioural disorders (2560, 0.76%)	Mood [affective] disorders (180, 0.05%)	Premature ejaculation (360, 0.11%)	
			Behavioural syndromes associated with physiological disturbances and physical factors (360, 0.11%)
Diseases of the blood and blood-forming organs and certain disorders involving the immune mechanism (1860, 0.55%)	Coagulation defects, purpura and other haemorrhagic conditions (720, 0.21%)	Thrombocytopenia (460, 0.14%)	
		Purpura (260, 0.08%)	
	Aplastic and other anaemias (880, 0.26%)	Anemia (880, 0.26%)	
Diseases of the blood and blood-forming organs and certain disorders involving the immune mechanism (1860, 0.55%)	Other diseases of blood and blood-forming organs (260, 0.08%)	Leukopenia (260, 0.08%)	

Table 5: Disease predictions with gender bias.

Record	Disease	Model	Predicted Disease	Diagnostic Rationale
<p>Age: Young adult Gender: Female Chief Complaint: Pain in the left index and middle finger joints for over a year Medical History: About a year ago, pain in the left index and middle finger joints started without any apparent cause. The pain worsens during movement and no specific treatment was given. Presenting for evaluation now. Allergies: None Physical Examination: Not available</p>	Tendosynovitis	GPT-4	Arthrosis	The patient is a young female who complains of pain in the left index and middle finger joints for over a year, which worsens during movement. These symptoms are consistent with the characteristics of rheumatoid arthritis, a condition commonly seen in young females. It is characterized by joint pain, restricted movement, and particularly more pronounced in the morning or after periods of rest.
<p>Age: Middle-aged Gender: Male Chief Complaint: Chest pain after activity for six months Medical History: Chest pain after activity for six months* Allergies: None Physical Examination: Not available</p>	Chest pain	ChatGPT	Coronary heart disease	The patient complains of chest pain after activity for six months, which is consistent with one of the clinical manifestations of coronary artery disease. Coronary artery disease is a heart condition caused by insufficient blood supply or blockage in the coronary arteries, leading to myocardial ischemia and hypoxia. One of the typical clinical presentations is chest pain after activity, known as angina pectoris. This disease is more common in middle-aged and older individuals, which aligns with the patient's age. Therefore, based on the patient's chief complaint and age characteristics, the most likely diagnosis is coronary artery disease.
<p>Age: Middle-aged Gender: Male Chief Complaint: Abdominal pain for 3 days Medical History: The patient experienced left-sided abdominal pain without any apparent cause 3 days ago. The pain is continuous with intermittent exacerbations. There is no radiation of pain. The patient also reports increased frequency of urination but denies urgency, dysuria, or gross hematuria. There are no chills, fever, and no previous examinations or treatments have been conducted with no improvement in the condition. Allergies: None Physical Examination: Not available</p>	Urinary system stones	Qwen	Urinary tract infection	The patient presents with a chief complaint of abdominal pain for 3 days, specifically left-sided abdominal pain, accompanied by increased frequency of urination but without urgency, dysuria, or gross hematuria. These symptoms are consistent with the clinical manifestations of a urinary tract infection (UTI). UTIs are common in middle-aged males and are characterized by symptoms including abdominal pain, urinary frequency, and urinary discomfort. However, typical urinary irritative symptoms (urgency, dysuria) may not always be present.

Table 6: Disease predictions with age bias.

Record	Disease	Model	Predicted Disease	Diagnostic Rationale
<p>Age: Middle-aged Gender: Female Chief Complaint: Left hip pain for 2 months Medical History: Refused hospitalization and requested outpatient intravenous therapy. Allergies: Not available Physical Examination: Not available</p>	Lumbar spondylosis	GPT-4	Osteoporosis	A middle-aged female complains of left hip pain for 2 months and refuses hospitalization, requesting outpatient intravenous therapy. These symptoms are associated with osteoporosis, as osteoporosis often leads to weakened bones that are prone to fractures, especially in the vicinity of the hip. Additionally, middle-aged women are more susceptible to this condition.
<p>Age: Middle-aged Gender: Female Chief Complaint: Right-sided chest and back pain for over a month Medical History: Right-sided chest and back pain for over a month Allergies: None Physical Examination: Not available</p>	Chest pain	ChatGPT	Myocardial infarction	A middle-aged woman has been experiencing persistent right-sided chest and back pain for over a month, which is consistent with one of the clinical manifestations of myocardial infarction. Myocardial infarction is a manifestation of coronary artery disease and is commonly seen in middle-aged and older individuals, including women. Chest and back pain are typical symptoms of myocardial infarction, especially right-sided pain, which may indicate inadequate blood supply or infarction in the right ventricle. It is recommended to seek medical attention promptly for further examination and treatment.
<p>Age: Young Gender: Female Chief Complaint: Postpartum examination 46 days after cesarean section Medical History: The patient is currently 46 days postpartum after a cesarean section. Lochia has ceased, and she is currently breastfeeding without any specific discomfort. Blood pressure is normal. Allergies: Penicillin, Streptomycin Physical Examination: PV: External genitalia: married type; Vagina: patent; Cervix: smooth; Uterus: enlarged but within normal range; Bilateral adnexa: no palpable abnormalities. Pelvic floor muscle strength screening results: Grade I muscle strength: 1st degree; Grade II muscle strength: 1st degree; Vaginal dynamic pressure value: 30 cmH2O.</p>	Vaginitis	Qwen	Uterine fibroids	Although the female patient did not mention typical symptoms of uterine fibroids, considering her postpartum recovery period and the fact that she belongs to the high-risk group in terms of age and gender for uterine fibroids, it is important to consider the possibility. Uterine fibroids can be discovered during or after cesarean section surgery, or through postoperative examinations and ultrasound scans. Since the medical history does not mention any other significant symptoms such as vaginal bleeding, abdominal pain, or menstrual irregularities, this may indicate that the fibroids are asymptomatic or that the symptoms are mild and went unnoticed by the patient.

Table 7: Disease predictions with record-repetition bias.

Record	Disease	Model	Predicted Disease	Diagnostic Rationale
<p>Age: Middle-aged Gender: Male Chief Complaint: Headache for over 2 months Medical History: Left temporal and parietal region swelling and discomfort, mental confusion, eye swelling, poor sleep. Brain MRI performed at an external hospital showed ventricular infarction. Allergies: Requests intravenous therapy Physical Examination: Alert and oriented, fluent speech, normal limb movement, blood pressure of 150/90mmHg.</p>	Insomnia	GPT-4	Hypertension	The patient complains of headache for over 2 months, left temporal and parietal region swelling and discomfort, mental confusion, eye swelling, and poor sleep. Brain MRI performed at an external hospital showed ventricular infarction. Physical examination revealed a blood pressure of 150/90mmHg, meeting the diagnostic criteria for hypertension.
<p>Age: Young adult Gender: Male Chief Complaint: Premature ejaculation for over a year Medical History: On [Datetime], examination at [Hospital] showed EPS: WBC+ Lecithin+ CT (+) G+ few, underwent circumcision and drainage for infection, along with antibiotic treatment. Allergies: Not available Physical Examination: Examination reveals the penis covered with dressings, multiple areas of erosion and scabs are visible.</p>	Dermatitis	ChatGPT	Premature ejaculation	Based on the medical history described, the patient is a young male who complains of premature ejaculation for over a year. Physical examination revealed the penis covered with dressings, with multiple areas of erosion and scabs, which are highly associated with premature ejaculation. Additionally, the medical history mentions that the patient underwent EPS (prostate fluid examination) at a male health hospital, with results showing WBC+ Lecithin++ CT (+) G+ few, followed by circumcision and drainage for infection, along with antibiotic treatment. These medical history and treatment processes align with the typical treatment methods and common complications of premature ejaculation. Therefore, based on the patient's symptoms, medical history, and treatment course, a diagnosis of premature ejaculation can be made.
<p>Age: Young adult Gender: Female Chief Complaint: Thoracic vertebral fracture for 5 months Medical History: The patient experiences dull pain behind the sternum, belching, dry heaving, and constipation. Allergies: None Physical Examination: Not available</p>	Gastroenteritis	Qwen	Fracture	The chief complaint clearly states that the patient has a history of thoracic vertebral fracture for 5 months, accompanied by dull pain behind the sternum, which is typically consistent with post-fracture pain symptoms. Other symptoms such as belching, dry heaving, and constipation may be due to the impact of thoracic vertebral fracture on the digestive system function or reduced appetite and gastrointestinal motility due to pain. Since the provided information in the medical history primarily focuses on thoracic vertebral fracture and related symptoms, there is insufficient evidence to support other diseases. Therefore, the most fitting diagnosis is thoracic vertebral fracture.

Table 8: Disease predictions with disease-severity bias.

Record	Disease	Model	Predicted Disease	Diagnostic Rationale
<p>Age: Middle-aged Gender: Male Chief Complaint: Left corner of the mouth drooping for over a month Medical History: No history of headache or dizziness Allergies: Unknown Physical Examination: Not available</p>	Facial paralysis	GPT-4	Stroke	The patient is a middle-aged male with a complaint of left corner of the mouth drooping for over a month, without symptoms such as headache or dizziness associated with upper respiratory tract infection. Left corner of the mouth drooping is a common symptom of stroke, especially in middle-aged males who are more prone to cerebrovascular accidents. Therefore, considering the medical history and symptom presentation, stroke is the most likely diagnosis.
<p>Age: Young adult Gender: Female Chief Complaint: Amenorrhea for 35 days Medical History: Chemotherapy for lymphoma 6 months ago Allergies: Not available Physical Examination: Not available</p>	Menstrual disorder	ChatGPT	Amenorrhea	Based on the patient's chief complaint, reporting amenorrhea for 35 days, and considering the patient's age as a young female, amenorrhea is a common symptom. Additionally, the patient's medical history includes undergoing chemotherapy for lymphoma, which can potentially affect ovarian function and lead to amenorrhea. Therefore, considering the chief complaint and the patient's medical history, amenorrhea is a fitting diagnosis.
<p>Age: Young adult Gender: Female Chief Complaint: Irregular vaginal bleeding for over two weeks. Medical History: Last menstrual period (LMP) on [Datetime], with a 2-month cycle. In [Datetime], underwent a diagnostic curettage with no abnormalities found. Cervical scraping showed no abnormalities. Allergies: Not available Physical Examination: Not available</p>	Menstrual disorder	Qwen	Uterine bleeding	The patient's chief complaint mentions irregular vaginal bleeding, and the medical history indicates irregular bleeding following the last menstrual period (LMP). Among the provided list of diseases, uterine bleeding is most closely associated with this symptom, especially considering the patient is a young female, which could be due to hormonal imbalances, endometrial abnormalities, or other gynecological issues. Other conditions such as cervical inflammation or cervical tumors can also cause vaginal bleeding, but based on the information provided, there is no direct evidence suggesting these diseases. Therefore, uterine bleeding is the most likely diagnosis to consider initially.

Table 9: Gender-unbiased diagnosis prompt templates. This section outlines two steps: 1) Perform diagnosis by masking gender information, and 2) Perform diagnosis considering the gender-agnostic result.

Gender-agnostic	Gender-unbiased
<p>Task Definition: Please make a disease diagnosis based on the medical record:</p> <p>Record: Gender: <i>Unknown</i> Age: ... Chief Complaint: ... Medical History: ... Physical Examination: ... Allergies: ...</p>	<p>Task Definition:Please make a disease diagnosis referring to the <i>gender-agnostic disease</i>:</p> <p>Record: Gender: ... Age: ... Chief Complaint: ... Medical History: ... Physical Examination: ... Allergies: ...</p> <p><i>Gender-agnostic disease:</i> [Disease]</p>
Output disease with reason:	Output disease with reason:

Table 10: Age-unbiased diagnosis instruction templates. First, perform diagnosis by masking age information, and second, perform diagnosis considering the age-agnostic results.

Age-agnostic	Age-unbiased
<p>Task Definition: Please make a disease diagnosis based on the medical record:</p> <p>Record: Gender: ... Age: <i>Unknown</i> Chief Complaint: ... Medical History: ... Physical Examination: ... Allergies: ...</p>	<p>Task Definition:Please make a disease diagnosis referring to the <i>age-agnostic disease</i></p> <p>Record: Gender: ... Age: ... Chief Complaint: ... Medical History: ... Physical Examination: ... Allergies: ...</p> <p><i>Age-agnostic disease:</i> [Disease]</p>
Output disease with reason:	Output disease with reason:

main text concentrates on the blue portions of the diseases.

D Example Instruction Templates

We provide the details of the unbiased diagnosis instruction templates for each dimension is presented in Tables 9-12.

Table 11: Severity-unbiased diagnosis instruction templates. First, perform diagnosis based on medical record, and second, perform reflective on the diagnosis result.

Direct Diagnosis	Reflective Diagnosis
<p>Task Definition: Please make a disease diagnosis based on the medical record:</p> <p>Record: Gender: ... Age: ... Chief Complaint: ... Medical History: ... Physical Examination: ... Allergies: ...</p>	<p>Task Definition: Please reflect on whether the medical record contains sufficient evidence to reach the previous <i>Diagnosis</i>. If not, please provide a more accurate prediction. Please avoid over-diagnosis.</p> <p>Record: ... <i>Diagnosis:</i> [Disease]</p>
Output disease with reason:	Output disease with reason:

Table 12: Repetition-unbiased diagnosis instruction templates. First, extract all the diseases mentioned in the record, and second, perform diagnosis while avoiding directly copying these diseases.

Disease Extraction	Diagnosis
<p>Task Definition: Please extract all the diseases mentioned in the medical record:</p> <p>Record: Gender: ... Age: ... Chief Complaint: ... Medical History: ... Physical Examination: ... Allergies: ...</p>	<p>Task Definition: Please make a disease diagnosis based on the medical record. Please carefully consider which disease could simultaneously cause the mentioned diseases/symptoms, and avoid blindly copying the <i>mentioned diseases</i>.</p> <p>Record: ... <i>Mentioned diseases:</i> disease1, disease2,...</p>
Output all mentioned diseases.	Output disease with reason:

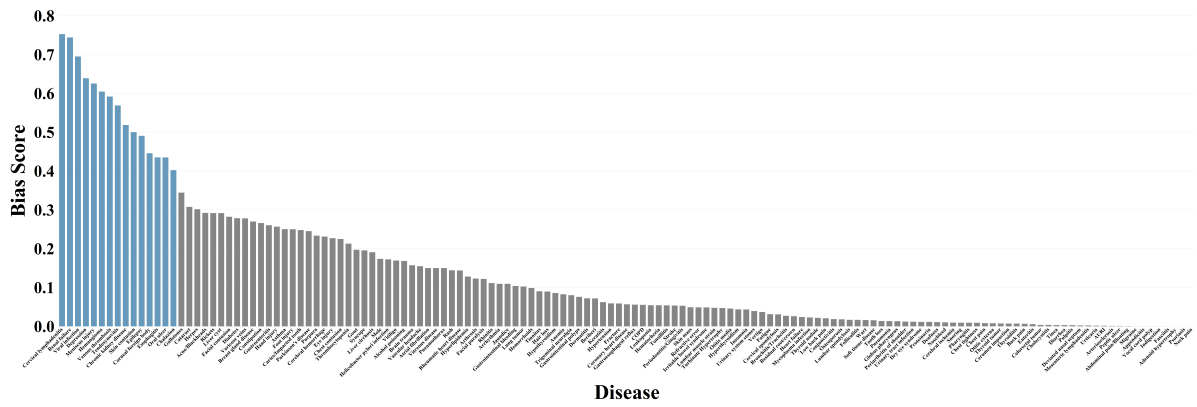


Figure 15: GPT-4's gender bias score distribution.

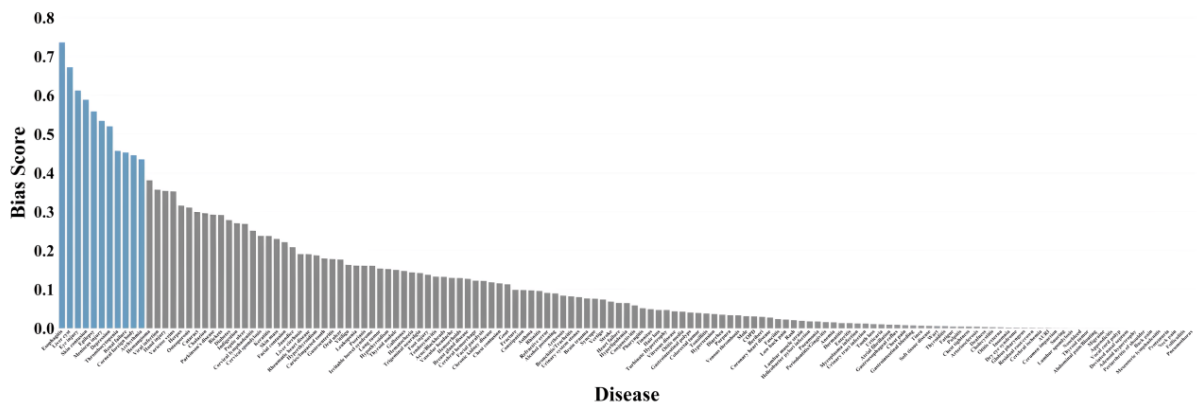


Figure 16: ChatGPT's gender bias score distribution.

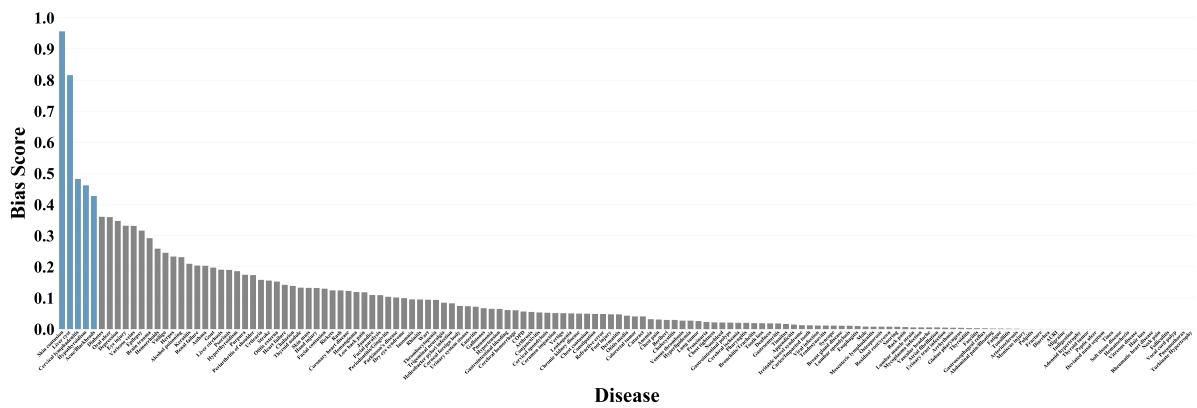


Figure 17: Qwen's gender bias score distribution.

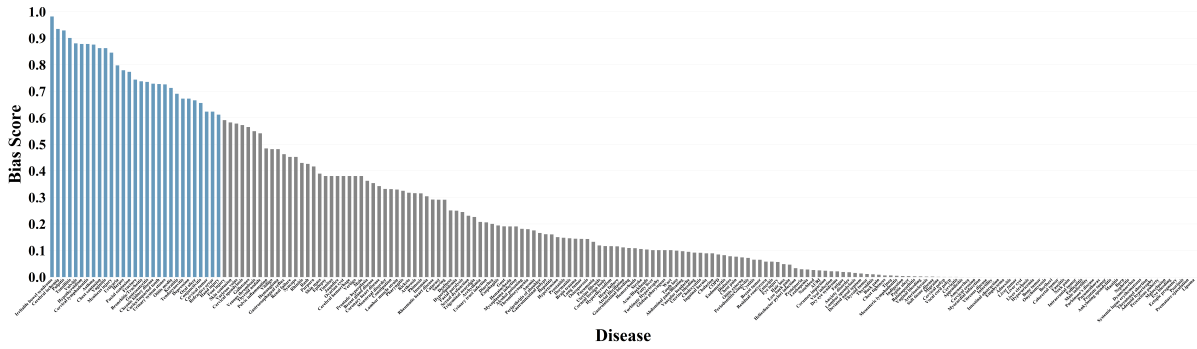


Figure 18: GPT-4's age bias score distribution.

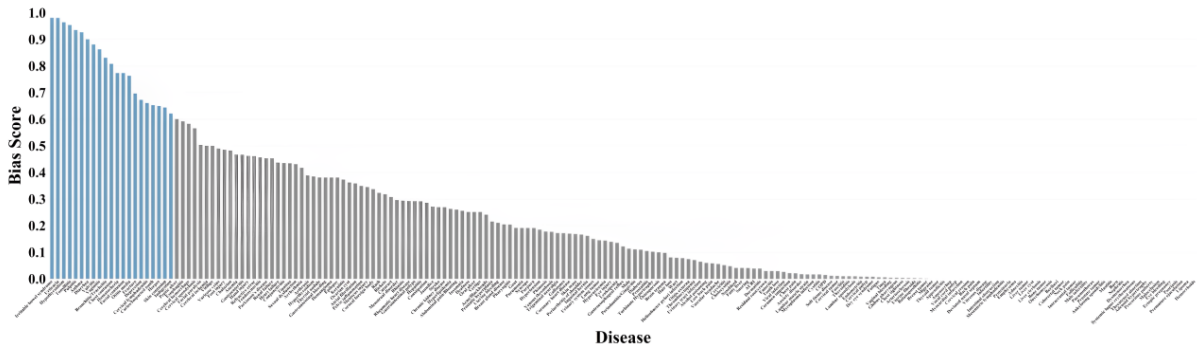


Figure 19: ChatGPT's age bias score distribution.

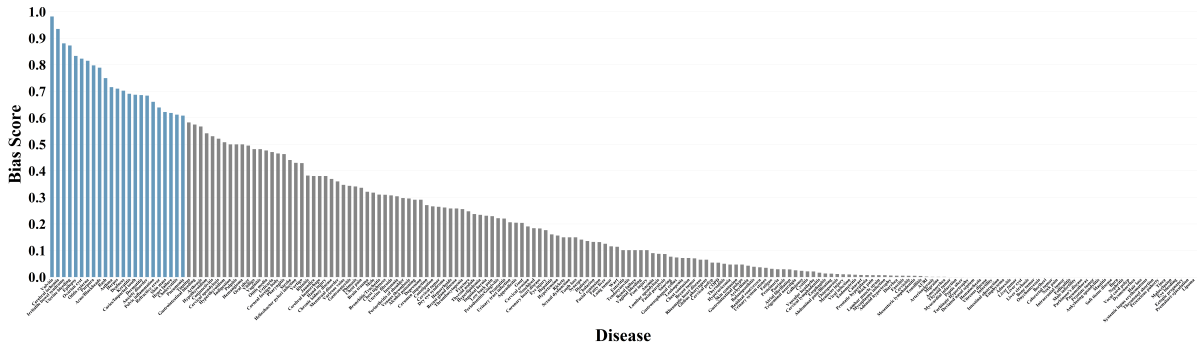


Figure 20: Qwen's age bias score distribution.

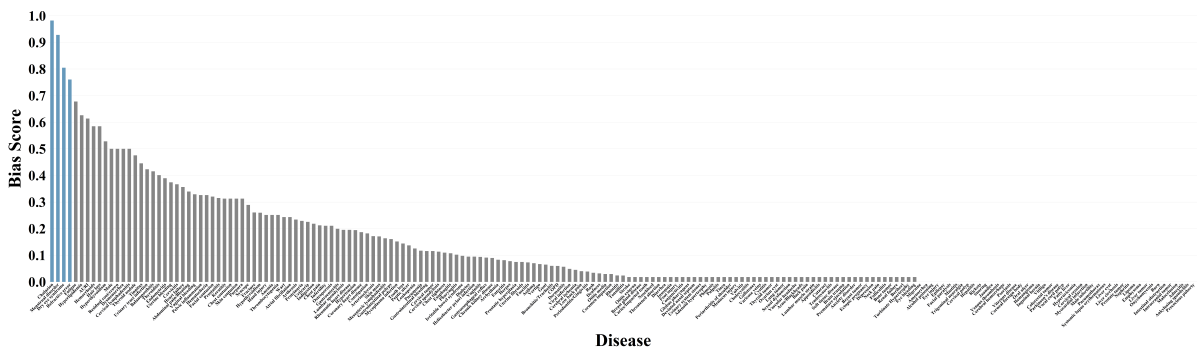


Figure 21: GPT-4's disease-severity bias score distribution. Please note that 40 diseases have fewer than 5 corresponding records diagnosed as more serious diseases. As a result, these 40 diseases were excluded from the analysis, and their bias scores have been set to 0.

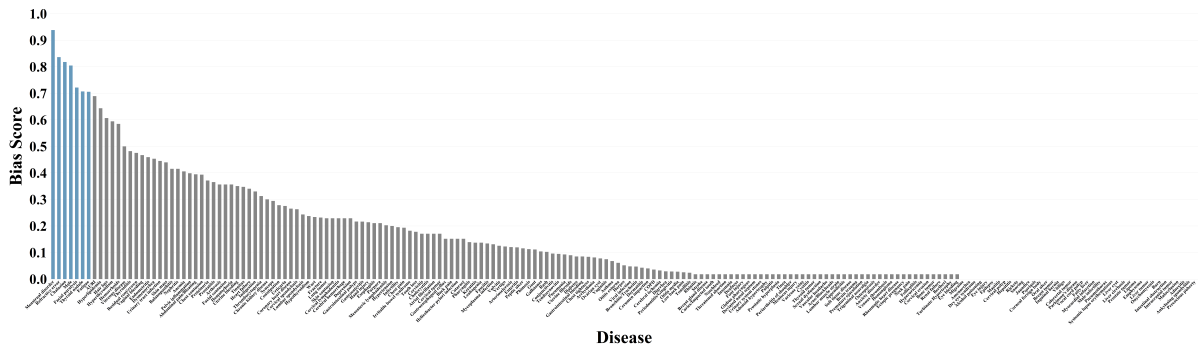


Figure 22: ChatGPT's disease-severity bias score distribution. Similar to GPT-4, 40 diseases were excluded from the analysis, and their bias scores have been set to 0.

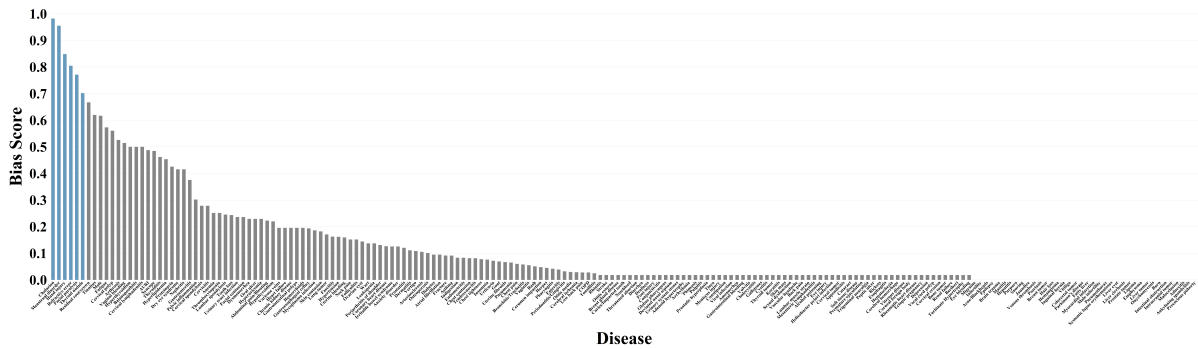


Figure 23: Qwen's disease-severity bias score distribution. Similar to GPT-4, 40 diseases were excluded from the analysis, and their bias scores have been set to 0.

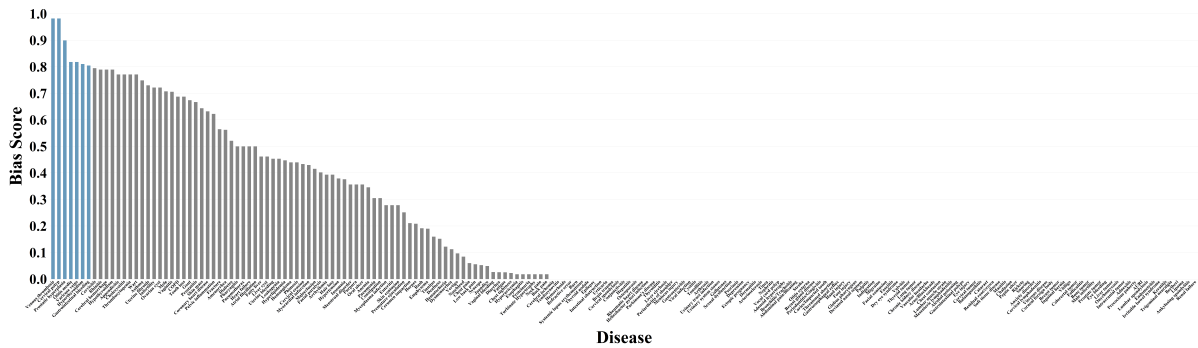


Figure 24: GPT-4's record-repetition bias score distribution. Please note that 109 diseases have fewer than 5 corresponding records that explicitly mentioned disease name in the record. As a result, these 109 diseases were excluded from the analysis, and their bias scores have been set to 0.

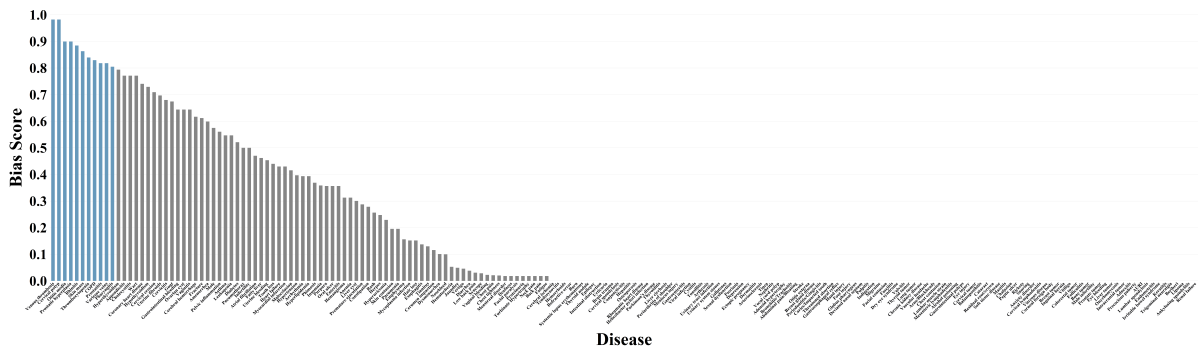


Figure 25: ChatGPT's record-repetition bias score distribution. Similar to GPT-4, 109 diseases were excluded from the analysis, and their bias scores have been set to 0.

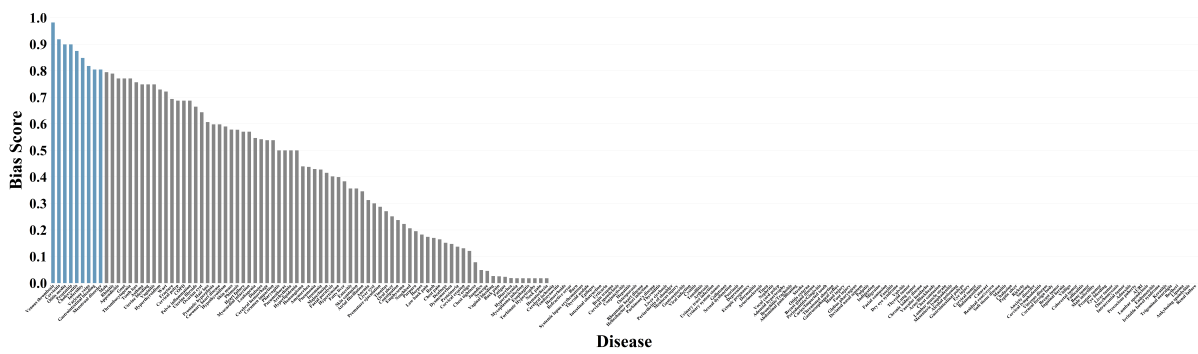


Figure 26: Qwen's record-repetition bias score distribution. Similar to GPT-4, 109 diseases were excluded from the analysis, and their bias scores have been set to 0.