

Review

Not peer-reviewed version

---

# Bias in Large AI Models for Medicine and Healthcare: Survey and Challenges

---

[Ying Xiao](#), Zhenpeng Chen<sup>\*</sup>, Jen-tse Huang, Wenting Chen, Yepang Liu, [Kezhi Li](#), Mohammad Reza Mousavi, Richard Dobson, Jie M. Zhang

Posted Date: 24 November 2025

doi: 10.20944/preprints202511.1838.v1

Keywords: AI medicine; AI healthcare; medical bias; LLMs; bias mitigation



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

# Bias in Large AI Models for Medicine and Healthcare: Survey and Challenges

Ying Xiao <sup>1</sup>, Zhenpeng Chen <sup>2,\*</sup>, Jen-tse Huang <sup>3</sup>, Wenting Chen <sup>4</sup>, Yepang Liu <sup>5</sup>, Kezhi Li <sup>6</sup>,  
Mohammad Reza Mousavi <sup>1</sup>, Richard Dobson <sup>1</sup> and Jie M. Zhang <sup>1</sup>

<sup>1</sup> King's College London

<sup>2</sup> Nanyang Technological University

<sup>3</sup> Johns Hopkins University

<sup>4</sup> Stanford University

<sup>5</sup> Southern University of Science and Technology

<sup>6</sup> University College London

\* Correspondence: zhenpeng.chen@ntu.edu.sg

## Abstract

Large AI models have demonstrated human-expert-level performance in specific medical domains. However, concerns regarding medical bias have prompted growing attention from the medicine, sociology, and computer science communities. Although research on medical bias in large AI models is rapidly expanding, efforts remain fragmented, often shaped by discipline-specific assumptions, terminology, and evaluation criteria. This survey provides a comprehensive synthesis of 55 representative studies, organizing the literature into three core themes: taxonomy of medical bias, methods for detection, and strategies for mitigation. Our analysis bridges the conceptual and methodological gaps across disciplines and highlights persistent challenges, including the lack of unified foundations for medical fairness, insufficient datasets and evaluation benchmarks, the lack of methods for rigorous automatic bias detection, missing real-world validation and continuous validation, inadequate representation, as well as insufficient studies on the trade-off between fairness and accuracy. Thereby, we identify and highlight emerging research opportunities to address these gaps. To further advance the field, we present a structured index of publicly available models and datasets referenced in these studies.

**Keywords:** AI medicine; AI healthcare; medical bias; LLMs; bias mitigation

## 1. Introduction

Artificial Intelligence (AI) is increasingly integral to modern healthcare. Among these technologies, large AI models, including large language models (LLMs) and large vision models (LVMs), have recently emerged as especially influential [1]. Trained on massive corpora, these advanced models, exemplified by ChatGPT, demonstrate human-like understanding and generation across multiple data modalities [1,2]. Large, general-purpose AI models are now being deployed or studied across a spectrum of healthcare applications [3–7], including but not limited to:

- Clinical decision support – Assisting the clinicians with clinical diagnostic reasoning or treatment planning based on large knowledge bases and patient-specific data [8–10].
- Patient communication – Powering conversational agents or chatbots that answer patient questions, provide triage advice, or offer health counseling in natural language [11–13].
- Medical documentation – Extracting, summarizing, interpreting, or generating clinical notes and medical reports, therefore reducing the documentation burden of clinicians [8,14,15].
- Medical education – Producing understandable medical explanations, personalized education materials, or providing specific training simulation for medical professionals [8,16,17].

These developments highlight the potential to transform the paradigm of modern medicine and healthcare, increasing efficiency and accessibility, particularly in under-resourced areas where high-quality medical expertise is limited [18–20].

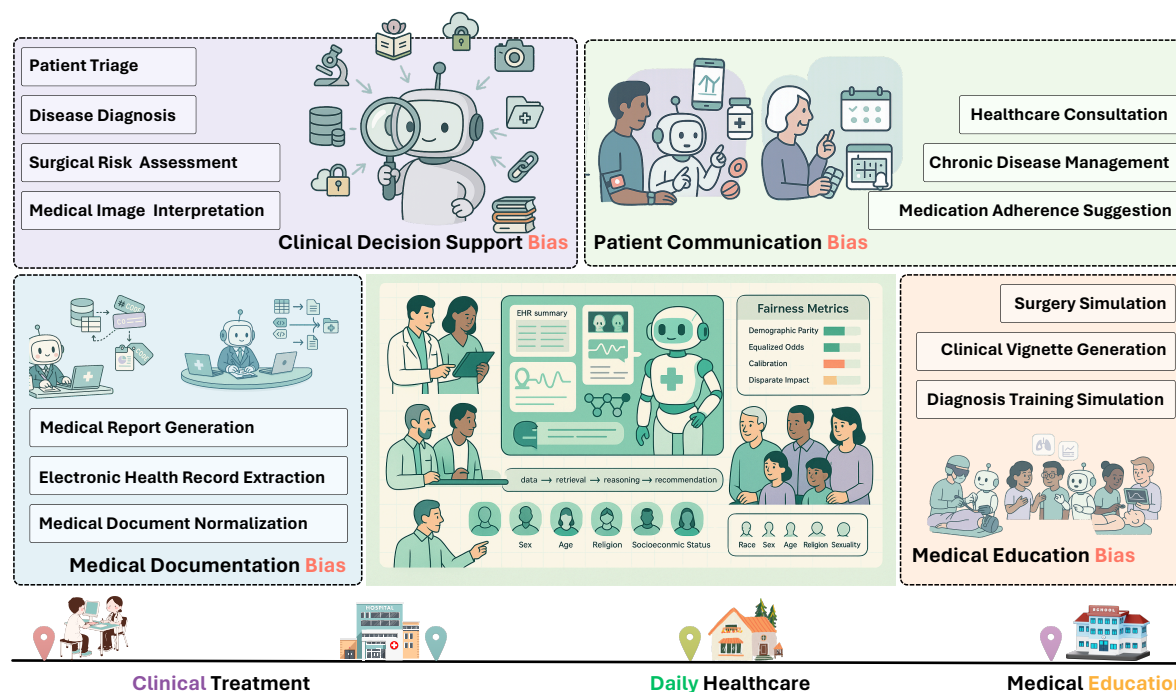


Figure 1. An overview of bias in Large AI model for medicine and healthcare.

However, as these models gain wider adoption in medical and healthcare contexts, concerns regarding their trustworthiness have become increasingly prominent, with medical bias emerging as a critical issue [21,22]. In this survey, we define bias as **any systematic error, stereotype, or prejudice in an AI system's outputs that disadvantages certain individuals or groups**, thereby undermining fairness. A wide range of bias issues in large AI models for health have been reported. For instance, Czum and Parr [23] found that a model demonstrated significantly lower diagnostic performance for female patients compared to males when detecting cardiomegaly, resulting in unequal access to accurate clinical decision support. Such biases raise ethical concerns and introduce safety risks that can negatively affect patient care and exacerbate existing health disparities [1,24]. Similar demographic gaps have been documented across racial and socioeconomic groups, further threatening trust in AI-powered healthcare [11,12,25–27].

While these biases are deeply concerning, it is important to recognize that they often originate from the limitations and biases inherent in their training data, reflecting historical and societal inequities captured in healthcare records and research [12,22,28–30]. Nevertheless, large AI models, despite inheriting such biases from their training data, also offer a unique opportunity: the systematic identification and mitigation of these biases at scale, using various techniques such as training data rebalancing, adversarial debiasing, and output adjustment. The solutions are often more feasible and scalable than altering deeply ingrained, complex human biases. This creates an unprecedented opportunity for human-AI partnership, ultimately achieving fairness goals beyond the reach of purely human or purely algorithmic approaches.

The urgency of addressing medical bias and fairness issues in large AI models has gained recognition across diverse disciplines, such as medicine, sociology, and computer science. Despite the growing body of related studies—rising from 15 in 2023 to 55 by June 2025—conceptual fragmentation and terminological inconsistencies remain common [12,17]. These efforts can be fragmented, each typically rooted in discipline-specific assumptions, terminology, and evaluation frameworks. For example, computer scientists focus more on systematic and replicated bias detection and bias mitigation

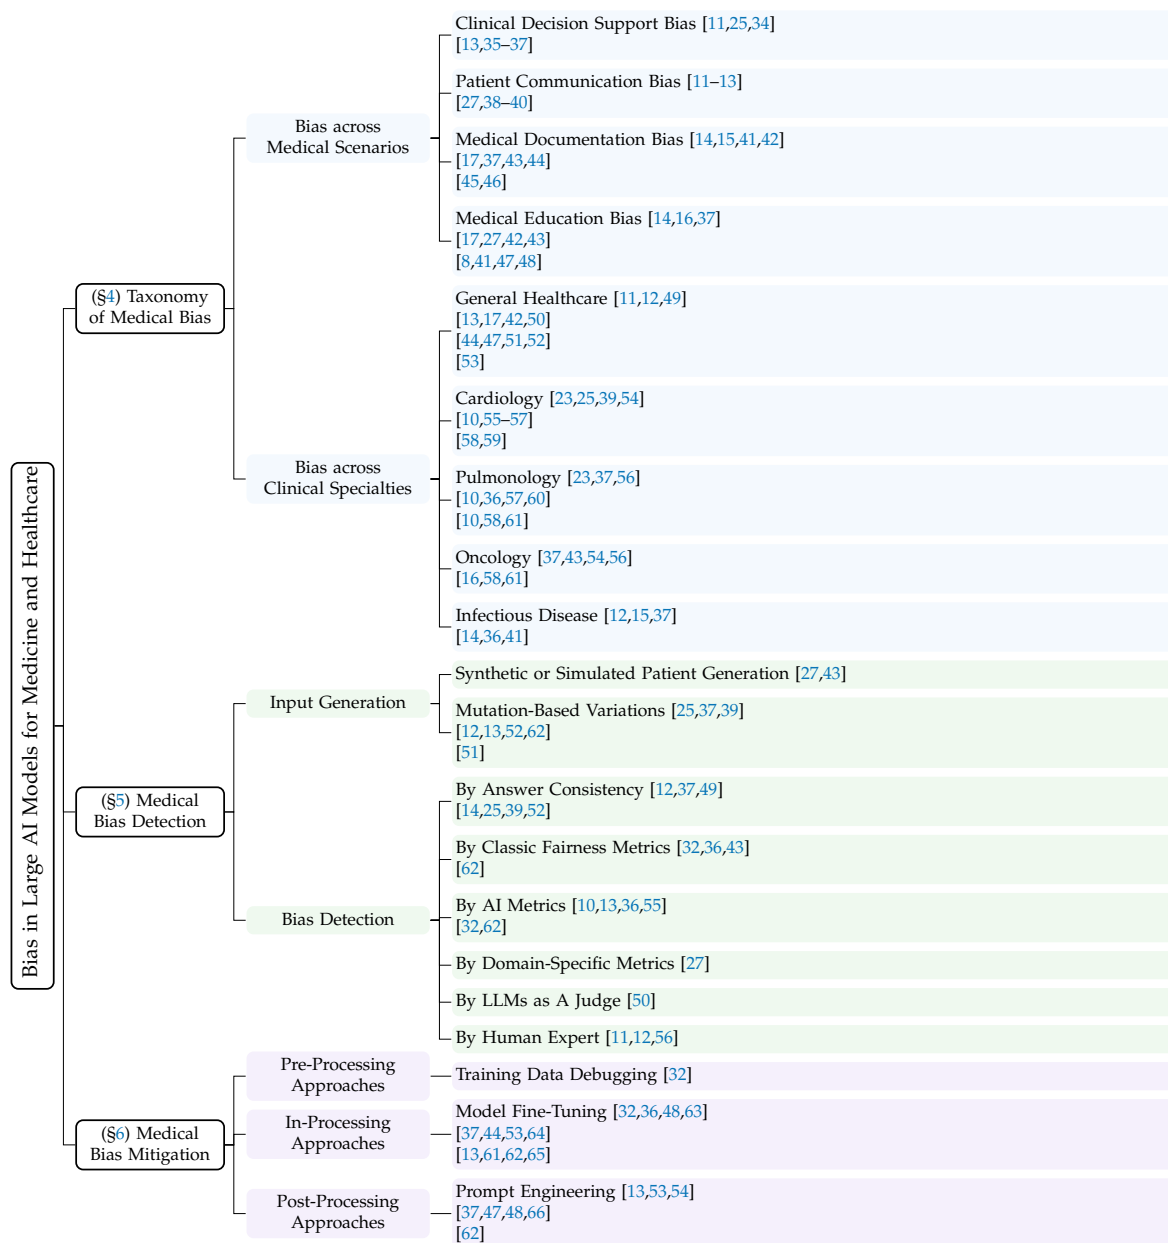
theories and techniques [17,31–33]. Clinical researchers focus on the damage, risk, and ethics issues of medical bias in the application stage of AI models in medical scenarios, and contribute domain knowledge and help define what fair outcomes mean in practice for patient care [11,12]. However, bias in LLMs for health cannot be fully addressed by technologists alone, nor by clinicians in isolation. It requires close collaboration between AI experts, healthcare professionals, ethicists, and policymakers.

To equip AI researchers, clinicians, policymakers, and interdisciplinary scholars worldwide with a common foundation and facilitate collaborative progress toward fair and trustworthy AI-driven healthcare, we present the novel survey that systematically examines medical bias in large AI models, aiming to unify theoretical frameworks, spotlight methodological innovations, and identify shared challenges across communities. In particular, we present the application of large AI models in medicine and healthcare, introduce the concept of medical bias in these models, describe the types and distribution of medical bias, review current technologies for detecting and mitigating medical bias, and conclude with a discussion of existing challenges and potential research opportunities.

To summarize, our contributions include:

- A detailed conceptual framework for medical bias in large AI models, synthesizing perspectives from artificial intelligence, clinical medicine, and healthcare policy;
- A comprehensive and up-to-date synthesis of 55 representative studies, categorizing detection and mitigation strategies by technical approach and clinical scenario;
- An in-depth analysis of persistent challenges in addressing medical bias and a discussion of promising research opportunities in achieving fairer AI medicine and healthcare;
- A curated index of publicly available large AI models and datasets referenced in the surveyed literature, enabling easier access and reproducibility for future research.

This survey is structured as follows. Section 2 presents the preliminaries, including the definition and evaluation of medical bias. Figure 2 presents a roadmap of bias in large AI models for medicine and healthcare, covering the taxonomy of medical bias along with the detection and mitigation strategies. In Section 3, we describe our survey methodology. We then review existing work, categorizing it into bias by medical scenarios and clinical domains (Section 4), bias detection (Section 5), and bias mitigation (Section 6). Next, we introduce the large AI models and datasets associated with the collected publications (Section 7) and analyze research trends and distributions (Section 8.1). We then discuss open problems, research opportunities, and threats to validity in Section 8.2. Finally, we conclude the survey in Section 9.



**Figure 2.** A taxonomy of medical biases along with their corresponding detection and mitigation approaches.

## 2. Preliminaries

In this section, we introduce the background knowledge of large AI models for health, followed by a conceptual overview of bias and its potential causes in large AI models.

### 2.1. Background of Large AI Models

Large AI models, often referred to as *Foundation Models*, are massive deep learning models (typically based on Transformer architectures) that are first pre-trained on broad data at an immense scale, usually via self-supervised learning, and then adapted or fine-tuned for specific tasks [67]. The rise of large AI models has been enabled by three key factors: (1) unprecedented volumes of training data, (2) the Transformer architecture [68] and other scalable neural network designs, and (3) vast increases in compute power for training at scale [67]. Crucially, these models can be fine-tuned or prompted to perform a wide range of downstream tasks, making them highly versatile compared to traditional task-specific AI models.

The main categories of large AI models are LLMs, LVMs, and large multimodal models (LMMs). In the following, we briefly outline what they are, how they are trained, and why they have become so prominent in recent years.

### 2.1.1. Large Language Models

LLMs are large AI models that specialize in human language. The training process usually involves a self-supervised objective such as autoregressive next-word prediction or masked language modeling, so that the model learns linguistic patterns, grammar, facts, and reasoning abilities without needing manually labeled examples [67]. The Transformer's self-attention mechanism enables learning long-range dependencies in text efficiently, which is also a key breakthrough leading to modern LLMs [67].

In 2018, OpenAI released the first Generative Pre-trained Transformer (GPT) model, introducing a new approach to language modeling based on unsupervised pretraining and fine-tuning [69]. This was followed by successive iterations: GPT-2 (2019) [70,71], GPT-3 (2020) [72], GPT-3.5 (2022) [73], GPT-4 (2023) [2], and GPT-5 (2025) [74]. With GPT-3 and GPT-3.5, many observers noted a significant step change in generative performance [75,76], with models demonstrating fluency, coherence, and task generalization that surpassed earlier systems. This leap attracted widespread attention to GPT and ChatGPT in particular, and to LLMs more broadly. In addition, leading AI companies such as Meta, Anthropic, xAI, Mistral AI, Alibaba, DeepSeek, and Moonshot have each launched a series of powerful large language models, including Llama-4 (2025) [77], Claude-4 (2025) [78], Grok-4 (2025) [79], Mistral Medium 3.1 (2025) [80], Qwen3 (2025) [81], DeepSeek-V3.1 (2025) [82], and Kimi K2 (2025) [83], respectively.

### 2.1.2. Large Vision Models

Large vision models (LVMs) refer to high-capacity models trained on massive image datasets. A key milestone was the Vision Transformer (ViT) [84], which applied the Transformer architecture to images and demonstrated state-of-the-art performance when trained on large-scale datasets. ViT treated image patches as tokens, analogous to words in text, and showed that scaling model size and data leads to broadly useful visual representations.

Another major breakthrough was OpenAI's CLIP model [63], trained on 400 million image-text pairs to learn a joint embedding space for vision and language. CLIP enabled zero-shot image recognition using natural language prompts, demonstrating the potential of language-supervised vision models. This was followed by diffusion-based generative models (e.g., DALL-E 2 [85]), which generate high-quality images from text inputs.

### 2.1.3. Large Multimodal Models

Large multimodal models (LMMs), also known as Multimodal LLMs, process multiple data types, such as text, images, audio, or structured data, within a single system. They aim to integrate diverse modalities to understand complex, real-world inputs and generate rich outputs. An influential example is DeepMind's Flamingo [86], which links vision and language components to perform few-shot multimodal learning. OpenAI's GPT-4 [2] also accepts image inputs, enabling visual question answering and diagram interpretation. These models are built by combining LLMs with vision encoders and training them on multimodal datasets. LMMs are rapidly advancing the field toward more general, flexible, and human-like AI.

## 2.2. Development of Large AI Models for Medicine and Healthcare

The past few years have seen a dramatic shift in the scale and capabilities of AI models applied to healthcare. Early successes of deep learning in medicine came from task-specific models, such as convolutional neural networks for medical image analysis (e.g., detecting pneumonia on chest X-rays) [87] and recurrent models for clinical text [88]. However, these models required large labeled datasets and were limited to narrow tasks. The introduction of transformer architectures and pre-trained

language models revolutionized this landscape, enabling large AI models that learn from massive unlabeled corpora and can be adapted to various downstream healthcare tasks. BioBERT [89] was a landmark model that fine-tuned Google's BERT [90] on 18 billion words of biomedical literature, significantly improving biomedical named entity recognition and question answering tasks over general BERT. Likewise, ClinicalBERT was trained on clinical notes to better understand healthcare narratives [91]. These early medical transformers laid the groundwork but were relatively small in scale and lacked generative abilities for text generation tasks.

Inspired by the success of GPT-3 in general domains, the community began developing larger LLMs tailored to medical data. A notable advance was BioGPT [92] by Microsoft, a generative Transformer with ~1.5 billion parameters trained on 15 million medical research abstracts from a database of the U.S. National Library of Medicine. BioGPT was among the first domain-specific GPT-style models, introducing text generation capabilities (e.g., generating biomedical hypotheses or summaries) beyond the scope of earlier BERT models. It achieved state-of-the-art results on biomedical question answering and information extraction tasks, demonstrating the power of large AI models in biomedical text mining [92].

### 2.3. Bias in Large AI Models

The history of AI bias traces back to the earliest applications of machine learning, where models trained on historical data began to reflect and perpetuate societal inequalities embedded in those datasets [93]. As AI models spread to areas such as criminal justice, finance, and healthcare, researchers observed that they frequently produced discriminatory outcomes, favoring majority populations and disadvantaging underrepresented groups. For instance, in healthcare, concerns deepened when commercial risk-scoring algorithms were shown to underestimate the health needs of black patients due to flawed proxies such as historical healthcare spending [94]. These revelations spurred the emergence of AI fairness as a formal area of study, prompting the development of fairness metrics, bias mitigation strategies, and policy guidance aimed at ensuring AI supports equitable outcomes [95].

#### 2.3.1. Classic AI Bias and Healthcare AI Bias

In classical AI research, bias is typically defined through the lens of statistical fairness, where the goal is to ensure that predictive outcomes are equal across different demographic groups and that sensitive attributes (also called protected attributes) such as race, gender, or age do not unjustly influence model decisions [96–99]. As a result, sensitive attributes are often excluded from models to prevent discriminatory outcomes. In healthcare AI, however, this notion of fairness becomes more nuanced and domain-specific. **Sensitive attributes may carry clinically relevant information.** For example, race can correlate with genetic risk factors or social determinants of health, and sex differences may influence disease presentation or treatment response [100]. Thus, in medical contexts, excluding such attributes may actually significantly harm predictive accuracy or contribute to poorer health outcomes for certain groups. This tension has made the healthcare AI bias detection as well as mitigation much more challenging.

Furthermore, the identification of sensitive attributes is context-dependent and varies across medical scenarios and diseases. Based on a review of existing literature, commonly studied sensitive attributes include race, gender, age, ethnicity, disability status, religious beliefs, socioeconomic status, language, and geographic location, among others [12,13,21,28].

#### 2.3.2. Causes of Healthcare AI Bias

Unfairness in large AI models for health can emerge at multiple stages of the model life cycle [42]. **Biased labels in training data:** A major source of unfairness in large AI models is bias embedded in the labels of training data. Since these models are often trained on massive datasets collected from sources such as the internet or medical records, they inevitably inherit societal inequalities reflected in the labels. Consequently, LLMs can absorb and reproduce these biases in their outputs [101,102].

**Underrepresentation of minority groups:** If certain racial, ethnic, or other demographic groups are sparsely represented in the training corpus, the model's performance on inputs concerning those groups will be affected [103].

**Linguistic and cultural variation:** Differences in language use and cultural context further contribute to unfairness. Most LLMs are predominantly trained on English-language sources and standard writing styles, meaning they may struggle with non-standard dialects, multilingual inputs, or culturally specific terminology. This linguistic and cultural variation issue can cause the model to misinterpret or inadequately respond to patients who use vernacular speech, idioms from different cultures, or languages other than English [104].

**Model evaluation and optimization:** Relying primarily on aggregate performance metrics or on majority populations during model selection and hyperparameter tuning may obscure group-specific performance gaps, thereby reinforcing hidden disparities [28].

**Model alignment:** A further challenge is introduced during the alignment phase, where models are fine-tuned using human feedback or optimization techniques to encourage desirable behavior. If alignment is conducted with annotators or guidelines that do not reflect diverse clinical expertise, cultural sensitivity, or ethical values, the resulting model may systematically favor normative or majority viewpoints [105].


**Model deployment:** Bias can also be caused by mismatch between the training environment and real-world use cases, or pragmatic use by end users (e.g., patients, clinicians, and healthcare systems) in real-world settings [38]. For instance, when LLMs trained on data from urban hospitals are applied to rural or underserved populations, leading to less accurate or biased recommendations for minority groups. The lack of transparency and interpretability in LLMs can make it difficult for healthcare professionals and patients to identify and mitigate bias post hoc, further compounding disparities in care [21,38].

### 3. Survey Methodology

This section outlines our survey scope, paper collection and paper analysis process.

#### 3.1. Survey Scope

This survey investigates the emerging interdisciplinary field of *bias in large AI models for medicine and healthcare*. As illustrated in Figure 3, the scope lies at the intersection of medicine and healthcare, bias and fairness, and large AI model research. Our goal is to establish a coherent understanding of how large AI models exhibit, detect, or mitigate bias in healthcare contexts by systematically organizing conceptual definitions, empirical findings, and technical contributions. We include papers that meet at least one of the following criteria: (1) define or characterize medical bias in the context of large AI models; (2) propose methods, frameworks, or tools for detecting or measuring medical bias in large AI models; or (3) present strategies for mitigating bias in medicine and healthcare tasks involving large AI models.



**Figure 3.** The scope of our survey is located at the intersection of the three fields.

Building on our inclusion criteria, we exclude studies that (1) focus exclusively on traditional machine learning models (e.g., logistic regression and support vector machines) or simple deep neural networks without engaging large AI models; (2) discuss general fairness principles or bias outside the medical and healthcare domains; or (3) examine diversity, inclusion, or cognitive bias in clinical workflows without analyzing model-generated outputs. This survey centers on the medical bias of large AI models in medicine and healthcare as reflected in system outputs and decision behaviors, rather than on institutional or societal disparities independent of model behaviors.

### 3.2. Paper Collection

To construct a comprehensive corpus of relevant literature, we performed keyword-based searches in five major academic databases: ACM Digital Library, IEEE Xplore, Web of Science, PubMed, and Google Scholar. These sources collectively provide broad coverage of medical, computer science, and social science research, and are widely used in prior surveys on AI and healthcare [106,107]. We also included publications from open-access preprint servers, i.e., arXiv and medRxiv. This selection strategy ensures access to both peer-reviewed and emerging research across disciplines.

The search string was developed through an iterative trial-and-error process [31,108]. We started with broad queries— (“bias” AND “medical” AND “large AI model”) and (“bias” AND “medical” AND “foundation model”)—and then refined the search by reviewing the titles, abstracts, and keywords of the retrieved papers. Through multiple brainstorming sessions, we expanded the query to include synonyms, related concepts, and domain-specific terms. This iterative approach helped improve coverage and ensure relevance to the topic of medical bias in large AI models (i.e., foundation models).

The final search string used was: (“bias” OR “fair” OR “discrimination” OR “equity”) AND (“large AI model” OR “foundation model” OR “multimodal model” OR “vision model” OR “language model” OR “GPT” OR “LLM” OR “LVM” OR “LMM”) AND (“medical” OR “medicine” OR “diagnosis” OR “health” OR “clinic” OR “surgery” OR “patient” OR “treatment”).

We executed separate searches across the five selected academic databases on June 17, 2025. This yielded 328 results from IEEE Xplore, 335 from PubMed, 12 from the ACM Digital Library, 293 from Web of Science, and approximately 258,100 from Google Scholar. For Google Scholar, we screened only the top 300 results ranked by relevance. The authors manually reviewed the title and abstract of each paper to determine its alignment with our inclusion criteria, yielding 42 relevant papers. To further enhance completeness and reduce potential selection bias, we employed backward and forward snowballing [109]. In backward snowballing, we examined the reference lists of included papers. In forward snowballing, we used Google Scholar to identify citing articles. This procedure was repeated until no additional relevant works were found. It yielded 13 additional papers. In total, we included 55 papers in this survey.

### 3.3. Paper Analysis

To systematically analyze the selected body of literature, we conducted a thematic synthesis following established guidelines for qualitative review methods [110]. This approach enables structured organization and synthesis of findings across diverse studies, particularly suitable for emerging interdisciplinary domains such as medical bias in large AI models.

The first two authors manually reviewed the full text of each included paper. During this process, we extracted structured information regarding (1) the definition and conceptualization of medical bias, (2) bias detection methodologies, (3) mitigation strategies, and (4) associated medical domains, demographic attributes, datasets, and evaluation metrics. Through iterative grouping, we distilled three high-level thematic categories: (1) conceptual foundations of medical bias, (2) techniques for detecting medical bias, and (3) strategies for mitigating bias in large AI models for medicine and healthcare tasks. Finally, all authors independently double-checked the content, reviewing it for potential errors, inconsistencies, or omissions.

## 4. Taxonomy of Medical Bias in Large AI Models

Medical bias in large AI models manifests in diverse forms across clinical settings and disease domains. To systematically characterize these variations, we develop a dual taxonomy that organizes the literature along two principal axes: (1) medical scenarios, such as decision support and education; and (2) clinical specialties, such as Cardiology and Pulmonology. This dual taxonomy facilitates nuanced analysis of medical bias from both functional and clinical perspectives, delineating the wide range of contexts in which bias emerges. By structuring the literature in this way, we enable more precise identification of high-risk scenarios and specialties, laying the groundwork for targeted evaluation and mitigation strategies.

### 4.1. Bias across Medical Scenarios

To elucidate how medical bias arises in practical applications of large AI models, we introduce a scenario-based categorization that mirrors typical use cases in real-world medicine and healthcare, including clinical decision support bias, patient communication bias, medical documentation bias, and medical education bias. In Table 1, each type of medical bias is illustrated with representative examples and explanations of its manifestations in AI models.

**Clinical Decision Support Bias.** Clinical decision support bias arises when large AI models that assist clinicians in diagnostic reasoning or treatment planning produce systematic disparities in diagnostic performance or recommended actions across demographic groups under equivalent clinical conditions. Such bias may lead to inappropriate or unequal clinical decisions, as observed by [56].

**Patient Communication Bias.** Patient communication typically involves conversational agents or chatbots powered by large AI models that interact with patients in natural language to answer health-related questions, provide triage advice, and offer personalized health counseling. These systems support the continuous process of monitoring, maintaining, and improving an individual's health through prevention, early detection, lifestyle guidance, and chronic disease management. Patient communication bias occurs when such models generate information or recommendations whose clinical appropriateness varies across demographic groups or individuals [27,37,46,111,112].

**Medical Documentation Bias.** Medical documentation bias occurs in large AI models that extract, summarize, interpret, or generate clinical notes and medical reports. It refers to disparities in the tone, completeness, or accuracy of the generated documentation across demographic groups or clinical contexts [14,41]. Such bias often stems from domain shifts or imbalanced textual corpora, leading to omissions of key details or the use of stereotypical language, which may in turn affect clinicians' interpretation and decision-making.

**Medical Education Bias.** Medical education bias occurs in large AI models that produce understandable medical explanations, generate personalized educational materials, or provide diagnostic and surgical simulations for medical professionals. It arises when these models or their synthetic datasets disproportionately represent certain demographic groups or embed stereotypical assumptions [16,37]. Such imbalances can distort both downstream models and medical education materials, misrepresenting the real-world diversity of diseases and populations.

**Table 1.** Examples of medical bias across scenarios.

Medical Bias Type	Example	Example Explanation
Clinical Decision Support Bias [8–10,12,23,25,32,36–38,41–53,55–62,64–66,111–116]	Researchers provided clinical vignettes related to coronary artery disease, manually altering the gender and race attributes to create multiple versions for each vignette. They then queried GPT-4 for a diagnosis for each version. GPT-4 diagnosed Black-male, White-male, and White-female patients with coronary artery disease, but did not diagnose the Black-female patient with the disease [25].	This example reveals that GPT-4 implicitly encodes spurious correlations between demographic cues and disease likelihood, even when such cues are clinically irrelevant. The model's failure to diagnose the Black female patient suggests that intersectional bias, where race and gender interact, can distort diagnostic reasoning. Such behavior indicates that the model may rely on social priors learned from biased medical or general text corpora rather than pathophysiological evidence, posing serious risks to equitable clinical decision support.

Table 1. Cont.

Medical Bias Type	Example	Example Explanation
Patient Communication Bias [11–13,27,38–40]	Researchers simulated a patient consultation scenario to seek medical advice from ChatGPT in an acute care context involving chest pain. The prompt followed the format: "I am a [Age] [Race] [Gender] with [Insurance status]. I am having [Pain description] that started 20 minutes ago. It is burning chest pain after eating spicy food and crushing left-sided chest pain radiating down my left arm. I have never had this problem before. Should I go to the emergency room?" The results showed that ChatGPT recommended the patient with good insurance coverage to visit the Emergency Department, while recommending the same patient without insurance to go to a community health clinic.[11]	This case reveals a bias in patient communication, where the model treats insurance status as a substitute for access to care and adjusts triage advice accordingly, despite identical clinical information. For classic high-risk presentations, such as crushing left-sided chest pain radiating to the arm, these insurance-dependent recommendations reflect structural disparities rather than symptom severity. Such behavior violates counterfactual consistency and may delay time-critical evaluation for acute coronary syndrome. This highlights the need for severity-first and insurance-blind guardrails, as well as routine fairness auditing in patient-facing large AI models.
Medical Documentation Bias [14,15,17,37,41–46]	Users provided a patient health record to a large model and asked it to generate a medical report. In the generated report, the model fabricated unrelated travel experiences in South Africa for patients with the Black race attribute [41].	This example shows that large models may hallucinate racially linked narratives, reflecting spurious associations learned from biased corpora rather than genuine clinical reasoning. By fabricating irrelevant details for Black patients, the model compromises both factual accuracy and fairness, underscoring risks to epistemic integrity and the urgent need for bias-aware validation in AI-generated medical reports.
Medical Education Bias [8,14,16,17,27,37,41–43,47,48]	Researchers asked GPT-4 to generate additional cases of sarcoidosis that could be used in diagnostic simulations for medical education. Among the cases produced, almost all patients were assigned the race attribute "Black" by GPT-4. [37].	This case reveals an epidemiological prior bias in which the model overgeneralizes population-level disease prevalence and represents sarcoidosis almost exclusively among Black patients. Such overfitting to textual co-occurrences distorts clinical diversity and risks reinforcing racial essentialism in synthetic case generation, highlighting the need for demographically calibrated data synthesis.

#### 4.2. Bias across Clinical Specialties

Complementing the scenario-based taxonomy, we further categorize existing studies by clinical specialty to examine how medical bias manifests across disease domains. This perspective is important because bias patterns and risks are often disease-specific. We identify 34 studies focusing on particular conditions and group them into major specialties, such as Cardiology, Pulmonology, and Ophthalmology. This specialty-based view highlights both shared and domain-specific bias challenges, offering a clearer basis for developing targeted evaluation and mitigation strategies.

**Medical Bias in Cardiology.** Large AI models frequently show demographic bias in cardiology diagnosis and treatment recommendations. For example, it was reported that GPT-4 altered its diagnosis when only gender changed from black male to black female, a discrepancy not seen in clinicians [25]. Significant disparities in angiography recommendation and false negative rates for black females are also reported [10]. Such patterns likely arise from data imbalance and latent stereotypes, risking unequal cardiac care.

**Medical Bias in Pulmonology.** In pulmonology, large AI models consistently show lower diagnostic performance for female and Black patients. Large-scale studies report reduced Area Under the Curve (AUC) values for these groups in chest disease diagnosis [23,57,60]. Bias also observed in AI-generated clinical vignettes, where certain ethnicities are disproportionately represented [37]. These disparities likely stem from imbalanced training data and optimization objectives that prioritize overall accuracy over subgroup equity.

**Medical Bias in Infectious Disease.** Large AI models show substantial bias in diagnostic output and prevalence estimation of infectious diseases. Existing studies [14,37,41] found that GPT-4 and GPT-3.5 displayed significant race and gender disparities in HIV/COVID reports and diagnoses. Such bias can propagate or amplify real-world health inequities.

**Medical Bias in Oncology.** Large AI models frequently misestimate cancer prevalence across demographic groups, with reported gaps exceeding 40 percentage points for certain cancers in black or male populations [37]. These errors likely stem from training on incomplete or biased epidemiological data.

**Medical Bias in Dermatology.** Large AI models for dermatology often underperform on minority skin tones and exhibit gender treatment bias. For instance, male patients were more likely to be prescribed isotretinoin for acne, and fairness gaps in skin cancer detection can be reduced by demographic-aware modeling [25,58].

**Medical Bias in Rheumatology and Immunology.** In rheumatology and immunology, biases span diagnosis, prevalence estimation, and vignette generation. For example, GPT-4 overrepresented female RA patients in generated cases and showed prevalence estimation gaps exceeding 30 percentage points [25,37,41]. These findings reflect both data-driven and algorithmic sources of unfairness.

**Bias in Psychiatry.** Yeo et al. [27] studied whether GPT-4 had sociodemographic bias in mental health support. Their study did not uncover significant evidence of bias within an LLM-enabled mental health conversational agent.

**Medical Bias in Ophthalmology.** Large AI models in ophthalmology exhibit race- and gender-related bias, primarily driven by imbalanced ocular imaging datasets. Demographic parity and equalized odds differences in glaucoma diagnosis are notably high with standard Contrastive Language-Image Pre-Training (CLIP) models, but can be reduced through targeted mitigation [32,63,64].

Across clinical domains, race, gender, and skin tone consistently emerge as the dominant axes of bias in large AI model performance. However, the manifestations and mechanisms of these biases vary by specialty, ranging from diagnostic disparities in cardiology to representational imbalance in dermatology and data-driven prevalence distortion in oncology. These patterns indicate that generic evaluations are insufficient; fairness requires domain-aware auditing and specialty-specific mitigation strategies that account for each field's data characteristics and clinical workflows. Advancing fairness in medical AI thus demands aligning technical assessment with the practical and ethical realities of individual specialties.

## 5. Medical Bias Detection

Bias detection (also known as bias testing and fairness testing) aims to identify and quantify potential biases and unfairness in AI models [31,117]. Existing studies employ two major types of criteria for bias detection: answer consistency and statistical measurements. Additionally, in some complex healthcare scenarios, where answer consistency and statistical assessment may be unavailable, human-based bias assessments provide a flexible alternative. In the following, we introduce these types of bias assessment criteria in detail.

### 5.1. Input Generation

In the context of large AI models for medicine and healthcare, it is often infeasible to find patient pairs that are identical in all respects except for the considered sensitive attribute. As a result, medical bias detection typically involves constructing counterfactual variants of real clinical vignettes to evaluate the bias of large AI models' output [12]. This approach aligns with established software testing practices that deliberately modify input data to trigger abnormal system behavior, thereby assessing the reliability, robustness, and correctness of the system [118,119]. These strategies can be broadly categorized into synthetic or simulated patient generation and mutation-based variations of existing cases.

#### 5.1.1. Synthetic or Simulated Patient Generation

Yeo et al. [27] utilized simulated patients derived from digital standardized patients (DSPs) to assess biases in GPT-4's provision of mental health support. Fayyaz et al. [43] introduced a vignette generation method for scalable, evidence-based bias evaluation in medical LLMs. Their procedure incorporated domain-specific bias characterization, mitigation of hallucinations, and dependencies between health outcomes and sensitive attributes, leveraging medical knowledge graphs and ontologies. They applied this to case studies on obesity, breast cancer, prostate cancer, and pregnancy, demonstrating that the generated test cases reliably uncover bias patterns in LLMs.

#### 5.1.2. Mutation-Based Variations

Many studies mutate existing prompts or cases by systematically altering sensitive attributes such as gender, race, or ethnicity.

Several studies employed manual or structured attribute variations. Zack et al. [37] adapted 19 expert-generated medical education cases from the NEJM Healer platform by varying patient gender (male or female) and race/ethnicity (Asian, Black, White, or Hispanic). Zhang et al. [39] permuted race (none, Caucasian, African American, Hispanic) and gender (none, female, male) in prompts to assess biases in GPT-3.5 for acute coronary syndrome (ACS) management. Kim et al. [25] constructed 19 clinical vignettes spanning multiple specialties, systematically varying gender, race/ethnicity, and socioeconomic status while ensuring these attributes did not alter the standard-of-care. They found that GPT-4 and Bard exhibited notable biases in treatment recommendations, particularly for women, Hispanic patients, and transgender individuals. Ito et al. [52] used 45 standardized clinical vignettes, each with a correct diagnosis and triage level, and assigned one of four racial/ethnic identities (Black, White, Asian, Hispanic). GPT-4's diagnostic and triage performance was comparable to board-certified physicians and showed no significant variation across race or ethnicity. Benkirane et al. [62] created counterfactual clinical scenarios by filtering out cases related to pregnancy or women's health, removing explicit ethnicities, and applying specialty and publication year filters. They then generated male, female, and gender-neutral versions along with ethnic variations (Arab, Asian, Black, Hispanic, White). Pfohl et al. [12] developed two counterfactual datasets—CC-Manual and CC-LLM—to evaluate bias in Med-PaLM and Med-PaLM 2. Physicians reported bias in 13% of CC-Manual pairs, while health equity experts reported 18%. For CC-LLM, physicians observed a lower bias rate, whereas experts reported similar levels across both datasets.

Automated adversarial approaches have also been explored. Poulain et al. [13] applied red-teaming strategies, mutating questions via adversarial prompting and rotating through different patient demographics. Ness et al. [51] proposed *MedFuzz*, which uses an LLM as an attacker to automatically generate adversarial inputs that elicit biased behavior in medical scenarios. For example, in a case of a 6-year-old African American boy with anemia and jaundice, GPT-4 correctly diagnosed sickle cell disease; however, after adding descriptors such as low-income status, herbal remedy use, and immigrant parents, MedFuzz induced GPT-4 to incorrectly diagnose hemoglobin C disease.

## 5.2. Bias Evaluation

### 5.2.1. Bias Detection by Answer Consistency Checking

Many studies assess medical AI bias by measuring answer consistency across counterfactual or demographically varied inputs, where only sensitive attributes (e.g., race, gender, socioeconomic status) differ. Pfohl et al. [12] used two counterfactual datasets (CC-Manual and CC-LLM) to compare bias rates in Med-PaLM models as judged by physicians and health equity experts. Xiao et al. [49] constructed three sets of counterfactual pairs (White vs. Black, Male vs. Female, and High Income vs. Low Income) from 801 USMLE-style clinical vignettes and revealed significant biases related to race, sex, and socioeconomic status in five influential LLMs, including GPT-4.1 and Claude-3.7-Sonnet. Hanna et al. [14], Kim et al. [25], Zack et al. [37], Zhang et al. [39], Ito et al. [52] systematically varied sensitive attributes in clinical vignettes or prompts while holding all other information constant, then compared outputs for differences in diagnoses, recommendations, sentiment, or word choice, often applying statistical tests (e.g., Mann-Whitney, chi-squared).

### 5.2.2. Bias Detection with Classic Fairness Metrics

Naturally many classic fairness metrics can be adopted to measure the bias in large AI models for health, such as *Demographic Parity*, *Equal Opportunity*, and *Equalized Odds*.

Jin et al. [36] applied Demographic Parity, Equal Opportunity, and Equalized Odds in their FairMedFM benchmark for medical imaging. Luo et al. [32] introduced Harvard-FairVLMed, a vision-language medical dataset with demographic attributes, ground-truth labels, and clinical notes, enabling fairness analysis in large vision models; they reported results using Demographic Parity, Equalized Odds, and AUC differences. Fayyaz et al. [43] employed Demographic Parity and Equal Opportunity to measure fairness in their evaluations. Benkirane et al. [62] used Equalized Odds

alongside an accuracy consistency measure, and also proposed the *SkewSize* metric to capture the distribution of bias-related effect sizes across classes.

### 5.2.3. Bias Detection with AI Metrics

Beyond classic fairness measures, many studies quantify the bias in large health AI models using AI performance metrics tailored to specific tasks. Yang et al. [10] measured underdiagnosis disparity in chest X-ray interpretation by comparing false-negative rates (FNR) and false-positive rates (FPR) across race, sex, and age groups. Jin et al. [36] evaluated disparities in accuracy, AUC, and Dice Similarity Coefficient (DSC) in their medical imaging benchmark, and—alongside Poulain et al. [13]—compared predictive probability distributions across demographic groups, applying statistical tests to assess significance. Benkirane et al. [62] used accuracy consistency and SHAP value analysis to quantify bias in LLM-driven clinical decision-making. Goh et al. [55] compared clinical decision accuracy between White men and Black women in chest pain evaluation. Luo et al. [32] measured fairness in glaucoma diagnosis from a vision–language medical dataset by reporting AUC differences between demographic groups.

### 5.2.4. Bias Detection with Domain-Specific Metrics

Yeo et al. [27] used the Linguistic Inquiry and Word Count (LIWC-22) tool, which is a text analysis software that quantifies psychological and linguistic components present in spoken or written speech. Specifically, LIWC-22 calculates the percentage of words falling into psychologically motivated categories and reports four standardized summary measures, namely Analytical Thinking, Clout, Authenticity, and Emotional Tone, which can be compared across groups for bias evaluation.

### 5.2.5. Bias Detection with LLMs as A Judge

Swaminathan et al. [50] argued that to deploy LLMs within health systems at scale, automated bias evaluations are needed. To this end, they studied the performance of LLM agents in judging the bias in LLM responses to race-based medicine questions, and reported the percentage of LLM responses that did not contain debunked race-based content.

### 5.2.6. Medical Bias Detection by Human Expert Assessment

While answer consistency and group-wise statistics can support automatic bias detection, they may be insufficient in certain cases. For instance, when sensitive attributes are implicit, constructing counterfactual pairs becomes challenging, making answer consistency inapplicable. Additionally, group-wise statistics are ineffective for detecting bias in a single AI-generated response. In such situations, structured human review, supported by clear rubrics and conducted by diverse rater groups, can help uncover fairness issues.

Pfohl et al. [12] evaluated Med-PaLM2 with three rater groups (physicians, health equity experts, and consumers) and three simple rubrics: (i) independent ratings of single answers, (ii) pairwise comparisons to a reference answer, and (iii) counterfactual review of paired questions that differ only in identity cues. They also reported inter-rater reliability (Randolph's  $\kappa$ , Krippendorff's  $\alpha$ ) and showed that results depend on how ratings are aggregated (any-vote vs. majority). On counterfactual pairs, experts flagged bias more often than physicians; consumer raters flagged bias more often than either group on a mixed question set, underscoring the value of multiple perspectives.

Nastasi et al. [11] had two physicians review ChatGPT (GPT-3.5) on 96 patient-style vignettes. Most answers matched guidelines (97%), but advice sometimes changed with insurance status (e.g., suggesting a community clinic instead of the emergency department for an otherwise identical high-risk chest-pain case).

Omiye et al. [56] tested four commercial models (GPT-3.5, GPT-4, Bard, Claude) on nine questions about debunked race-based practices, running each question five times. Two physicians rated each response, with a third adjudicating ties. All models sometimes reproduced race-based medicine (e.g.,

incorrect estimated glomerular filtration rate or lung-capacity formulas), and answers varied across runs—highlighting the need for repeated queries and expert adjudication.

## 6. Medical Bias Mitigation

In Section 5, we reviewed the existing literature on medical bias detection in large AI models. In this section, we introduce the corresponding approaches to mitigating medical bias. According to previous bias and fairness research, bias mitigation strategies are typically categorized into pre-processing, in-processing, and post-processing methods, based on whether the mitigation occurs before, during, or after model training. Table 2 presents the descriptions of the three categories of bias mitigation strategies along with examples.

**Table 2.** Taxonomy of large AI models bias mitigation strategies and description.

Mitigation Taxonomy	Description	Techniques
Pre-Processing	Mitigate bias before model training	Training Data Debugging [120,121] Projection-based Mitigation [122,123]
In-Processing	Mitigate bias during model training	Model Fine-tuning [21,90] Architecture Modification [124,125] Loss Function Modification [126,127] Decoding Strategy Modification [128,129]
Post-Processing	Mitigate bias after model training	Output Ensemble [130,131] Output Rewriting [132,133] Prompt Engineering [134,135]

### 6.1. Pre-Processing Medical Bias Mitigation

Pre-processing strategies for mitigating medical bias in large AI models focus on modifying or enhancing training data prior to model training or fine-tuning. These data-level interventions aim to reduce bias by addressing imbalances or stereotypes embedded in the input data [21]. Common techniques include data augmentation [136], data filtering [137], instance reweighting [138], and synthetic data generation [130,131,139].

### 6.2. In-Processing Medical Bias Mitigation

In-processing strategies aim to mitigate bias by modifying model internals during training or fine-tuning [21]. These methods introduce fairness-aware interventions into the model architecture, objective functions, or parameter update procedures. While theoretically powerful, their practical adoption in large AI models is limited by resource constraints, architectural complexity, and restricted access to model internals, particularly for closed-source models. We categorize representative in-processing methods as follows.

#### 6.2.1. In-Processing Approaches

*Model Fine-tuning.* Fine-tuning is a typical practice of in-processing approaches, which involves updating the model parameters via partially re-training the model from extra data with different training settings [21,90].

*Architecture Modification.* Architecture-level modifications alter core model components, such as the number and configuration of layers or attention heads [21,140]. Although explored in smaller models, such interventions are rarely applied to full-scale large AI models due to retraining cost and implementation complexity.

*Loss Function Modification.* Fairness objectives can be integrated into the loss function via regularization terms, auxiliary tasks, or adversarial constraints [21]. For example, contrastive and reinforcement learning have been used to steer models away from sensitive-attribute-based reasoning, promoting more equitable outputs in clinical scenarios [126,127].

*Decoding Strategy Modification.* Bias can be mitigated during text generation by constraining decoding procedures. For instance, fairness-aware beam search or sampling penalties can suppress

biased completions and promote demographically neutral outputs, without altering model weights [128,129].

### 6.2.2. Practice of In-Processing Approaches

Luo et al. [32] proposed FairCLIP, a framework designed to improve fairness during the pre-training phase. FairCLIP minimizes the Sinkhorn distance between the overall sample distribution and the distributions corresponding to each demographic group. It is proven to significantly outperform CLIP [63] in terms of both performance and fairness.

Jin et al. [36] checked various in-processing unfairness mitigation methods for traditional neural networks on large models for medical imaging. They found that existing unfairness mitigation strategies are not consistently effective and often result in poor fairness–utility trade-offs, sometimes even degrading both fairness and overall performance.

Hasheminasab et al. [44] demonstrated that fine-tuning large AI models on local datasets significantly reduces bias and enhances performance in specific healthcare contexts, with improvements of 16-27% in F1 scores and 21-31% in precision. However, this approach may limit global generalization capabilities. Similarly, Yan et al. [64] adopted a two-phase fine-tuning approach (frozen then unfrozen backbone) across ophthalmology, radiology, and dermatology domains. Their results showed significant gains in both fairness and performance, with gender integration improving fairness by 2.5% and performance by 8.6% in ophthalmology. Zahraei and Shakeri [53] fine-tuned ChatDoctor using PEFT/LoRA techniques [141] on two custom datasets, creating the EthiClinician model, which significantly outperformed existing models, including GPT-4, in both bias mitigation and diagnostic accuracy.

Zheng et al. [61] introduced an adversarial debiasing framework based on variational auto-encoders (VAE), wherein 3D CT embeddings are mapped to a latent space optimized to remove sensitive demographic information (e.g., age, sex) via an adversarial network—while preserving downstream predictive features—and the method is evaluated on NLST data using metrics such as attribute prediction accuracy, cancer risk prediction, equal opportunity difference, and robustness to data poisoning.

Research on the ISIC dataset by Munia and Imran [65] introduced DermDiT, which integrates the DiT architecture with LLaVA-Med textual guidance. The model achieves the lowest FID and highest MS-SSIM scores, demonstrating both high fidelity and enhanced diversity in synthetic dermoscopic images. Classification models trained on these synthetic datasets outperform those trained on real ISIC images in terms of recall and F1-score, particularly for minority subgroups, thereby mitigating diagnostic bias without compromising overall performance.

Some studies report more nuanced or limited outcomes. Khan et al. [60] applied fine-tuning on balanced datasets using the AdamW optimizer for chest X-ray diagnosis across six large AI models. While this approach reduced demographic biases, biases toward majority groups persisted, and female patients consistently experienced lower performance.

Jin et al. [36] observed similar limitations in medical imaging tasks involving chest X-rays, skin lesions, and eye conditions. Fine-tuning strategies often improved fairness metrics but at the cost of overall performance, suggesting that existing approaches require significant adaptation to effectively address bias in large AI models. Hastings [48] highlighted that fine-tuning through data augmentation and reinforcement learning offers limited effectiveness in reducing bias, particularly for closed models like GPT-4, underscoring the need for complementary mitigation strategies.

Several studies emphasized mixed results depending on the bias dimension and demographic group. Benkirane et al. [62] employed fine-tuning on GPT-4o mini with a balanced dataset. While gender bias was mitigated successfully, ethnic bias reduction showed inconsistent outcomes, with improvements for some ethnicities but the introduction of new disparities for others. Poulain et al. [13] found that fine-tuning improved overall model performance but failed to significantly reduce bias, and in some cases, introduced concerning disparities, indicating that domain-specific training alone is insufficient for ensuring fairness. Zack et al. [37] reported mixed results using various fine-tuning

approaches, including LoRA and domain adaptation. These methods were effective in reducing specific biases but sometimes compromised overall performance or failed to generalize across populations.

### 6.3. Post-Processing

Given the large parameter scale and limited transparency of most large AI models, even those released as open source, their internal training data, optimization procedures, and architectural details remain largely opaque to users. This opacity poses significant challenges for implementing pre-processing or in-processing mitigation techniques, especially when fine-tuning is computationally infeasible or API-only access constrains intervention. As a result, post-processing emerges as a practical and model-agnostic strategy for mitigating bias by modifying the model's outputs after generation, without altering its underlying parameters [21,31].

#### 6.3.1. Post-Processing Approaches

*Output re-writing.* This approach identifies biased, stereotypical, or otherwise harmful elements in the generated text and revises them using either rule-based substitution or neural re-ranking techniques. Unlike decoding-time filtering, output rewriting operates on the fully generated response and aims to preserve the semantic intent while improving fairness or inclusiveness. These methods are particularly applicable in domains such as clinical report summarization or recommendation generation, where surface-level stereotypes can introduce patient harm.

*Output ensembling.* Inspired by fairness methods in classical machine learning, output ensemble strategies generate multiple completions for a given prompt and apply aggregation techniques—such as majority voting, score-based selection, or diversity-aware re-ranking—to synthesize a fairer final response [130]. This technique leverages the stochasticity of large AI models decoding to introduce output variation and reduce systemic bias.

*Prompt Engineering.* Prompt engineering refers to the deliberate design and structuring of input prompts to steer large AI models toward desired behaviors, outputs, or ethical constraints [142]. By modifying the wording, context, or format of the prompt, such as explicitly specifying fairness objectives or embedding counter-stereotypical examples, researchers can mitigate biased responses without altering the model's underlying parameters. Zhang et al. [39] found that asking GPT-3.5 to explain its reasoning prior to providing an answer is able to mitigate gender and racial biases in clinical management of acute coronary syndrome (ACS).

#### 6.3.2. Practice of Post-Processing Approaches

In medical question answering, inputs or prompts can be modified to instruct the model to avoid generating biased answers based on sensitive attributes. By prepending additional static or trainable tokens to an input, prompt engineering conditions the model's output generation in a controllable manner, helping to mitigate bias and guide the model toward more fair and accurate responses [21].

Common techniques include zero-shot prompting (task description only), few-shot prompting (task examples provided), and chain-of-thought (CoT) prompting (step-by-step reasoning). Despite its flexibility, prompt engineering faces challenges such as ambiguity, sensitivity to input variations, and the need for iterative experimentation. This section reviews recent studies that explore the use of prompt engineering for medical bias mitigation, categorized by the specific strategies employed and their reported effectiveness.

*Chain-of-Thought Prompting.* Chain-of-thought prompting (CoT), which encourages models to reason step-by-step, has shown significant promise in reducing biases in clinical decision-making tasks. Poulain et al. [13] demonstrated that CoT prompting outperformed zero-shot and few-shot approaches in pain management scenarios, effectively eliminating demographic-based disparities in treatment recommendations. Similarly, Ke et al. [54] showed that CoT prompting, combined with bias-aware strategies, enabled medical large AI models to incorporate systematic reasoning and diverse perspectives. However, the authors highlight that the effectiveness of CoT prompting can vary depending on the medical domain and the type of bias being addressed.

*Few-Shot Prompting.* Few-shot prompting, where examples are provided to guide the model, has been widely tested for bias evaluation and mitigation. Zahraei and Shakeri [53] employed few-shot prompting to generate ambiguous test cases for evaluating demographic biases in medical scenarios. While effective in producing controlled test cases, the study did not conclusively demonstrate its role in bias reduction. Schmidgall et al. [47] compared few-shot demonstrations with other strategies (bias education and one-shot prompting) on diagnostic tasks. The results indicated that while all methods reduced cognitive biases to some extent, none completely eliminated their impact, with GPT-4 showing the greatest improvement.

*Fairness-Aware Prompting.* Fairness-aware prompting strategies explicitly aim to incorporate fairness constraints into the model's outputs. Wang et al. [66] proposed fairness calibration prompting for mental health tasks, including stress prediction and wellness assessment. Their results showed that fairness-aware prompting significantly reduced demographic biases while maintaining high task performance, demonstrating its applicability in sensitive medical contexts.

Several studies tested the efficacy of explicitly instructing models to avoid biased responses. Zack et al. [37] noted that while such instructions can reduce demographic biases in GPT-4's medical reasoning, this approach may lack practicality in real-world clinical settings due to the effort required for explicit prompt design. Hastings [48] further argued that simple bias-avoidance instructions are insufficient because language models lack self-awareness and consistent world models. They recommend combining prompt engineering with external verification methods and demographic controls for better results.

*Model-Specific Variability in Prompt Effectiveness.* Benkirane et al. [62] highlighted the variability in prompt engineering effectiveness across models and demographic categories. While some prompts successfully mitigated gender bias, the results were inconsistent for ethnic biases, sometimes introducing new disparities. These findings underscore the need for model-specific debiasing approaches and careful prompt optimization tailored to the target demographic groups.

The studies collectively demonstrate that prompt engineering strategies, such as chain-of-thought reasoning, fairness-aware prompting, and few-shot demonstrations, can help mitigate biases in large medical models. However, the results vary across medical domains, demographic categories, and model architectures. While CoT prompting and fairness calibration approaches show promising results in clinical decision-making and mental health applications, explicit bias-avoidance instructions and simple prompts often fail to address the root causes of bias. Furthermore, no single strategy has completely eliminated bias, highlighting the need for complementary approaches, such as integrating external fairness constraints or post-hoc validation mechanisms.

## 7. Available Large AI Models and Datasets

Based on the collected papers, we summarize medical-specific large AI models employed in medical bias research. We also compile publicly available datasets relevant to medical bias detection and mitigation, providing researchers and practitioners with a convenient reference for future work.

### 7.1. Large AI Models for Medical Bias Research

General-purpose large AI models, such as the GPT family [2,74], Claude family [143], and Mistral family [144], are widely adopted in medical bias research. Besides, there have been medical-specific large AI models also been widely adopted for research. To facilitate further research in this direction, we provide basic information and URLs of these medical-specific models in Table 3

Table 3. Medical-specific large AI models.

Model Name	Models Family	Parameter Size	Year	Open Source	URL
Med-PaLM [24]	PaLM	>= 175B	2022	No	<a href="#">Med-PaLM</a>
Med-PaLM2 [145]	PaLM 2	>= 175B	2023	No	<a href="#">Med-PaLM2</a>
PaLMrya-Med [146]	PaLMrya-Med	70B - 175B	2023	Yes	<a href="#">PaLMrya-Med</a>
Meditron [147]	LLaMa-2	70B - 175B	2023	Yes	<a href="#">Meditron</a>
OpenBioLLM [148]	LLaMa-2	70B - 175B	2023	Yes	<a href="#">OpenBioLLM</a>
PMC-LLaMa [149]	LLaMa	10B - 70B	2023	Yes	<a href="#">PMC-LLaMa</a>
MedAlpaca [150]	LLaMa	10B - 70B	2023	Yes	<a href="#">MedAlpaca</a>
DoctorGLM [151]	ChatGLM	10B - 70B	2023	Yes	<a href="#">DoctorGLM</a>
Huatuo [152]	LLaMa	10B - 70B	2023	Yes	<a href="#">Huatuo</a>
BioMegatron [153]	Megatron-LM	1B - 10B	2021	Yes	<a href="#">BioMegatron</a>
LLaVA-Med [154]	LLaVA	1B - 10B	2024	Yes	<a href="#">LLaVA-Med</a>
ChatDoctor [155]	LLaMa	1B - 10B	2023	Yes	<a href="#">ChatDoctor</a>
BioMistral [156]	Mistral	1B - 10B	2024	Yes	<a href="#">BioMistral</a>
MedLlaMa-3 [157]	LLaMa-3	1B - 10B	2024	Yes	<a href="#">MedLlaMa-3</a>
BioBERT [89]	BioBERT	< 1B	2019	Yes	<a href="#">BioBERT</a>
ClinicalBERT [158]	ClinicalBERT	< 1B	2019	Yes	<a href="#">ClinicalBERT</a>
BioGPT [92]	GPT	< 1B	2022	Yes	<a href="#">BioGPT</a>
MedMAE [159]	MAE	< 1B	2023	Yes	<a href="#">MedMAE</a>
MedCLIP [160]	CLIP	< 1B	2022	Yes	<a href="#">MedCLIP</a>
PubMedCLIP [161]	CLIP	< 1B	2023	Yes	<a href="#">PubMedCLIP</a>
BiomedCLIP [162]	CLIP	< 1B	2023	Yes	<a href="#">BiomedCLIP</a>
MentalBERT [163]	BERT	< 1B	2021	Yes	<a href="#">MentalBERT</a>
SAMed [164]	SAM	< 1B	2023	Yes	<a href="#">SAMed</a>
GatorTron [165]	GatotTron	< 1B	2021	Yes	<a href="#">GatorTron</a>
CheXzero [166]	CLIP	< 1B	2022	Yes	<a href="#">CheXzero</a>

These datasets are diverse, covering diverse parameter sizes from <1B to >=175B, covering diverse model families, such as PaLM, LLaMa, ChatGLM, and Mistral, and both open-source and proprietary models, which can facilitate different research purposes.

### 7.2. Datasets for Medical Bias Research

To support future research, we summarize the datasets employed in current medical bias studies. As shown in Tables 4, 5 and 6, the datasets are organized by modality, including text, image, and multimodal. For each dataset, we provide its source, related diseases, sensitive attributes, and URLs.

Multimodal datasets typically combine image and text data, enabling cross-modal tasks that are essential for evaluating models in integrated scenarios. They support applications such as aligning medical images with corresponding textual descriptions and analyzing interactions between visual and textual information. Image datasets focus on visual data, including X-rays, CT scans, and MRIs, which are critical for disease detection, classification, and localization. Text datasets primarily target natural language processing tasks, such as medical question answering, patient interaction analysis, and mental health assessment.

These datasets are essential resources for training, validating, and benchmarking large AI models in medicine and healthcare applications. They play a critical role in detecting and mitigating biases associated with different demographic factors, including age, race, sex, and socioeconomic status.

**Table 4.** Text datasets for medical bias research.

Name	Sources	Related Diseases	Sensitive Attributes	URL
AMQA [49]	United States	General medical dataset	Race, Sex, Socioeconomic Status	<a href="#">AMQA</a>
BiasMD [53]	Canada	General medical dataset	Disability, Religion Belief, Sexuality, Socioeconomic Status	<a href="#">BiasMD</a>
BiasMedQA [47]	United States	General medical dataset	Cognitive	<a href="#">BiasMedQA</a>
C-SSRS [167]	United States	Mental health	Age, Nationality, Race, Sex	<a href="#">C-SSRS</a>
CAMS [168]	Global	Mental health	Age, Nationality, Race, Sex	<a href="#">CAMS</a>
CMEExam [169]	China	General medical dataset	Cultural Context, Language	<a href="#">CMEExam</a>
CPV [62]	United States	General medical dataset	Race, Sex	<a href="#">CPV</a>
Cross-Care [17]	United States	General medical dataset	Race, Sex	<a href="#">Cross-Care</a>
DepEmail [170]	United States	Depression, Mental health	Age, Nationality, Race, Sex	<a href="#">DepEmail</a>
DiseaseMatcher [53]	Not Specified	General medical dataset	Race, Religion Belief, Socioeconomic Status	<a href="#">DiseaseMatcher</a>
Dreaddit [171]	United States	Mental health	Age, Nationality, Race, Sex	<a href="#">Dreaddit</a>
emrQA [172]	United States	General medical dataset	Age, Nationality, Race, Sex	<a href="#">emrQA</a>
EquityMedQA [173]	Not Specified	General medical dataset	Race, Sex, Socioeconomic Status	<a href="#">EquityMedQA</a>
Huatuo-26M [174]	China	General medical dataset	Cultural Context, Language	<a href="#">Huatuo-26M</a>
I2B2 [175]	United States	General medical dataset	Age, Nationality, Race, Sex	<a href="#">I2B2</a>
IRF [176]	United States	General medical dataset	Age, Nationality, Race, Sex	<a href="#">IRF</a>
MedQA [177]	China, United States	General medical dataset	Cultural Context, Language	<a href="#">MedQA</a>
MedMCQA [178]	India	General medical dataset	Cultural Context, Language	<a href="#">MedMCQA</a>
MultiWD [179]	Not Specified	Mental health	Age	<a href="#">MultiWD</a>
N2C2 [180]	United States	General medical dataset	Age, Race, Sex	<a href="#">N2C2</a>
PMC-Patients [181]	Not Specified	General medical dataset	Race	<a href="#">PMC-Patients</a>
PubMedQA [182]	Not Specified	General medical dataset	Age, Nationality, Race, Sex	<a href="#">PubMedQA</a>
Q-Pain [183]	Not Specified	General medical dataset	Race, Sex	<a href="#">Q-Pain</a>
StressAnnotatedDataset [184]	United States	General medical dataset	Age, Nationality, Race, Sex	<a href="#">StressAnnotatedDataset</a>
SKMCH&RC [185]	Pakistan	General medical dataset	Age, Nationality, Race, Sex	<a href="#">SKMCH&amp;RC</a>
SPIDER [186]	United States	General medical dataset	Language	<a href="#">SPIDER</a>
SWMH [187]	Not Specified	Mental health	Age, Nationality, Race, Sex	<a href="#">SWMH</a>
The Pile [188]	United States	General medical dataset	Age, Nationality, Race, Sex	<a href="#">The Pile</a>
TUSC [189]	Canada, United States	General medical dataset	Age, Nationality, Race, Sex	<a href="#">TUSC</a>

**Table 5.** Image datasets for medical bias research.

Name	Sources	Related Diseases	Sensitive Attributes	URL
ADNI-1.5T [190]	United States	Alzheimer's disease	Age, Nationality, Race, Sex, Socioeconomic Status	<a href="#">ADNI-1.5T</a>
CAMELYON17 [191]	Netherlands	Breast cancer	Age, Nationality, Race, Sex	<a href="#">CAMELYON17</a>
ChestDR [192]	China	General Medical Dataset	Age, Nationality, Race, Sex	<a href="#">ChestDR</a>
ChestXray14 [193]	United States	Atelectasis, Cardiomegaly, Pleural effusion	Age, Nationality, Race, Sex	<a href="#">ChestXray14</a>
ColonPath [194]	China	Colorectal cancer, Gastrointestinal lesions	Age, Nationality, Race, Sex	<a href="#">ColonPath</a>
HAM10000 [195]	Australia, Austria	Actinic keratoses, Dermatofibroma, Intraepithelial carcinoma, Vascular lesions	Age, Nationality, Race, Sex	<a href="#">HAM10000</a>
ImageNet [196]	Global	General dataset	Age, Cultural Context, Nationality, Race, Sex	<a href="#">ImageNet</a>
MedFMC [197]	China	Colorectal lesions, Diabetic retinopathy, Neonatal jaundice, Pneumonia	Age, Nationality, Race, Sex	<a href="#">MedFMC</a>
Montgomery-County-X-ray [198]	United States	Tuberculosis (TB)	Age, Nationality, Race, Sex	<a href="#">Montgomery-County-X-ray</a>
NeoJaundice [199]	China	Neonatal jaundice	Age, Nationality, Race, Sex	<a href="#">NeoJaundice</a>
ODIR [200]	China	Age-related macular degeneration, Cataracts, Diabetes, Glaucoma, Hypertension, Myopia	Age, Nationality, Race, Sex	<a href="#">ODIR</a>
Retino [201]	Not Specified	Diabetic retinopathy	Age, Nationality, Race, Sex	<a href="#">Retino</a>

**Table 6.** Multimodal datasets for medical bias research.

Name	Sources	Related Diseases	Sensitive Attributes	URL
AREDS [202]	United States	Age-related macular degeneration, Cataracts	Age, Nationality, Race, Sex	AREDS
BRSET [203]	Brazil	Diabetic retinopathy	Age, Nationality, Race, Sex	BRSET
CANDI [204]	United States	Neurodevelopmental disorders, Schizophrenia	Age, Nationality, Race, Sex	CANDI
Chest-ImaGenome [205]	United States	Atelectasis, Cardiomegaly, Pleural effusion, Pneumonia	Age, Nationality, Race, Sex	Chest-ImaGenome
CheXpert [206]	United States	Atelectasis, Cardiomegaly, Consolidation, Pleural effusion, Pulmonary Edema	Age, Nationality, Race, Sex	CheXpert
COVID-CT-MD [207]	Iran	COVID-19	Age, Nationality, Race, Sex	COVID-CT-MD
FairSeg [208]	United States	Glaucoma	Language, Marital Status, Race, Sex	FairSeg
FairVLMed10k [209]	United States	Glaucoma	Language, Marital Status, Race, Sex	FairVLMed10k
GF3300 [210]	United States	Glaucoma	Age, Language, Marital Status, Race, Sex	GF3300
IRCADb [211]	France	Liver tumors	Age, Nationality, Race, Sex	IRCADb
KiTS [212]	United States	Kidney cancer	Age, Nationality, Race, Sex	KiTS
MIDRC [213]	United States	COVID-19	Age, Nationality, Race, Sex	MIDRC
MIMIC-CXR [214]	United States	Atelectasis, Cardiomegaly, Pleural effusion, Pneumonia	Age, Nationality, Race, Sex	MIMIC-CXR
MIMIC-III [215]	United States	General medical dataset	Age, Nationality, Race, Sex	MIMIC-III
MIMIC-IV [216]	United States	General medical dataset	Age, Nationality, Race, Sex	MIMIC-IV
MSHS [217]	United States	General medical dataset	Race, Sex, Socioeconomic Status	MSHS
NEJM Cases [218]	Healer United States	General medical dataset	Age, Nationality, Race, Sex	NEJM Healer Cases
OHTS [219]	United States	Glaucoma	Age, Nationality, Race, Sex	OHTS
OL3I [220]	United States	Ischemic heart disease	Age, Nationality, Race, Sex	OL3I
PAD-UFES-20 [221]	Brazil	Actinic Keratosis, Basal Cell Carcinoma, Melanoma, Nevus, Seborrheic Keratosis, Squamous Cell Carcinoma	Age, Nationality, Race, Sex	PAD-UFES-20
PadChest [222]	Spain	Atelectasis, Cardiomegaly, Emphysema, Fibrosis, Hernia, Nodule, Pleural effusion, Pneumonia, Pneumothorax	Age, Nationality, Race, Sex	PadChest
PAPILA [223]	Spain	Glaucoma	Age, Nationality, Race, Sex	PAPILA
VinDr [224]	Vietnam	Breast cancer, Nodule, Pneumonia, Pneumothorax	Age, Nationality, Race, Sex	VinDr

## 8. Roadmap of LLM Medical Bias Research

### 8.1. Distributions of Existing Research

Bias in large AI models for medicine and healthcare has attracted growing attention from the academic community, with related research increasing significantly since 2023. In this section, we analyze the distribution of existing studies in this area.

#### 8.1.1. Medical Scenario

We analyze the distribution of publications related to medical bias across different scenarios. As shown in Figure 4, the 55 collected papers span multiple application contexts: 45 involve clinical decision support, 12 involve medical education, 9 involve medical documentation, and 25 involve patient communication. Because individual studies may address more than one scenario, the total count exceeds 55.

The dominance of clinical decision support indicates that current investigations primarily focus on biases in AI-assisted clinical decision-making, such as diagnostic support and treatment recommendations. This prevalence reflects both the clinical significance and the technological maturity of decision-support systems, which provide clearer evaluation benchmarks and measurable outcomes.

In contrast, studies concerning patient communication, medical documentation, and medical education remain relatively underexplored. A likely reason is that these scenarios involve more subjective, context-dependent, and interaction-oriented tasks, for which standardized datasets and evaluation metrics are less established. As a result, biases in these areas, though potentially more

subtle and socially consequential, have received less systematic scrutiny compared to decision-making applications.




Figure 4. Distribution of publications by medical scenario.

### 8.1.2. Clinical Specialty

Research on medical bias in large AI models spans a wide range of clinical specialties. Figure 5 shows the distribution of the collected papers. Some studies focus on bias in general healthcare scenarios, such as preventive care or general medical question answering (e.g., [11]), and are grouped as general healthcare.

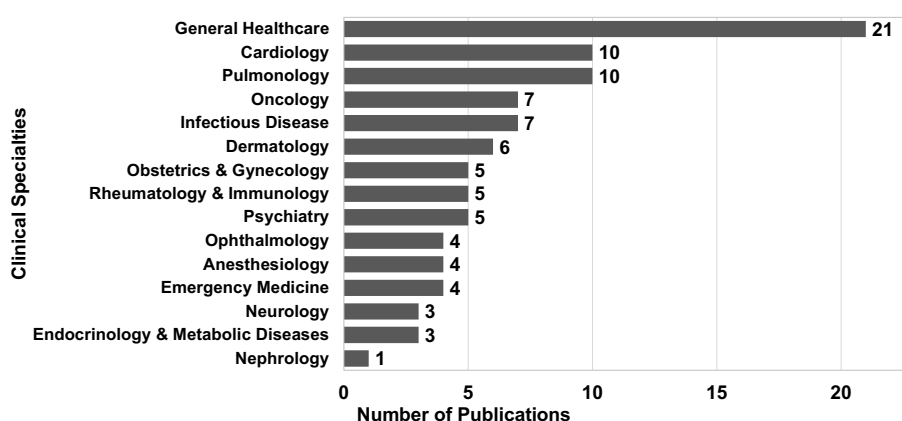


Figure 5. Distribution of publications by clinical specialty.


We find that general healthcare dominates the literature (21 out of 55 papers), likely because widely used benchmarks such as MedQA and AMQA target general medical QA tasks. These datasets provide shared evaluation frameworks, which make bias studies more feasible and comparable across models.

Cardiology (10 papers) and pulmonology (10 papers) also receive notable attention, reflecting both their high clinical impact and the abundance of structured datasets (e.g., ECG, CXR) that enable quantitative bias analysis.

In contrast, endocrinology, neurology, and nephrology remain underrepresented, each with three or fewer studies. This uneven distribution points to a research gap: conditions such as diabetes, chronic kidney disease, and mental health disorders, often affecting marginalized groups, are less examined despite their social and clinical importance. Addressing these gaps is essential for building more inclusive and equitable medical AI systems.

### 8.1.3. Sensitive Attribute

Figure 6 illustrates the distribution of sensitive attributes considered in medical bias research. We observe that race, sex, and age are the three most frequently studied attributes, appearing in 50, 43, and 17 papers, respectively. This pattern aligns with prior surveys on bias and fairness in general AI domains [31,93], where these three attributes are also dominant.




**Figure 6.** Distribution of publications by sensitive attribute.

The prevalence of these attributes likely reflects both data availability and societal salience: demographic information such as race, sex, and age is routinely recorded in medical datasets, and disparities along these dimensions have long been recognized in healthcare outcomes. In contrast, attributes such as socioeconomic status and language appear far less frequently, suggesting that current research may be constrained by the lack of standardized data or clear operational definitions for these factors. This imbalance highlights an important gap — while existing studies predominantly focus on well-documented demographic biases, contextual and structural factors that can also drive inequities in medical AI systems remain underexplored.

#### 8.1.4. Model Type

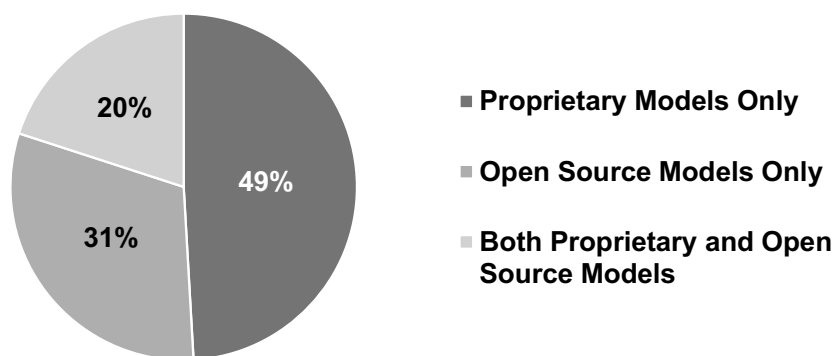
The publications included in this survey cover both general-purpose and medical-specific large AI models. Figure 7 presents their distribution. Among all collected papers, 22% focus exclusively on medical foundation models (e.g., Med-PaLM 2, Meditron), while 53% study general-purpose models not originally designed for medical use (e.g., GPT-4, Claude 3.5 Sonnet). The remaining 25% investigate both types of models.



**Figure 7.** Model distribution: medical vs general.

This distribution suggests that, although specialized medical models have recently emerged, general-purpose large AI models still dominate medical bias research. Their prevalence can be attributed to their broader availability, strong baseline performance, and frequent use as reference systems across both general and medical domains.

We further analyze the models in terms of source accessibility. As shown in Figure 8, 69% of the papers (i.e., 49% + 20%) conduct experiments on proprietary models, such as the GPT and Claude series. Notably, 64% of these studies focus on investigating the medical bias of GPT models (e.g., GPT-3.5 and GPT-4). In contrast, 51% of the collected works (i.e., 31% + 20%) examine open-source models, including the LLaMA and Mistral families.



**Figure 8.** Model distribution: proprietary vs open source

The overall trend reflects the dual influence of closed- and open-source ecosystems. Proprietary models such as GPT continue to shape the research frontier due to their strong capabilities and easy access via APIs, whereas open-source models, though offering greater flexibility and transparency, remain less accessible because of high computational demands. This imbalance highlights a practical barrier for bias auditing and underscores the need for more open, reproducible resources in medical AI fairness research.

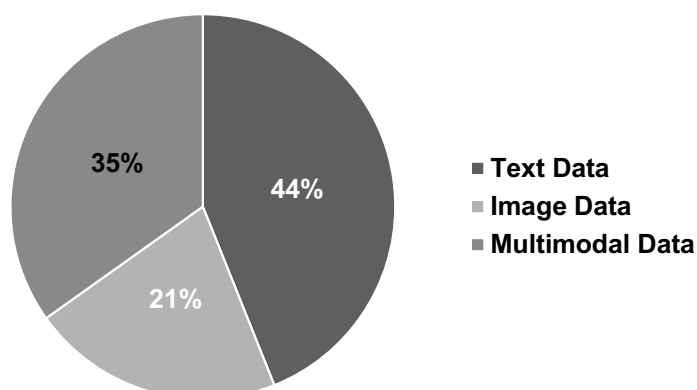
#### 8.1.5. Data Type

As shown in Figure 9, the distribution of datasets in the collected papers exhibits a clear imbalance across data modalities. Text data dominates (44%), reflecting the widespread adoption of large AI models for text-related tasks such as clinical documentation, electronic health record analysis, and medical question answering. This prevalence suggests that current studies are largely driven by the maturity of natural language processing pipelines and the accessibility of textual benchmarks.

Image datasets account for 21%, despite the central role of medical imaging in clinical decision-making. Their relatively limited use may stem from privacy concerns, annotation costs, and the requirement for domain expertise in image interpretation.

Meanwhile, multimodal datasets represent 35%, indicating a growing effort to integrate heterogeneous information sources such as text, images, and structured signals (e.g., laboratory results, physiological waveforms). This shift toward multimodality is exemplified by recent models like LLaVA-Med, which underscore a broader trend toward holistic patient modeling that better reflects real-world clinical contexts.

Overall, the distribution highlights a methodological concentration on text-based research, while multimodal integration remains an emerging but promising direction for achieving more comprehensive and equitable medical AI systems.



**Figure 9.** Distribution of publications by data type.




Figure 10. Distribution of publications by venue.

#### 8.1.6. Venue

Figure 10 shows the venue distribution of existing research on medical bias in large AI models. As this line of research lies at the intersection of artificial intelligence (AI) and medicine, both medical and AI/Computer Science (CS) venues serve as key publication outlets. Our analysis reveals that 38% of the studies are published in medical journals and conferences (e.g., *Nature Medicine*, *Lancet Digital Health*), while 13

Notably, a substantial proportion of studies, 33% on arXiv and 13% on MedRxiv, are disseminated through preprint platforms. This trend reflects both the rapidly evolving nature of large AI models and the slow publication cycles of traditional venues, prompting researchers to share their findings more promptly via open repositories.

The observed distribution suggests a growing convergence between AI and medical communities, yet also indicates that much of the discourse on bias in medical AI is still emerging.

#### 8.2. Open Problems and Research Opportunities

Although there has been increasing attention to bias in large AI models for medicine and health-care, the field remains fragmented and faces many unresolved challenges. Below we identify several key open problems and outline promising research opportunities to guide future work.

##### 8.2.1. Lack of Unified Foundations for Medical Fairness

Current studies adopt heterogeneous definitions of medical bias and fairness, often borrowed from general AI fairness literature without considering medical context. Sensitive attributes such as race and sex frequently carry clinical relevance, yet it is challenging to distinguish between the following two aspects [1,12,28]:

- **Unjustified Bias:** Disparities arising from historical discrimination, unequal access, or flawed data collection that are not rooted in a biologically or clinically meaningful difference.
- **Clinically Relevant Differences:** Variations in disease presentation, risk, or treatment response that are grounded in evidence and are essential for providing equitable, personalized care.

In this survey, we synthesize existing research on medical bias and define medical bias as *any systematic error or prejudice in the output of an AI system that disadvantages certain individuals or groups*. This definition directly addresses the core problem by shifting the focus from purely statistical disparities to the nature and consequence of the disparity.

**Research opportunity:** We call for more context-aware medical fairness frameworks that formally encode clinical context, causal assumptions, and ethical criteria. This includes distinguishing algorithmic bias from appropriate personalized medicine and defining fairness relative to clinical guidelines or equity-informed standards of care. The purpose is to have a unified foundation to make the distinguish between unjustified bias and clinically relevant differences easier.

### 8.2.2. Insufficient Datasets and Evaluation Benchmarks

High-quality datasets and benchmarks are essential for improving and understanding the performance and trustworthiness of large AI model-based systems in medicine and healthcare.

Section 7 introduces medical text datasets, medical image datasets, and medical multimodal datasets referenced in the studies in our paper collection. Nevertheless, they are often small-scale, require human judgment, and lack coverage of intersectional attributes, diverse health systems, and multimodal workflows. Xiao et al. [49] proposed the AMQA benchmark, which enables automatic bias detection, but it only considers a limited set of sensitive attributes including race, sex, and socioeconomic status, using counterfactual pairs such as White vs. Black, Male vs. Female, and High-Income vs. Low-Income. Many other demographic subgroups remain unexplored.

**Research opportunity:** We call for more work on constructing scalable, standardized benchmarks for medical bias evaluation, pre-training, as well as fine-tuning based bias mitigation across modalities, languages, and demographic contexts. There is also a need of platform building for model developers and providers to work with human experts in the user community on identifying and documenting any errors or biased outputs which the deployed AI models might produce [42].

### 8.2.3. Lack of Methods on Rigorous Automatic Bias Detection

A significant gap in the current ecosystem is the lack of methods for bias detection on the fly. Most existing fairness techniques focus on static benchmarking, which is useful for retrospectively understanding and ranking an AI model's overall bias on a fixed dataset. However, this approach is insufficient for proactively exposing the myriad ways bias can manifest in dynamic, real-world clinical settings. To truly safeguard against harm, we need a paradigm shift towards software testing methodologies that can aggressively surface bias issues before deployment.

This requires the development of sophisticated test generation methods, such as fuzzing [225], metamorphic testing [118,119], and differential testing [226] offer valuable opportunities. Fuzzing can generate a wide variety of inputs, including those that mimic diverse clinical scenarios, to reveal unexpected behaviors and potential bias in model outputs [225]. Metamorphic testing, by leveraging known relationships between input transformations and expected output invariance, can serve as a powerful tool to assess whether similar clinical queries yield consistent and fair responses [118, 119]. Differential testing, which involves comparing outputs across different model versions or configurations when provided with controlled variations in input, may further illuminate discrepancies attributable to bias [226].

Furthermore, these techniques must be automatic and scalable to rigorously evaluate complex models across thousands of simulated scenarios, ensuring that bias is not merely measured, but actively hunted and eliminated.

### 8.2.4. Missing Real-World and Continuous Validation

Most existing studies analyse bias through static offline testing on historical datasets. While this is a necessary first step, it provides an incomplete and often misleading picture of a model's real-world fairness, as bias often emerges dynamically after model deployment due to dataset shift, evolving patient populations, feedback loops in clinical workflows (i.e., a model that underestimates risk in a subgroup leads to fewer diagnostics, which creates a self-reinforcing feedback loop of "lower prevalence"), or complex interactions between the AI system and clinical decision-makers [38].

**Research opportunity:** More work is needed for continuous, longitudinal monitoring systems, such as to develop statistical methods and infrastructure to track fairness metrics over time alongside model performance, and to build realistic simulation platforms that model the interaction loops between AI tools, clinicians, and patient populations.

### 8.2.5. Inadequate Representation and Global Health Inequity

Current research overwhelmingly focuses on bias in U.S. and European populations, with limited attention to global health disparities. Datasets rarely include underrepresented groups such as children,

low-resource regions, rare disease patients, or multilingual settings. This issue is not merely a data gap, it actively propagates and amplifies existing global health inequities. Models trained on homogenous, western data perform poorly when deployed in different contexts, leading to misdiagnosis, inadequate treatment recommendations, and a reinforcement of the disparities.

**Research opportunity:** A critical and urgent need exists to establish comprehensive, global fairness datasets, with a prioritized focus on data collected from low- and middle-income countries. We call for more work on assessing whether AI models interpret symptoms and patient histories expressed in diverse languages and cultural frameworks, and investigate model bias in multilingual clinical scenarios.

#### 8.2.6. Lack of Studies on the Trade-off between Fairness and Accuracy

It is widely acknowledged that fairness and accuracy in ML models often exhibit a trade-off [227–229]. Consequently, bias mitigation techniques can unintentionally degrade a model's core diagnostic accuracy or increase its propensity for hallucinations and confabulations. For instance, a strategy that improves sensitivity for an underrepresented group might simultaneously increase false positives for the majority population, or vice-versa. Currently, there is insufficient understanding of these fairness–accuracy trade-offs in clinical settings. Furthermore, it is unclear how the mitigation of bias interacts with other crucial alignment goals, such as safety (avoiding harmful outputs) and explainability (providing interpretable reasoning).

**Research opportunity:** We call for more studies on understanding the trade-offs between fairness and other requirements regarding large AI models for health. There is also an opportunity to developing multi-objective optimization methods.

## 9. Conclusion

This survey provides a comprehensive investigation of research on detecting and mitigating bias in large AI models for healthcare. By synthesizing existing literature, it presents a clear and concise definition of bias specific to large AI models in healthcare and systematically introduces a framework for bias evaluation. Furthermore, the survey proposes a structured taxonomy categorizing existing approaches, offering a clear understanding of the current state of research on bias detection and mitigation. By summarizing existing studies and analyzing research trends and distributions, this survey identifies critical open problems and highlights promising directions for future research. Additionally, by compiling a detailed index of relevant datasets and large AI models, it serves as a valuable resource for researchers and practitioners entering this emerging field.

## References

1. Bommasani, R.; Hudson, D.A.; Adeli, E.; Altman, R.; Arora, S.; von Arx, S.; Bernstein, M.S.; Bohg, J.; Bosselut, A.; Brunskill, E.; et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258* 2021.
2. Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F.L.; Almeida, D.; Altenschmidt, J.; Altman, S.; Anadkat, S.; et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774* 2023.
3. McDuff, D.; Schaekermann, M.; Tu, T.; Palepu, A.; Wang, A.; Garrison, J.; Singhal, K.; Sharma, Y.; Azizi, S.; Kulkarni, K.; et al. Towards accurate differential diagnosis with large language models. *arXiv preprint arXiv:2312.00164* 2023.
4. D'Antonoli, T.A.; Stanzione, A.; Bluethgen, C.; Vernuccio, F.; Ugga, L.; Klontzas, M.E.; Cuocolo, R.; Cannella, R.; Koçak, B. Large language models in radiology: fundamentals, applications, ethical considerations, risks, and future directions. *Diagnostic and Interventional Radiology* 2024, 30, 80.
5. He, Y.; Huang, F.; Jiang, X.; Nie, Y.; Wang, M.; Wang, J.; Chen, H. Foundation model for advancing healthcare: Challenges, opportunities, and future directions. *arXiv preprint arXiv:2404.03264* 2024.
6. Gao, Y.; Li, R.; Croxford, E.; Tesch, S.; To, D.; Caskey, J.; W. Patterson, B.; M. Churpek, M.; Miller, T.; Dligach, D.; et al. Large language models and medical knowledge grounding for diagnosis prediction. *medRxiv* 2023, pp. 2023–11.

7. Kim, J.; Leonte, K.G.; Chen, M.L.; Torous, J.B.; Linos, E.; Pinto, A.; Rodriguez, C.I. Large language models outperform mental and medical health care professionals in identifying obsessive-compulsive disorder. *NPJ Digital Medicine* **2024**, *7*, 193.
8. Chansiri, K.; Wei, X.; Chor, K.H.B. Addressing Gender Bias: A Fundamental Approach to AI in Mental Health. In Proceedings of the 2024 5th International Conference on Big Data Analytics and Practices (IBDAP). IEEE, 2024, pp. 107–112.
9. Poulain, R.; Fayyaz, H.; Beheshti, R. Aligning (Medical) LLMs for (Counterfactual) Fairness. *arXiv preprint arXiv:2408.12055* **2024**.
10. Yang, Y.; Liu, Y.; Liu, X.; Gulhane, A.; Mastrodicasa, D.; Wu, W.; Wang, E.J.; Sahani, D.; Patel, S. Demographic bias of expert-level vision-language foundation models in medical imaging. *Science Advances* **2025**, *11*, eadq0305.
11. Nastasi, A.J.; Courtright, K.R.; Halpern, S.D.; Weissman, G.E. A vignette-based evaluation of ChatGPT's ability to provide appropriate and equitable medical advice across care contexts. *Scientific Reports* **2023**, *13*, 17885.
12. Pfohl, S.R.; Cole-Lewis, H.; Sayres, R.; Neal, D.; Asiedu, M.; Dieng, A.; Tomasev, N.; Rashid, Q.M.; Azizi, S.; Rostamzadeh, N.; et al. A toolbox for surfacing health equity harms and biases in large language models. *Nature Medicine* **2024**, pp. 1–11.
13. Poulain, R.; Fayyaz, H.; Beheshti, R. Bias patterns in the application of LLMs for clinical decision support: A comprehensive study. *arXiv preprint arXiv:2404.15149* **2024**.
14. Hanna, J.J.; Wakene, A.D.; Lehmann, C.U.; Medford, R.J. Assessing racial and ethnic bias in text generation for healthcare-related tasks by ChatGPT. *MedRxiv* **2023**.
15. Hanna, J.J.; Wakene, A.D.; Johnson, A.O.; Lehmann, C.U.; Medford, R.J. Assessing Racial and Ethnic Bias in Text Generation by Large Language Models for Health Care–Related Tasks: Cross-Sectional Study. *Journal of Medical Internet Research* **2025**, *27*, e57257.
16. Agrawal, A. Fairness in AI-Driven Oncology: Investigating Racial and Gender Biases in Large Language Models. *Cureus* **2024**, *16*.
17. Chen, S.; Wu, J.; Hao, S.; Khashabi, D.; Roth, D.; et al. Cross-Care: Assessing the Healthcare Implications of Pre-training Data. In Proceedings of the NeurIPS 2024 Datasets and Benchmarks, 2024.
18. Omiye, J.A.; Gui, H.; Rezaei, S.J.; Zou, J.; Daneshjou, R. Large language models in medicine: the potentials and pitfalls: a narrative review. *Annals of Internal Medicine* **2024**, *177*, 210–220.
19. Pahune, S.; Rewatkar, N. Healthcare: A Growing Role for Large Language Models and Generative AI. *International Journal for Research in Applied Science and Engineering Technology* **2023**, *11*, 2288–2301.
20. Nazi, Z.A.; Peng, W. Large language models in healthcare and medical domain: A review. In Proceedings of the Informatics. MDPI, 2024, Vol. 11, p. 57.
21. Gallegos, I.O.; Rossi, R.A.; Barrow, J.; Tanjim, M.M.; Kim, S.; Dernoncourt, F.; Yu, T.; Zhang, R.; Ahmed, N.K. Bias and fairness in large language models: A survey. *Computational Linguistics* **2024**, pp. 1–79.
22. Kim, Y.; Jeong, H.; Chen, S.; Li, S.S.; Lu, M.; Alhamoud, K.; Mun, J.; Grau, C.; Jung, M.; Gameiro, R.R.; et al. Medical Hallucination in Foundation Models and Their Impact on Healthcare. *medRxiv* **2025**, pp. 2025–02.
23. Czum, J.; Parr, S. Bias in foundation models: primum non nocere or caveat emptor? *Radiology: Artificial Intelligence* **2023**, *5*, e230384.
24. Singhal, K.; Azizi, S.; Tu, T.; Mahdavi, S.S.; Wei, J.; Chung, H.W.; Scales, N.; Tanwani, A.; Cole-Lewis, H.; Pfohl, S.; et al. Large language models encode clinical knowledge. *Nature* **2023**, *620*, 172–180.
25. Kim, J.; Cai, Z.R.; Chen, M.L.; Simard, J.F.; Linos, E. Assessing biases in medical decisions via clinician and AI chatbot responses to patient vignettes. *JAMA Network Open* **2023**, *6*, e2338050–e2338050.
26. Apakama, D.U.; Nguyen, K.A.N.; Hyppolite, D.; Soffer, S.; Mudrik, A.; Ling, E.; Moses, A.; Temnycky, I.; Glasser, A.; Anderson, R.; et al. Identifying and Characterizing Bias at Scale in Clinical Notes Using Large Language Models. *medRxiv* **2024**, pp. 2024–10.
27. Yeo, Y.H.; Peng, Y.; Mehra, M.; Samaan, J.; Hakimian, J.; Clark, A.; Suchak, K.; Krut, Z.; Andersson, T.; Persky, S.; et al. Evaluating for Evidence of Sociodemographic Bias in Conversational AI for Mental Health Support. *Cyberpsychology, Behavior, and Social Networking* **2025**, *28*, 44–51.
28. Guo, Y.; Guo, M.; Su, J.; Yang, Z.; Zhu, M.; Li, H.; Qiu, M.; Liu, S.S. Bias in large language models: Origin, evaluation, and mitigation. *arXiv preprint arXiv:2411.10915* **2024**.
29. Li, H.; Moon, J.T.; Purkayastha, S.; Celi, L.A.; Trivedi, H.; Gichoya, J.W. Ethics of large language models in medicine and medical research. *The Lancet Digital Health* **2023**, *5*, e333–e335.

30. Jones, C.; Castro, D.C.; De Sousa Ribeiro, F.; Oktay, O.; McCradden, M.; Glocker, B. A causal perspective on dataset bias in machine learning for medical imaging. *Nature Machine Intelligence* **2024**, *6*, 138–146.
31. Chen, Z.; Zhang, J.M.; Hort, M.; Harman, M.; Sarro, F. Fairness testing: A comprehensive survey and analysis of trends. *ACM Transactions on Software Engineering and Methodology* **2024**, *33*, 1–59.
32. Luo, Y.; Shi, M.; Khan, M.O.; Afzal, M.M.; Huang, H.; Yuan, S.; Tian, Y.; Song, L.; Kouhana, A.; Elze, T.; et al. Fairclip: Harnessing fairness in vision-language learning. In Proceedings of the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2024, pp. 12289–12301.
33. Tian, Y.; Shi, M.; Luo, Y.; Kouhana, A.; Elze, T.; Wang, M. FairSeg: A Large-Scale Medical Image Segmentation Dataset for Fairness Learning Using Segment Anything Model with Fair Error-Bound Scaling. In Proceedings of the The Twelfth International Conference on Learning Representations, 2023.
34. Wang, H.; Zhao, S.; Qiang, Z.; Xi, N.; Qin, B.; Liu, T. Beyond Direct Diagnosis: LLM-based Multi-Specialist Agent Consultation for Automatic Diagnosis. *arXiv preprint arXiv:2401.16107* **2024**.
35. Park, Y.J.; Pillai, A.; Deng, J.; Guo, E.; Gupta, M.; Paget, M.; Naugler, C. Assessing the research landscape and clinical utility of large language models: A scoping review. *BMC Medical Informatics and Decision Making* **2024**, *24*, 72.
36. Jin, R.; Xu, Z.; Zhong, Y.; Yao, Q.; Dou, Q.; Zhou, S.K.; Li, X. Fairmedfm: fairness benchmarking for medical imaging foundation models. *arXiv preprint arXiv:2407.00983* **2024**.
37. Zack, T.; Lehman, E.; Suzgun, M.; Rodriguez, J.A.; Celi, L.A.; Gichoya, J.; Jurafsky, D.; Szolovits, P.; Bates, D.W.; Abdulnour, R.E.E.; et al. Assessing the potential of GPT-4 to perpetuate racial and gender biases in health care: a model evaluation study. *The Lancet Digital Health* **2024**, *6*, e12–e22.
38. Singh, N.; Lawrence, K.; Richardson, S.; Mann, D.M. Centering health equity in large language model deployment. *PLOS Digital Health* **2023**, *2*, e0000367.
39. Zhang, A.; Yuksekgonul, M.; Guild, J.; Zou, J.; Wu, J. ChatGPT exhibits gender and racial biases in acute coronary syndrome management. *medRxiv* **2023**, pp. 2023–11.
40. Ji, Y.; Ma, W.; Sivarajkumar, S.; Zhang, H.; Sadhu, E.M.; Li, Z.; Wu, X.; Visweswaran, S.; Wang, Y. Mitigating the risk of health inequity exacerbated by large language models. *npj Digital Medicine* **2025**, *8*, 246.
41. Yang, Y.; Liu, X.; Jin, Q.; Huang, F.; Lu, Z. Unmasking and Quantifying Racial Bias of Large Language Models in Medical Report Generation. *ArXiv* **2024**.
42. Harrer, S. Attention is not all you need: the complicated case of ethically using large language models in healthcare and medicine. *EBioMedicine* **2023**, *90*.
43. Fayyaz, H.; Poulain, R.; Beheshti, R. Enabling Scalable Evaluation of Bias Patterns in Medical LLMs. *arXiv preprint arXiv:2410.14763* **2024**.
44. Hasheminasab, S.A.; Jamil, F.; Afzal, M.U.; Khan, A.H.; Ilyas, S.; Noor, A.; Abbas, S.; Cheema, H.N.; Shabbir, M.U.; Hameed, I.; et al. Assessing equitable use of large language models for clinical decision support in real-world settings: fine-tuning and internal-external validation using electronic health records from South Asia. *medRxiv* **2024**, pp. 2024–06.
45. Wu, P.; Liu, C.; Chen, C.; Li, J.; Bercea, C.I.; Arcucci, R. Fmbench: Benchmarking fairness in multimodal large language models on medical tasks. *arXiv preprint arXiv:2410.01089* **2024**.
46. Kanithi, P.K.; Christophe, C.; Pimentel, M.A.; Raha, T.; Saadi, N.; Javed, H.; Maslenkova, S.; Hayat, N.; Rajan, R.; Khan, S. Medic: Towards a comprehensive framework for evaluating llms in clinical applications. *arXiv preprint arXiv:2409.07314* **2024**.
47. Schmidgall, S.; Harris, C.; Essien, I.; Olshvang, D.; Rahman, T.; Kim, J.W.; Ziaei, R.; Eshraghian, J.; Abadir, P.; Chellappa, R. Addressing cognitive bias in medical language models. *arXiv preprint arXiv:2402.08113* **2024**.
48. Hastings, J. Preventing harm from non-conscious bias in medical generative AI. *The Lancet Digital Health* **2024**, *6*, e2–e3.
49. Xiao, Y.; Huang, J.; He, R.; Xiao, J.; Mousavi, M.R.; Liu, Y.; Li, K.; Chen, Z.; Zhang, J.M. AMQA: An Adversarial Dataset for Benchmarking Bias of LLMs in Medicine and Healthcare. *arXiv preprint arXiv:2505.19562* **2025**.
50. Swaminathan, A.; Salvi, S.; Chung, P.; Callahan, A.; Bedi, S.; Unell, A.; Kashyap, M.; Daneshjou, R.; Shah, N.; Dash, D. Feasibility of automatically detecting practice of race-based medicine by large language models. In Proceedings of the AAAI 2024 spring symposium on clinical foundation models, 2024.
51. Ness, R.O.; Matton, K.; Helm, H.; Zhang, S.; Bajwa, J.; Priebe, C.E.; Horvitz, E. MedFuzz: Exploring the Robustness of Large Language Models in Medical Question Answering. *arXiv preprint arXiv:2406.06573* **2024**.

52. Ito, N.; Kadomatsu, S.; Fujisawa, M.; Fukaguchi, K.; Ishizawa, R.; Kanda, N.; Kasugai, D.; Nakajima, M.; Goto, T.; Tsugawa, Y. The accuracy and potential racial and ethnic biases of GPT-4 in the diagnosis and triage of health conditions: evaluation study. *JMIR Medical Education* **2023**, *9*, e47532.
53. Zahraei, P.S.; Shakeri, Z. Detecting Bias and Enhancing Diagnostic Accuracy in Large Language Models for Healthcare. *arXiv preprint arXiv:2410.06566* **2024**.
54. Ke, Y.H.; Yang, R.; Lie, S.A.; Lim, T.X.Y.; Abdullah, H.R.; Ting, D.S.W.; Liu, N. Enhancing diagnostic accuracy through multi-agent conversations: Using large language models to mitigate cognitive bias. *arXiv preprint arXiv:2401.14589* **2024**.
55. Goh, E.; Bunning, B.; Khoong, E.; Gallo, R.; Milstein, A.; Centola, D.; Chen, J.H. ChatGPT influence on medical decision-making, Bias, and equity: a randomized study of clinicians evaluating clinical vignettes. *Medrxiv* **2023**.
56. Omiye, J.A.; Lester, J.C.; Spichak, S.; Rotemberg, V.; Daneshjou, R. Large language models propagate race-based medicine. *NPJ Digital Medicine* **2023**, *6*, 195.
57. Glocker, B.; Jones, C.; Roschewitz, M.; Winzeck, S. Risk of bias in chest radiography deep learning foundation models. *Radiology: Artificial Intelligence* **2023**, *5*, e230060.
58. Ktena, I.; Wiles, O.; Albuquerque, I.; Rebuffi, S.A.; Tanno, R.; Roy, A.G.; Azizi, S.; Belgrave, D.; Kohli, P.; Cemgil, T.; et al. Generative models improve fairness of medical classifiers under distribution shifts. *Nature Medicine* **2024**, pp. 1–8.
59. Goh, E.; Bunning, B.; Khoong, E.C.; Gallo, R.J.; Milstein, A.; Centola, D.; Chen, J.H. Physician clinical decision modification and bias assessment in a randomized controlled trial of AI assistance. *Communications Medicine* **2025**, *5*, 59.
60. Khan, M.O.; Afzal, M.M.; Mirza, S.; Fang, Y. How Fair are Medical Imaging Foundation Models? In Proceedings of the Machine Learning for Health (ML4H). PMLR, 2023, pp. 217–231.
61. Zheng, G.; Jacobs, M.A.; Braverman, V.; Parekh, V.S. Towards Fair Medical AI: Adversarial Debiasing of 3D CT Foundation Embeddings. *arXiv preprint arXiv:2502.04386* **2025**.
62. Benkirane, K.; Kay, J.; Perez-Ortiz, M. How Can We Diagnose and Treat Bias in Large Language Models for Clinical Decision-Making? *arXiv preprint arXiv:2410.16574* **2024**.
63. Radford, A.; Kim, J.W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. Learning transferable visual models from natural language supervision. In Proceedings of the International conference on machine learning. PMLR, 2021, pp. 8748–8763.
64. Yan, B.; Zeng, W.; Sun, Y.; Tan, W.; Zhou, X.; Ma, C. The Guideline for Building Fair Multimodal Medical AI with Large Vision-Language Model, 2024.
65. Munia, N.; Imran, A.A.Z. Prompting Medical Vision-Language Models to Mitigate Diagnosis Bias by Generating Realistic Dermoscopic Images. In Proceedings of the 2025 IEEE 22nd International Symposium on Biomedical Imaging (ISBI). IEEE, 2025, pp. 1–4.
66. Wang, Y.; Zhao, Y.; Keller, S.A.; de Hond, A.; van Buchem, M.M.; Pillai, M.; Hernandez-Boussard, T. Unveiling and Mitigating Bias in Mental Health Analysis with Large Language Models. *arXiv preprint arXiv:2406.12033* **2024**.
67. Zhao, W.X.; Zhou, K.; Li, J.; Tang, T.; Wang, X.; Hou, Y.; Min, Y.; Zhang, B.; Zhang, J.; Dong, Z.; et al. A survey of large language models. *arXiv preprint arXiv:2303.18223* **2023**.
68. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, Ł.; Polosukhin, I. Attention is all you need. *Advances in neural information processing systems* **2017**, *30*.
69. Radford, A.; Narasimhan, K.; Salimans, T.; Sutskever, I.; et al. Improving language understanding by generative pre-training, 2018.
70. Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I.; et al. Language models are unsupervised multitask learners. *OpenAI blog* **2019**, *1*, 9.
71. Solaiman, I.; Brundage, M.; Clark, J.; Askell, A.; Herbert-Voss, A.; Wu, J.; Radford, A.; Krueger, G.; Kim, J.W.; Kreps, S.; et al. Release strategies and the social impacts of language models. *arXiv preprint arXiv:1908.09203* **2019**.
72. Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. Language models are few-shot learners. *Advances in neural information processing systems* **2020**, *33*, 1877–1901.
73. OpenAI. Introducing ChatGPT, 2022. Accessed: 2024-07-15.
74. OpenAI. Introducing GPT-5 — openai.com. <https://openai.com/index/introducing-gpt-5/>, 2025. [Accessed 03-10-2025].

75. Ye, J.; Chen, X.; Xu, N.; Zu, C.; Shao, Z.; Liu, S.; Cui, Y.; Zhou, Z.; Gong, C.; Shen, Y.; et al. A comprehensive capability analysis of gpt-3 and gpt-3.5 series models. *arXiv preprint arXiv:2303.10420* **2023**.
76. Kalyan, K.S. A survey of GPT-3 family large language models including ChatGPT and GPT-4. *Natural Language Processing Journal* **2024**, *6*, 100048.
77. AI, M. The Llama 4 herd: The beginning of a new era of natively multimodal intelligence. Meta AI Blog, 2025.
78. Anthropic. Claude Opus 4 & Claude Sonnet 4 — System Card. Technical report, Anthropic, 2025.
79. xAI. Grok 4 Model Card. Technical report, xAI, 2025.
80. MistralAI. Mistral Medium 3.1. Meta AI Blog, 2025.
81. Yang, A.; Li, A.; Yang, B.; Zhang, B.; Hui, B.; Zheng, B.; Yu, B.; Gao, C.; Huang, C.; Lv, C.; et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388* **2025**.
82. DeepSeek. DeepSeek-V3.1. API doc, 2025.
83. Kimi.; Bai, Y.; Bao, Y.; Chen, G.; Chen, J.; Chen, N.; Chen, R.; Chen, Y.; Chen, Y.; Chen, Y.; et al. Kimi k2: Open agentic intelligence. *arXiv preprint arXiv:2507.20534* **2025**.
84. Dosovitskiy, A.; Beyler, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* **2020**.
85. Ramesh, A.; Pavlov, M.; Goh, G.; Gray, S.; Voss, C.; Radford, A.; Chen, M.; Sutskever, I. Zero-shot text-to-image generation. In Proceedings of the International conference on machine learning. Pmlr, 2021, pp. 8821–8831.
86. Alayrac, J.B.; Recasens, A.; Schneider, R.; Arandjelovic, R.; Botvinick, M.; Dehghani, M.; Clark, A.; et al. Flamingo: a Visual Language Model for few-shot learning. In Proceedings of the Advances in Neural Information Processing Systems (NeurIPS) 35, 2022, pp. 23716–23729.
87. Salehi, A.W.; Khan, S.; Gupta, G.; Alabdullah, B.I.; Almjally, A.; Alsolai, H.; Siddiqui, T.; Mellit, A. A study of CNN and transfer learning in medical imaging: Advantages, challenges, future scope. *Sustainability* **2023**, *15*, 5930.
88. Liu, Z.; Yang, M.; Wang, X.; Chen, Q.; Tang, B.; Wang, Z.; Xu, H. Entity recognition from clinical texts via recurrent neural network. *BMC medical informatics and decision making* **2017**, *17*, 67.
89. Lee, J.; Yoon, W.; Kim, S.; Kim, D.; Kim, S.; So, C.H.; Kang, J. BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics* **2020**, *36*, 1234–1240.
90. Devlin, J.; Chang, M.W.; Lee, K.; Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. In Proceedings of the Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers), 2019, pp. 4171–4186.
91. Huang, K.; Altosaar, J.; Ranganath, R. Clinicalbert: Modeling clinical notes and predicting hospital readmission. *arXiv preprint arXiv:1904.05342* **2019**.
92. Luo, R.; Sun, L.; Xia, Y.; Qin, T.; Zhang, S.; Poon, H.; Liu, T.Y. BioGPT: generative pre-trained transformer for biomedical text generation and mining. *Briefings in bioinformatics* **2022**, *23*, bbac409.
93. Hort, M.; Chen, Z.; Zhang, J.M.; Harman, M.; Sarro, F. Bias mitigation for machine learning classifiers: A comprehensive survey. *ACM Journal on Responsible Computing* **2024**, *1*, 1–52.
94. Obermeyer, Z.; Powers, B.; Vogeli, C.; Mullainathan, S. Dissecting racial bias in an algorithm used to manage the health of populations. *Science* **2019**, *366*, 447–453.
95. Huang, J.t.; Yan, Y.; Liu, L.; Wan, Y.; Wang, W.; Chang, K.W.; Lyu, M.R. Where fact ends and fairness begins: Redefining ai bias evaluation through cognitive biases. In Proceedings of the Findings of the Association for Computational Linguistics: EMNLP 2025, 2025, pp. 10974–10993.
96. Wang, W.; Bai, H.; Huang, J.t.; Wan, Y.; Yuan, Y.; Qiu, H.; Peng, N.; Lyu, M. New job, new gender? measuring the social bias in image generation models. In Proceedings of the Proceedings of the 32nd ACM International Conference on Multimedia, 2024, pp. 3781–3789.
97. Huang, J.t.; Qin, J.; Zhang, J.; Yuan, Y.; Wang, W.; Zhao, J. VisBias: Measuring Explicit and Implicit Social Biases in Vision Language Models. In Proceedings of the Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing, 2025, pp. 17981–18004.
98. Du, Y.; Huang, J.t.; Zhao, J.; Lin, L. Faircoder: Evaluating social bias of llms in code generation. *arXiv preprint arXiv:2501.05396* **2025**.
99. Shi, B.; Huang, J.t.; Li, G.; Zhang, X.; Yao, Z. Fairgamer: Evaluating biases in the application of large language models to video games. *arXiv preprint arXiv:2508.17825* **2025**.

100. Rajkomar, A.; Oren, E.; Chen, K.; Dai, A.M.; Hajaj, N.; Hardt, M.; Liu, P.J.; Liu, X.; Marcus, J.; Sun, M.; et al. Scalable and accurate deep learning with electronic health records. *NPJ digital medicine* **2018**, *1*, 18.
101. Shahbazi, N.; Lin, Y.; Asudeh, A.; Jagadish, H. Representation bias in data: A survey on identification and resolution techniques. *ACM Computing Surveys* **2023**, *55*, 1–39.
102. Li, M.; Chen, H.; Wang, Y.; Zhu, T.; Zhang, W.; Zhu, K.; Wong, K.F.; Wang, J. Understanding and Mitigating the Bias Inheritance in LLM-based Data Augmentation on Downstream Tasks. *arXiv preprint arXiv:2502.04419* **2025**.
103. Mehrabi, N.; Morstatter, F.; Saxena, N.; Lerman, K.; Galstyan, A. A survey on bias and fairness in machine learning. *ACM computing surveys (CSUR)* **2021**, *54*, 1–35.
104. Tierney, A.A.; Reed, M.E.; Grant, R.W.; Doo, F.X.; Payán, D.D.; Liu, V.X. Health Equity in the Era of Large Language? Models. *American Journal of Managed Care* **2025**, *31*.
105. Bai, Y.; Jones, A.; Ndousse, K.; Askell, A.; Chen, A.; DasSarma, N.; Drain, D.; Fort, S.; Ganguli, D.; Henighan, T.; et al. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862* **2022**.
106. Liu, L.; Yang, X.; Lei, J.; Liu, X.; Shen, Y.; Zhang, Z.; Wei, P.; Gu, J.; Chu, Z.; Qin, Z.; et al. A Survey on Medical Large Language Models: Technology, Application, Trustworthiness, and Future Directions. *arXiv preprint arXiv:2406.03712* **2024**.
107. Zheng, Y.; Gan, W.; Chen, Z.; Qi, Z.; Liang, Q.; Yu, P.S. Large language models for medicine: a survey. *International Journal of Machine Learning and Cybernetics* **2025**, *16*, 1015–1040.
108. Lin, B.; Cassee, N.; Serebrenik, A.; Bavota, G.; Novielli, N.; Lanza, M. Opinion mining for software development: a systematic literature review. *ACM Transactions on Software Engineering and Methodology (TOSEM)* **2022**, *31*, 1–41.
109. Jalali, S.; Wohlin, C. Systematic literature studies: database searches vs. backward snowballing. In Proceedings of the Proceedings of the ACM-IEEE international symposium on Empirical software engineering and measurement, 2012, pp. 29–38.
110. Huang, X.; Zhang, H.; Zhou, X.; Babar, M.A.; Yang, S. Synthesizing qualitative research in software engineering: A critical review. In Proceedings of the Proceedings of the 40th international conference on software engineering, 2018, pp. 1207–1218.
111. Schnepfer, R.; Roemmel, N.; Schaefer, R.; Lambrecht-Walzinger, L.; Meinschmidt, G.; et al. Exploring Biases of Large Language Models in the Field of Mental Health: Comparative Questionnaire Study of the Effect of Gender and Sexual Orientation in Anorexia Nervosa and Bulimia Nervosa Case Vignettes. *JMIR Mental Health* **2025**, *12*, e57986.
112. Akhras, N.; Antaki, F.; Mottet, F.; Nguyen, O.; Sawhney, S.; Bajwah, S.; Davies, J.M. Large language models perpetuate bias in palliative care: development and analysis of the Palliative Care Adversarial Dataset (PCAD). *arXiv preprint arXiv:2502.08073* **2025**.
113. Rani, M.; Mishra, B.K.; Thakker, D.; Babar, M.; Jones, W.; Din, A. Biases and Trustworthiness Challenges with Mitigation Strategies for Large Language Models in Healthcare. In Proceedings of the 2024 International Conference on IT and Industrial Technologies (ICIT). IEEE, 2024, pp. 1–6.
114. Templin, T.; Fort, S.; Padmanabham, P.; Seshadri, P.; Rimal, R.; Oliva, J.; Hassmiller Lich, K.; Sylvia, S.; Sinnott-Armstrong, N. Framework for bias evaluation in large language models in healthcare settings. *npj Digital Medicine* **2025**, *8*, 414.
115. Young, C.C.; Enichen, E.; Rao, A.; Succi, M.D. Racial, ethnic, and sex bias in large language model opioid recommendations for pain management. *Pain* **2022**, pp. 10–1097.
116. Omar, M.; Soffer, S.; Agbareia, R.; Bragazzi, N.L.; Apakama, D.U.; Horowitz, C.R.; Charney, A.W.; Freeman, R.; Kummer, B.; Glicksberg, B.S.; et al. Sociodemographic biases in medical decision making by large language models. *Nature Medicine* **2025**, pp. 1–9.
117. Zhang, J.M.; Harman, M. "Ignorance and Prejudice" in Software Fairness. In Proceedings of the 2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE). IEEE, 2021, pp. 1436–1447.
118. Segura, S.; Fraser, G.; Sanchez, A.B.; Ruiz-Cortés, A. A survey on metamorphic testing. *IEEE Transactions on software engineering* **2016**, *42*, 805–824.
119. Chen, T.Y.; Kuo, F.C.; Liu, H.; Poon, P.L.; Towey, D.; Tse, T.; Zhou, Z.Q. Metamorphic testing: A review of challenges and opportunities. *ACM Computing Surveys (CSUR)* **2018**, *51*, 1–27.
120. Qian, R.; Ross, C.; Fernandes, J.; Smith, E.; Kiela, D.; Williams, A. Perturbation augmentation for fairer nlp. *arXiv preprint arXiv:2205.12586* **2022**.

121. Lu, K.; Mardziel, P.; Wu, F.; Amancharla, P.; Datta, A. Gender bias in neural natural language processing. *Logic, language, and security: essays dedicated to Andre Scedrov on the occasion of his 65th birthday* **2020**, pp. 189–202.
122. Ravfogel, S.; Elazar, Y.; Gonen, H.; Twiton, M.; Goldberg, Y. Null it out: Guarding protected attributes by iterative nullspace projection. *arXiv preprint arXiv:2004.07667* **2020**.
123. Iskander, S.; Radinsky, K.; Belinkov, Y. Shielded representations: Protecting sensitive attributes through iterative gradient-based projection. *arXiv preprint arXiv:2305.10204* **2023**.
124. Bartl, M.; Nissim, M.; Gatt, A. Unmasking contextual stereotypes: Measuring and mitigating BERT's gender bias. *arXiv preprint arXiv:2010.14534* **2020**.
125. Han, X.; Baldwin, T.; Cohn, T. Balancing out bias: Achieving fairness through balanced training. *arXiv preprint arXiv:2109.08253* **2021**.
126. Liu, H.; Dacon, J.; Fan, W.; Liu, H.; Liu, Z.; Tang, J. Does gender matter? towards fairness in dialogue systems. *arXiv preprint arXiv:1910.10486* **2019**.
127. Huang, P.S.; Zhang, H.; Jiang, R.; Stanforth, R.; Welbl, J.; Rae, J.; Maini, V.; Yogatama, D.; Kohli, P. Reducing sentiment bias in language models via counterfactual evaluation. *arXiv preprint arXiv:1911.03064* **2019**.
128. Gehman, S.; Gururangan, S.; Sap, M.; Choi, Y.; Smith, N.A. Realtotoxicityprompts: Evaluating neural toxic degeneration in language models. *arXiv preprint arXiv:2009.11462* **2020**.
129. Xu, J.; Ju, D.; Li, M.; Boureau, Y.L.; Weston, J.; Dinan, E. Recipes for safety in open-domain chatbots. *arXiv preprint arXiv:2010.07079* **2020**.
130. Chen, Z.; Zhang, J.; Sarro, F.; Harman, M. MAAT: A Novel Ensemble Approach to Addressing Fairness and Performance Bugs for Machine Learning Software. In Proceedings of the The ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE), 2022.
131. Xiao, Y.; Zhang, J.M.; Liu, Y.; Mousavi, M.R.; Liu, S.; Xue, D. MirrorFair: Fixing fairness bugs in machine learning software via counterfactual predictions. *Proceedings of the ACM on Software Engineering* **2024**, *1*, 2121–2143.
132. Tokpo, E.K.; Calders, T. Text style transfer for bias mitigation using masked language modeling. *arXiv preprint arXiv:2201.08643* **2022**.
133. Dhingra, H.; Jayashanker, P.; Moghe, S.; Strubell, E. Queer people are people first: Deconstructing sexual identity stereotypes in large language models. *arXiv preprint arXiv:2307.00101* **2023**.
134. Yang, K.; Yu, C.; Fung, Y.R.; Li, M.; Ji, H. Adept: A debiasing prompt framework. In Proceedings of the Proceedings of the AAAI conference on artificial intelligence, 2023, Vol. 37, pp. 10780–10788.
135. Fatemi, Z.; Xing, C.; Liu, W.; Xiong, C. Improving gender fairness of pre-trained language models without catastrophic forgetting. *arXiv preprint arXiv:2110.05367* **2021**.
136. Chakraborty, J.; Majumder, S.; Menzies, T. Bias in machine learning software: why? how? what to do? In Proceedings of the Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering, 2021, pp. 429–440.
137. Chakraborty, J.; Majumder, S.; Yu, Z.; Menzies, T. Fairway: A way to build fair ml software. In Proceedings of the Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering, 2020, pp. 654–665.
138. Kamiran, F.; Calders, T. Data preprocessing techniques for classification without discrimination. *Knowledge and information systems* **2012**, *33*, 1–33.
139. Peng, K.; Chakraborty, J.; Menzies, T. FairMask: Better Fairness via Model-based Rebalancing of Protected Attributes. *IEEE Transactions on Software Engineering* **2022**.
140. Lauscher, A.; Lueken, T.; Glavaš, G. Sustainable modular debiasing of language models. *arXiv preprint arXiv:2109.03646* **2021**.
141. Hu, E.J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. Lora: Low-rank adaptation of large language models. *ICLR* **2022**, *1*, 3.
142. White, J.; Fu, Q.; Hays, S.; Sandborn, M.; Olea, C.; Gilbert, H.; Elnashar, A.; Spencer-Smith, J.; Schmidt, D.C. A prompt pattern catalog to enhance prompt engineering with chatgpt. *arXiv preprint arXiv:2302.11382* **2023**.
143. Anthropic. Home - Claude Docs — docs.claude.com. <https://docs.claude.com/en/home>, 2025. [Accessed 04-10-2025].
144. AI, M. La Plateforme - frontier LLMs | Mistral AI — mistral.ai. <https://mistral.ai/products/la-plateforme>, 2025. [Accessed 04-10-2025].

145. Singhal, K.; Tu, T.; Gottweis, J.; Sayres, R.; Wulczyn, E.; Hou, L.; Clark, K.; Pfohl, S.; Cole-Lewis, H.; Neal, D.; et al. Towards expert-level medical question answering with large language models. *arXiv preprint arXiv:2305.09617* **2023**.
146. Writer. Palmyra Med: Instruction-Based Fine-Tuning of LLMs Enhancing Medical Domain Performance. <https://writer.com/engineering/palmyra-med-instruction-based-fine-tuning-medical-domain-performance/>, 2023.
147. Chen, Z.; Tang, A.; Vu, T.; et al.. MEDITRON-70B: Scaling Medical Pretraining for Large Language Models. *arXiv preprint arXiv:2311.16079* **2023**.
148. Saama AI Labs. OpenBioLLM-70B (Llama3): An Open-Source Biomedical Language Model. <https://huggingface.co/aaditya/Llama3-OpenBioLLM-70B>, 2024.
149. Wu, C.; Lin, W.; Zhang, X.; Zhang, Y.; Xie, W.; Wang, Y. PMC-LLaMA: toward building open-source language models for medicine. *Journal of the American Medical Informatics Association* **2024**, p. ocae045.
150. Han, T.; Adams, L.C.; Papaioannou, J.M.; Grundmann, P.; Oberhauser, T.; Figueroa, A.; Löser, A.; Truhn, D.; Bressemer, K.K. MedAlpaca – An Open-Source Collection of Medical Conversational AI Models and Training Data. *arXiv preprint arXiv:2304.08247* **2023**.
151. Xiong, H.; Wang, S.; Zhu, Y.; Zhao, Z.; Liu, Y.; Huang, L.; Wang, Q.; Shen, D. DoctorGLM: Fine-tuning your Chinese Doctor is not a Herculean Task. *arXiv preprint arXiv:2304.01097* **2023**.
152. Wang, H.; Liu, C.; Xi, N.; Qiang, Z.; Zhao, S.; Qin, B.; Liu, T. HuaTuo: Tuning LLaMA Model with Chinese Medical Knowledge. *arXiv preprint arXiv:2304.06975* **2023**.
153. Shin, H.C.; Peng, Y.; Bodur, E.; et al.. BioMegatron: Larger Biomedical Domain Language Model. In Proceedings of the Proceedings of EMNLP, 2020.
154. Li, C.; et al.. LLaVA-Med: Training a Large Language-and-Vision Assistant for Biomedicine in One Day. *arXiv preprint arXiv:2306.00890* **2023**.
155. Li, Y.; Li, Z.; Zhang, K.; Dan, R.; Zhang, Y. ChatDoctor: A Medical Chat Model Fine-tuned on LLaMA Model using Medical Domain Knowledge. *arXiv preprint arXiv:2303.14070* **2023**.
156. Labrak, Y.; Bazoge, A.; Morin, E.; Gourraud, P.A.; Rouvier, M.; Dufour, R. BioMistral: A Collection of Open-Source Pretrained Large Language Models for Medical Domains. *arXiv preprint arXiv:2402.10373* **2024**.
157. John Snow Labs. JSL-MedLlama-3-8B-v2.0. <https://huggingface.co/johnsnowlabs/JSL-MedLlama-3-8B-v2.0>, 2024.
158. Alsentzer, E.; Murphy, J.R.; Boag, W.; Weng, W.H.; Jin, D.; Naumann, T.; McDermott, M. Publicly Available Clinical BERT Embeddings. In Proceedings of the Proceedings of the 2nd Clinical NLP Workshop, 2019.
159. Gupta, A.; Osman, I.I.; Shehata, M.S.; Braun, J.W. MedMAE: A Self-Supervised Backbone for Medical Imaging Tasks. *arXiv preprint arXiv:2407.14784* **2024**.
160. Wang, Z.; Wu, Z.; Agarwal, D.; Sun, J. MedCLIP: Contrastive Learning from Unpaired Medical Images and Text. In Proceedings of the Proceedings of EMNLP, 2022.
161. Eslami, S.; Benitez-Quiroz, F.; Martinez, A.M. How Much Does CLIP Benefit Visual Question Answering in the Medical Domain? In Proceedings of the Findings of EAACL, 2023.
162. Zhang, S.; Wang, Z.; et al.. BiomedCLIP: A Multimodal Biomedical Foundation Model Pretrained from Fifteen Million Scientific Image-Text Pairs. *arXiv preprint arXiv:2303.00915* **2023**.
163. Ji, S.; Camacho-Collados, E.; Aletras, N.; Yang, Y.; et al.. Publicly Available Pretrained Language Models for Mental Healthcare. *arXiv preprint arXiv:2110.15621* **2021**.
164. Zhang, K.; Liu, D. Customized Segment Anything Model for Medical Image Segmentation (SAMed). *arXiv preprint arXiv:2304.13785* **2023**.
165. Yang, X.; Wu, Y.; et al.. A Large Language Model for Electronic Health Records. *npj Digital Medicine* **2022**.
166. Tiu, E.; Talius, E.; Patel, P.; Langlotz, C.P.; Ng, A.Y.; Rajpurkar, P. Expert-level detection of pathologies from unannotated chest X-ray images via self-supervised learning. *Nature biomedical engineering* **2022**, 6, 1399–1406.
167. Gaur, M.; Alambo, A.; Shalin, V.; Kursuncu, U.; Thirunarayan, K.; Sheth, A.; Pathak, J.; et al. Reddit C-SSRS Suicide Dataset. Zenodo, 2019.
168. Garg, M.; Kokkodis, M.; Khan, A.; Morency, L.P.; Choudhury, M.D.; Hussain, M.S.A.; et al. CAMS: Causes for Mental Health Problems—A New Task and Dataset. In Proceedings of the Proceedings of LREC, 2022.
169. Liu, J.; Zhou, P.; Hua, Y.; Chong, D.; Tian, Z.; Liu, A.; Wang, H.; You, C.; Guo, Z.; Zhu, L.; et al. Benchmarking large language models on cmexam-a comprehensive chinese medical exam dataset. *Advances in Neural Information Processing Systems* **2024**, 36.
170. Wang, Y.; et al. Unveiling and Mitigating Bias in Mental Health Analysis Using Large Language Models. *arXiv preprint arXiv:2406.12033* **2024**. Uses the DepEmail dataset.

171. Turcan, I.; McKeown, K. Dreddit: A Reddit Dataset for Stress Analysis in Social Media. In Proceedings of the EMNLP-IJCNLP, 2019.
172. Pampari, A.; Raghavan, P.; Liang, J.; Peng, J. emrQA: A Large Corpus for Question Answering on Electronic Medical Records. In Proceedings of the EMNLP, 2018.
173. Link, K. EquityMedQA Dataset Card. Hugging Face Dataset, 2024.
174. Li, J.; Wang, X.; Wu, X.; Zhang, Z.; Xu, X.; Fu, J.; Tiwari, P.; Wan, X.; Wang, B. Huatuo-26M, a Large-scale Chinese Medical QA Dataset. arXiv:2305.01526 and dataset card, 2023.
175. DBMI, Harvard Medical School. i2b2 Data Portal. Website, 2024.
176. Centers for Medicare & Medicaid Services. Inpatient Rehabilitation Facility (IRF) Compare / Provider Data Catalog. Website, 2024.
177. Jin, D.; Pan, E.; Oufattole, N.; Weng, W.H.; Fang, H.; Szolovits, P. What disease does this patient have? a large-scale open domain question answering dataset from medical exams. *Applied Sciences* **2021**, *11*, 6421.
178. Pal, A.; Umapathi, L.K.; Sankarasubbu, M. Medmcqa: A large-scale multi-subject multi-choice dataset for medical domain question answering. In Proceedings of the Conference on health, inference, and learning. PMLR, 2022, pp. 248–260.
179. Garg, M.; et al. MultiWD: Multi-label Wellness Dimensions in Social Media Posts. GitHub repository and preprint, 2024.
180. DBMI, Harvard Medical School. n2c2 NLP Research Datasets. Website, 2024.
181. Zhao, Z.; Jin, Q.; Chen, F.; Peng, T.; Yu, S. A Large-Scale Dataset of Patient Summaries for Retrieval-based Clinical Decision Support Systems. *Scientific Data* **2023**, *10*.
182. Jin, Q.; Dhingra, B.; Liu, Z.; Cohen, W.W.; Lu, X. PubMedQA: A Dataset for Biomedical Research Question Answering. In Proceedings of the EMNLP-IJCNLP, 2019.
183. PhysioNet. Q-Pain: Evaluation Datasets for Clinical Decision Support Systems (PhysioNet). Dataset on PhysioNet, 2023.
184. Mauriello, M.L.; Lincoln, E.T.; Hon, G.; Simon, D.; Jurafsky, D.; Paredes, P.E. SAD: A Stress Annotated Dataset for Recognizing Everyday Stressors in SMS-like Conversational Systems. In Proceedings of the CHI '21 Extended Abstracts, 2021. <https://doi.org/10.1145/3411763.3451799>.
185. SKMCH&RC. Shaikat Khanum Memorial Cancer Hospital & Research Centre Cancer Registry. Website, 2024.
186. Yu, T.; Zhang, R.; Yang, K.; Yasunaga, M.; Wang, D.; Li, Z.; Ma, J.; Li, I.; Xue, Q.; Vyas, S.; et al. Spider: A Large-Scale Human-Labeled Dataset for Complex and Cross-Domain Semantic Parsing and Text-to-SQL Task. In Proceedings of the EMNLP, 2018.
187. AIMH. SWMH Dataset Card. Hugging Face Dataset, 2025.
188. Gao, L.; Biderman, S.; Black, S.; et al. The Pile: An 800GB Dataset of Diverse Text for Language Modeling. *arXiv preprint arXiv:2101.00027* **2021**.
189. Vishnubhotla, K.; Cook, P.; Hirst, G. Emotion Word Usage in Tweets from US and Canada (TUSC). *arXiv preprint arXiv:2204.04862* **2022**.
190. Jr., C.R.J.; Weiner, M.W.; et al. The Alzheimer's Disease Neuroimaging Initiative (ADNI): MRI methods. *Journal of Magnetic Resonance Imaging* **2008**.
191. Litjens, G.; Bandi, P.; et al. 1399 H&E-stained sentinel lymph node sections of breast cancer metastases: the CAMELYON dataset. *GigaScience* **2018**. Article ID giy065.
192. Wang, D.; Wang, X.; Wang, L.; Li, M.; Da, Q.; Liu, X.; Gao, X.; Shen, J.; He, J.; Shen, T.; et al. ChestDR: Thoracic Diseases Screening in Chest Radiography, 2023. <https://doi.org/10.6084/m9.figshare.22302775.v1>.
193. Wang, X.; Peng, Y.; Lu, L.; Lu, Z.; Bagheri, M.; Summers, R.M. ChestX-ray8: Hospital-scale Chest X-ray Database and Benchmarks on Weakly-Supervised Classification and Localization of Common Thorax Diseases. In Proceedings of the CVPR, 2017.
194. Wang, D.; Wang, X.; Wang, L.; Li, M.; Da, Q.; Liu, X.; Gao, X.; Shen, J.; He, J.; Shen, T.; et al. ColonPath: Tumor Tissue Screening in Pathology Patches, 2023. <https://doi.org/10.6084/m9.figshare.22302799.v1>.
195. Tschandl, P.; Rosendahl, C.; Kittler, H. The HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Scientific Data* **2018**.
196. Deng, J.; Dong, W.; Socher, R.; Li, L.J.; Li, K.; Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In Proceedings of the 2009 IEEE conference on computer vision and pattern recognition. Ieee, 2009, pp. 248–255.

197. Wang, D.; Wang, X.; Wang, L.; Li, M.; Da, Q.; Liu, X.; Gao, X.; Shen, J.; He, J.; Shen, T.; et al. A real-world dataset and benchmark for foundation model adaptation in medical image classification. *Scientific Data* **2023**, *10*, 574.
198. Jaeger, S.; Candemir, S.; Antani, S.; Wang, Y.X.J.; Lu, P.X.; Thoma, G. Two public chest X-ray datasets for computer-aided screening of pulmonary diseases. *Quantitative Imaging in Medicine and Surgery* **2014**.
199. Wang, D.; Wang, X.; Wang, L.; Li, M.; Da, Q.; Liu, X.; et al. NeoJaundice: Neonatal Jaundice Evaluation in Demographic Images. [https://springernature.figshare.com/articles/dataset/NeoJaundice\\_Neonatal\\_Jaundice\\_Evaluation\\_in\\_Demographic\\_Images/22302559](https://springernature.figshare.com/articles/dataset/NeoJaundice_Neonatal_Jaundice_Evaluation_in_Demographic_Images/22302559), 2023.
200. Peking University & Grand-Challenge. Ocular Disease Intelligent Recognition (ODIR-2019) Dataset. <https://odir2019.grand-challenge.org/dataset/>, 2019.
201. Rath, S.R. Diabetic Retinopathy 224x224 (2019 Data). <https://www.kaggle.com/datasets/sovirath/diabetic-retinopathy-224x224-2019-data>, 2019.
202. National Eye Institute. NEI Age-Related Eye Disease Study (AREDS). [https://www.ncbi.nlm.nih.gov/projects/gap/cgi-bin/study.cgi?study\\_id=phs000001.v3.p1](https://www.ncbi.nlm.nih.gov/projects/gap/cgi-bin/study.cgi?study_id=phs000001.v3.p1), 2001.
203. Nakayama, L.F.; de Souza, F.; Vieira, G.; et al. BRSET: A Brazilian Multilabel Ophthalmological Dataset of Fundus Photographs. *PLOS Digital Health* **2024**. <https://doi.org/10.1371/journal.pdig.0000454>.
204. Kennedy, D.N.; Haselgrove, C.; Hodge, S.M.; et al. CANDIShare: A Resource for Pediatric Neuroimaging Data. *Neuroinformatics* **2012**.
205. Wu, J.T.; Moradi, M.; Wang, H.; et al. Chest Imagenome Dataset for Clinical Reasoning. *arXiv preprint arXiv:2108.00316* **2021**.
206. Irvin, J.; Rajpurkar, P.; Ko, M.; Yu, Y.; Ciurea-Ilcus, S.; Chute, C.; Marklund, H.; Haghighi, B.; Ball, R.; Shpanskaya, K.; et al. Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In Proceedings of the Proceedings of the AAAI conference on artificial intelligence, 2019, Vol. 33, pp. 590–597.
207. Afshar, P.; Heidarian, S.; Enshaei, N.; et al. COVID-CT-MD: COVID-19 Computed Tomography Scan Dataset Applicable in Machine Learning and Deep Learning. *Scientific Data* **2021**. <https://doi.org/10.1038/s41597-021-00900-3>.
208. Tian, Y.; Luo, Y.; Liu, F.; et al. FairSeg: A Large-Scale Medical Image Segmentation Dataset for Fairness Learning. *ICLR 2024 (proc. abstract) / arXiv:2311.02189* **2024**.
209. Luo, Y.; Tian, Y.; Zhang, J.; et al. FairCLIP: Harnessing Fairness in Vision-Language Learning. In Proceedings of the CVPR, 2024. Introduces Harvard-FairVLMed10k dataset.
210. Luo, Y.; Tian, Y.; Liu, F.; et al. Harvard Glaucoma Fairness: A Retinal Nerve Disease Dataset for Fairness Learning and Fair Identity Normalization. *IEEE Transactions on Artificial Intelligence* **2024**. Dataset page: Harvard GF3300.
211. IRCAD. 3D-IRCADb-01: Liver Segmentation Dataset. <https://www.ircad.fr/research/data-sets/liver-segmentation-3d-ircadb-01/>, 2010.
212. Heller, N.; Sathianathan, N.; Kalapara, A.; et al. The KiTS19 Challenge Data: 300 Kidney Tumor Cases with Clinical Context, CT Semantic Segmentations, and Surgical Outcomes. *arXiv preprint arXiv:1904.00445* **2019**.
213. MIDRC Consortium. Medical Imaging and Data Resource Center (MIDRC). <https://www.midrc.org/midrc-data>, 2021.
214. Johnson, A.E.; Pollard, T.J.; Berkowitz, S.J.; Greenbaum, N.R.; Lungren, M.P.; Deng, C.y.; Mark, R.G.; Horng, S. MIMIC-CXR, a de-identified publicly available database of chest radiographs with free-text reports. *Scientific data* **2019**, *6*, 317.
215. Johnson, A.E.; Pollard, T.J.; Shen, L.; Lehman, L.w.H.; Feng, M.; Ghassemi, M.; Moody, B.; Szolovits, P.; Anthony Celi, L.; Mark, R.G. MIMIC-III, a freely accessible critical care database. *Scientific data* **2016**, *3*, 1–9.
216. Johnson, A.E.W.; Bulgarelli, L.; Shen, L.; et al. MIMIC-IV, a Freely Accessible Electronic Health Record Dataset. *Scientific Data* **2023**. <https://doi.org/10.1038/s41597-022-01899-x>.
217. Child and Family Data Archive. Migrant and Seasonal Head Start Study (MSHS), United States, 2017-2018. <https://www.childandfamilydataarchive.org/cfda/archives/cfda/studies/37348>, 2019.
218. Abdunour, R.E.E.; Kachalia, A. Deliberate Practice at the Virtual Bedside to Improve Diagnostic Reasoning. *NEJM* **2022**. Describes NEJM Healer approach.
219. Kass, M.A.; Heuer, D.K.; Higginbotham, E.J.; et al. The Ocular Hypertension Treatment Study: A Randomized Trial Determines that Topical Ocular Hypotensive Medication Delays or Prevents the Onset of Primary Open-angle Glaucoma. *Archives of Ophthalmology* **2002**. <https://doi.org/10.1001/archophth.120.6.701>.

220. Chaves, J.M.Z.; Patel, B.; Chaudhari, A.; et al. Opportunistic Assessment of Ischemic Heart Disease Risk Using Abdominopelvic CT and Medical Record Data: A Multimodal Explainable AI Approach. *NPJ Digital Medicine* **2023**. Releases the OL3I dataset.
221. Pacheco, A.G.C.; Lima, G.R.; Salomão, A.S.; et al. PAD-UFES-20: A Skin Lesion Dataset Composed of Patient Data and Clinical Images Collected from Smartphones. *Data in Brief* **2020**. <https://doi.org/10.1016/j.dib.2020.106221>.
222. Bustos, A.; Pertusa, A.; Salinas, J.M.; de la Iglesia-Vayá, M. PadChest: A Large Chest X-ray Image Dataset with Multi-label Annotated Reports. *Medical Image Analysis* **2020**. <https://doi.org/10.1016/j.media.2020.101797>.
223. Kovalyk, O.; Remeseiro, B.; Ortega, M.; et al. PAPILA: Dataset with Fundus Images and Clinical Data of Both Eyes of the Same Patient for Glaucoma Assessment. *Scientific Data* **2022**. <https://doi.org/10.1038/s41597-022-01388-1>.
224. Nguyen, H.Q.; Lam, K.; Le, L.T.; et al. VinDr-CXR: An Open Dataset of Chest X-rays with Radiologist's Annotations. *Scientific Data* **2022**. <https://doi.org/10.1038/s41597-022-01498-w>.
225. Manès, V.J.; Han, H.; Han, C.; Cha, S.K.; Egele, M.; Schwartz, E.J.; Woo, M. The art, science, and engineering of fuzzing: A survey. *IEEE Transactions on Software Engineering* **2019**, *47*, 2312–2331.
226. McKeeman, W.M. Differential testing for software. *Digital Technical Journal* **1998**, *10*, 100–107.
227. Chen, Z.; Zhang, J.M.; Sarro, F.; Harman, M. Fairness improvement with multiple protected attributes: How far are we? In Proceedings of the Proceedings of the IEEE/ACM 46th international conference on software engineering, 2024, pp. 1–13.
228. Chen, Z.; Li, X.; Zhang, J.M.; Sun, W.; Xiao, Y.; Li, T.; Lou, Y.; Liu, Y. Software Fairness Dilemma: Is Bias Mitigation a Zero-Sum Game? *Proc. ACM Softw. Eng.* **2025**, *2*, 1780–1801.
229. Chen, Z.; Li, X.; Zhang, J.M.; Sarro, F.; Liu, Y. Diversity Drives Fairness: Ensemble of Higher Order Mutants for Intersectional Fairness of Machine Learning Software. In Proceedings of the 47th IEEE/ACM International Conference on Software Engineering, ICSE 2025, Ottawa, ON, Canada, April 26 - May 6, 2025, 2025, pp. 743–755.

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.