



 Latest updates: <https://dl.acm.org/doi/10.1145/3793542>

SURVEY

Fair Foundation Models for Medical Image Analysis: Challenges and Perspectives



PDF Download
3793542.pdf
26 January 2026
Total Citations: 0
Total Downloads: 0

Accepted: 13 January 2026
Revised: 30 November 2025
Received: 28 March 2025

[Citation in BibTeX format](#)

Fair Foundation Models for Medical Image Analysis: Challenges and Perspectives

DILERMANDO QUEIROZ, Federal University of São Paulo, Brazil

ANDERSON CARLOS, Federal Institute of Goiás, Brazil

ANDRÉ ANJOS, Idiap Research Institute, Switzerland

LILIAN BERTON, Federal University of São Paulo, Brazil

Ensuring equitable Artificial Intelligence (AI) in healthcare demands systems that make unbiased decisions across all demographic groups, bridging technical innovation with ethical principles. Foundation Models (FMs), trained on vast datasets through self-supervised learning, enable efficient adaptation across medical imaging tasks while reducing dependency on labeled data. These models demonstrate potential for enhancing fairness, though significant challenges remain in achieving consistent performance across demographic groups. Our review indicates that effective bias mitigation in FMs requires systematic interventions throughout all stages of development. While previous approaches focused primarily on model-level bias mitigation, our analysis reveals that fairness in FMs requires integrated interventions throughout the development pipeline, from data documentation to deployment protocols. This comprehensive framework advances current knowledge by demonstrating how systematic bias mitigation, combined with policy engagement, can effectively address both technical and institutional barriers to equitable AI in healthcare. The development of equitable FMs represents a critical step toward democratizing advanced healthcare technologies, particularly for underserved populations and regions with limited medical infrastructure and computational resources.

CCS Concepts: • Applied computing → Life and medical sciences.

Additional Key Words and Phrases: Foundation Models, Fairness, World Health

Acknowledgments

We thank Fundação de Amparo à Pesquisa do Estado de São Paulo (FAPESP), grants 21/14725-3 and 23/12493-3, Conselho Nacional de Desenvolvimento Científico e Tecnológico (CNPQ), Swiss National Science Foundation (SNSF) under Grant No. 200021E_214653, Santos Dumont supercomputer at the LNCC. We would like thank Lucas Tosta for the designs.

1 Introduction

AI in medicine offers transformative potential through improving access to diagnostics and enhancing the speed and quality of medical care [111, 170]. AI tools extend healthcare services, particularly in resource-limited regions, thereby making care more efficient and accessible [152]. However, these advancements raise critical ethical concerns that emphasize the need for a fair and equitable distribution of benefits across all populations [105, 126, 130]. We must develop and apply these technologies responsibly to uphold bioethical principles of

Authors' Contact Information: Dilermando Queiroz, dilermando.queiroz@unifesp.br, Federal University of São Paulo, São José dos Campos, São Paulo, Brazil; Anderson Carlos, Federal Institute of Goiás, Brazil; André Anjos, Idiap Research Institute, Switzerland; Lilian Berton, Federal University of São Paulo, Brazil.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM 2637-8051/2026/1-ART

<https://doi.org/10.1145/3793542>

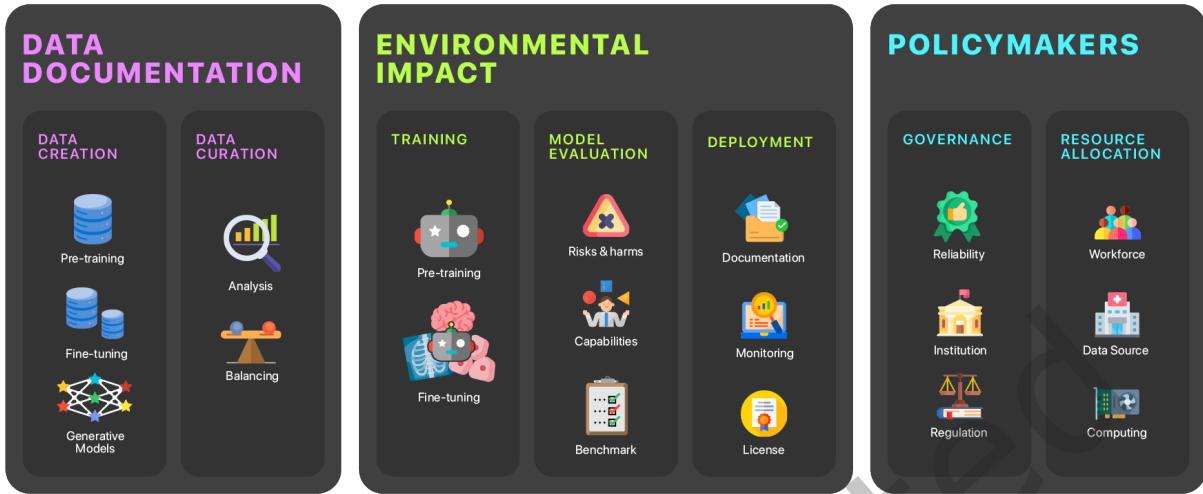


Fig. 1. **Conceptual framework.** This figure delineates sequential phases in FMs development, illustrating principal challenges for achieving fairness. Three focal domains are emphasized: data documentation, curation dataset ensuring diversity and quality to detect and mitigate bias; environmental impact, spanning training, deployment, and resource efficiency; and policymakers, establish governance, standards, and resource distribution to ensure ethical, equitable FMs access.

justice, autonomy, beneficence, and non-maleficence to prevent discrimination and promote inclusive healthcare for all [14].

Governments worldwide are establishing regulatory frameworks for AI across critical sectors. The EU AI Act [47], the first comprehensive regulation of its kind, has introduced a risk-based classification system for AI applications. In parallel, the U.S. Office for Civil Rights has enacted specific protections against algorithmic bias in healthcare through the Affordable Care Act [144]. These regulatory initiatives reflect a coordinated global effort to implement evidence-based guidelines that ensure the fairness, safety, and equity of AI systems [35, 64, 86, 93, 95, 117, 140, 146].

FMs serve as essential components in AI by enabling diverse tasks across text, image, video, audio, and other domains through their versatility and scalability [17]. Through training on massive datasets, these models capture broad patterns within and across domains. However, significant challenges persist in developing ethical models that can effectively identify and reduce inherent biases [54, 78, 89]. While bias persists across AI systems, FMs demonstrate significant potential for their mitigation [124, 145], thus creating opportunities for unified models that equitably serve diverse populations while driving greater inclusion in AI.

The production of FMs requires substantial resources, including specialized labor, large datasets, and significant computational infrastructure. These high costs restrict development capabilities to select countries, institutions, and companies, thereby increasing the risk of global inequalities. Stakeholders in regions such as Africa have indicated their inability to develop such models, which potentially widens the disparity between populations that benefit from AI and those that do not [3]. Addressing this challenge requires developing strategies that facilitate the creation of globally accessible FMs while enabling broad participation in their production.

Bias mitigation must occur throughout the proposed development process of FMs, as illustrated in Figure 1. This process requires thorough data documentation during the curation and creation phases to ensure dataset diversity and representativeness. The consideration of environmental impacts during training, model evaluation, and deployment phases enhances fairness. Policymakers fulfill a crucial role through the enactment of laws

and the allocation of resources that support equitable AI practices. Through the integration of bias mitigation strategies at each developmental stage, researchers can develop inclusive and responsible AI models that serve diverse populations effectively.

A substantial body of research has empirically demonstrated inherent biases in FMs and established initial frameworks for fairness evaluation [54, 72, 78, 89]. The FairMedFM benchmark, for instance, provides a standardized basis for these empirical assessments [72]. While comprehensive investigations of trustworthiness in medical imaging FMs have been conducted [134], these studies primarily address fairness as one component within broader trustworthiness considerations. Additionally, existing analyses provide thorough technical perspectives on medical imaging FMs [88, 170] but lack systematic examination of fairness considerations. Current reviews of AI fairness in healthcare [31, 45, 106, 130, 162] have not adequately addressed the distinct challenges posed by FMs. To our knowledge, no study has adopted a broader conceptual lens to address the interconnected ethical, environmental, and governance challenges in ensuring fairness in FMs. To address this gap, this review investigates the development lifecycle from a narrative perspective, identifying key opportunities and challenges for fair FMs.

Our contributions are:

- An investigation of potentials and gaps in bias mitigation methods across all stages of FMs development.
- A global assessment of dataset distribution patterns reveals inequalities in representation across different regions.
- An analysis of potential policy interventions for ensuring and promoting the development of more equitable FMs.

2 Background and Taxonomy

This section delineates the fundamental principles underlying fairness and FMs. We also present a table summarizing key concepts and strategies relevant to our review (Table 1).

2.1 Fairness

Principles of Trustworthy AI. The development of medical image analysis AI systems, to achieve trustworthiness, is guided by six fundamental principles: fairness, universality, traceability, usability, robustness, and explainability [86]. Among these, fairness, which ensures non-discriminatory outcomes across diverse patient populations, represents a critical determinant of ethical AI implementation. Our analysis centers on fairness and examines its interdependencies with other principles, specifically how it interacts with robustness to ensure reliable system performance and with traceability to maintain systematic accountability, thus demonstrating how these principles collectively contribute to trustworthy AI in clinical settings. We specifically center hallucinations as a critical reliability issue, as these generated false but plausible outputs can undermine fairness by causing biased or unequal clinical outcomes.

Healthcare Disparities. Differences in national healthcare delivery capacities result in variable health outcomes across populations [29, 155]. Within countries that maintain public healthcare systems, disparities persist because of multiple factors, including race, gender, age, ethnicity, body mass index, education, insurance status, and geographic location [12, 153]. These determinants influence an individual's capacity to access treatment, obtain quality care, and achieve favorable health outcomes.

AI Bias Manifestation. AI models frequently assimilate and reproduce biases inherent in their training data, subsequently reflecting and intensifying societal inequalities. These models often depend on spurious correlations, employing computational shortcuts that amplify existing biases [51, 53, 178]. A well-designed fair algorithm produces impartial decisions by ensuring equitable outcomes across demographic groups without discrimination based on protected attributes such as race, gender, or age. Despite the implementation of controlled

Table 1. **Taxonomy.** Key concepts and strategies encompassing fairness, protected attributes, foundation models, vision-language models, and learning approaches relevant to scalable and equitable AI systems.

| Term | Summary | Key Concepts |
|-------------------------------|--|---|
| Fairness | The principle that ensures equitable treatment and non-discriminatory outcomes across diverse groups defined by protected attributes in algorithmic decision-making. | Equity, bias mitigation, performance parity, protected groups. |
| Hallucination | The generation of plausible but false or fabricated information not grounded in real-world facts. | Misinformation, clinical risk, reliability, uncertainty, inaccurate knowledge. |
| Protected Attributes | Patient demographic and clinical metadata that require safeguarding to prevent bias in algorithmic decision-making, ensuring fairness across groups defined by these attributes (e.g., age, gender). | Protected attributes, bias prevention, demographic metadata, group fairness. |
| Foundation Models (FMs) | Large-scale pretrained models adaptable to diverse tasks, with emphasis on self-supervised and weakly-supervised training strategies. | Pre-training, downstream task, scalability, robustness. |
| Vision Language Models (VLMs) | Models that integrate visual and linguistic data processing, typically trained using weakly supervised learning to leverage large-scale paired image-text data. | Multimodal learning, weakly supervised learning, visual-linguistic integration, FMs. |
| World Models | Models that simulate and predict the dynamics of the real world, enabling anticipation of future states and informed decision-making. | Environment Simulation, latent space prediction, planning. |
| Generative Models | Models that learn data patterns to create new, original content resembling the training data. | Content generation, hallucinations, synthetic data. |
| Knowledge Agglomeration | Knowledge agglomeration synthesizes knowledge by distilling diverse representations from multiple FMs into a single, unified agglomerative model. | Knowledge distillation, Model merging, agglomerative models. |
| Self-supervised Learning | Training method using large unlabeled data by leveraging inherent data structures. | Unlabeled data, Masked Autoencoder (MAE) [61], Simple Framework for Contrastive Learning of Visual Representations (SimCLR) [33], Image-based Joint-Embedding Predictive Architecture (I-JEPA) [9]. |
| Weakly-supervised Learning | Learning approach that leverages limited, noisy, or imprecise labels to guide model training, commonly used in large-scale and multimodal datasets. | Noisy labels, limited supervision, semi-supervised methods, VLMs, Contrastive Language-Image Pre-training (CLIP) [125], Sigmoid Loss for Language Image Pre-Training (SigLIP) [166]. |

datasets and balanced groups, bias can manifest through various factors, including image complexity and labeling inconsistencies [37]. Consequently, unintended biases may emerge even under optimal conditions, highlighting the necessity for continuous evaluation and mitigation procedures to maintain algorithmic fairness.

Bias Mitigation Strategies. Previous research has established three primary classifications for bias mitigation strategies: pre-processing, in-processing, and post-processing [31, 45, 106, 162]. Pre-processing approaches modify datasets through demographic representation balancing, protected feature removal (such as race and gender), or synthetic data augmentation to enhance diversity [22, 24, 116]. In-processing methods alter the training process by integrating fairness constraints into the loss function or implementing adversarial training to prevent bias acquisition [26, 154, 164]. Post-processing techniques adjust model predictions after training through methods such as prediction calibration to satisfy fairness criteria [34, 74, 123].

Framework Integration. Our research extends traditional bias mitigation strategies (pre-processing, in-processing, and post-processing) into a comprehensive framework encompassing all stages of FMs development. Through the integration of bias mitigation efforts across all phases and the inclusion of policymakers in technical discussions, as illustrated in Figure 1, this framework systematically addresses biases throughout the development process. The approach responds to significant concerns regarding the potential of these models to amplify global economic inequalities. The incorporation of policymakers ensures the integration of ethical considerations and regulatory measures into FMs development, thus mitigating potential adverse effects on global economic disparities.

2.2 Foundation Models

FMs are trained on extensive datasets that can be efficiently adapted to multiple downstream tasks through fine-tuning, thereby eliminating the necessity of training specialized models from scratch. The Vision Transformer (ViT) [44] serves as a leading FMs architecture in computer vision for extracting fundamental image features. The integration of self-supervised learning techniques, such as MAE [61] and contrastive learning methods like SimCLR [33], enables these models to learn directly from large volumes of unlabeled data, thus largely mitigating the need for costly manual labeling processes. The I-JEPA introduces a self-supervised learning approach that predicts semantic embeddings directly, offering enhanced efficiency compared to pixel-reconstruction methods [9].

Capabilities and Applications. Through the utilization of massive unlabeled datasets, advanced architectures, and self-supervised learning techniques, FMs acquire comprehensive image representations within their training domains. The fine-tuning process facilitates task-specific adaptation through minimal adjustments and limited labeled data requirements. Several key capabilities establish their essential role in scalable, real-world applications: rapid task adaptation, robustness to distribution shifts, and efficient utilization of labeled data [10, 143, 156, 175]. For example, FMs improve diagnostic accuracy by 11.5% relative to supervised methods and, in out-of-distribution settings, achieve comparable performance using only 1–33% of the labeled data [10].

Multimodal Integration. These models facilitate training across diverse data types within a unified architectural framework. VLMs demonstrate this capability by integrating visual and linguistic processing through weakly supervised learning that combines image captions with contrastive losses, such as CLIP [125]. This integration supports simultaneous interpretation and generation of information across modalities. Subsequent work has improved training by replacing contrastive loss with a sigmoid-based loss, as in SigLIP [167], while other studies highlight the crucial role of data quality over model architecture or pretraining objectives [161].

Domain Specialization. Domain-specific FMs constitute specialized architectures tailored to address the distinctive characteristics of particular domains, ranging from natural image processing [118], medical imaging applications [10, 156] to VLMs tailored for medical domain [133, 150]. In contrast to general-purpose models, these specialized frameworks concentrate on the unique data structures and analytical challenges inherent to their respective domains. Medical imaging presents a notable example, where image characteristics vary substantially in resolution, ranging from whole-organ visualization to cellular-level structures, with data originating from diverse acquisition modalities including X-ray radiography, computed tomography (CT), magnetic resonance imaging (MRI), and ultrasonography [170].

3 Data Documentation

This section delineates the essential components of data documentation, encompassing both the data creation phase and subsequent curation processes.

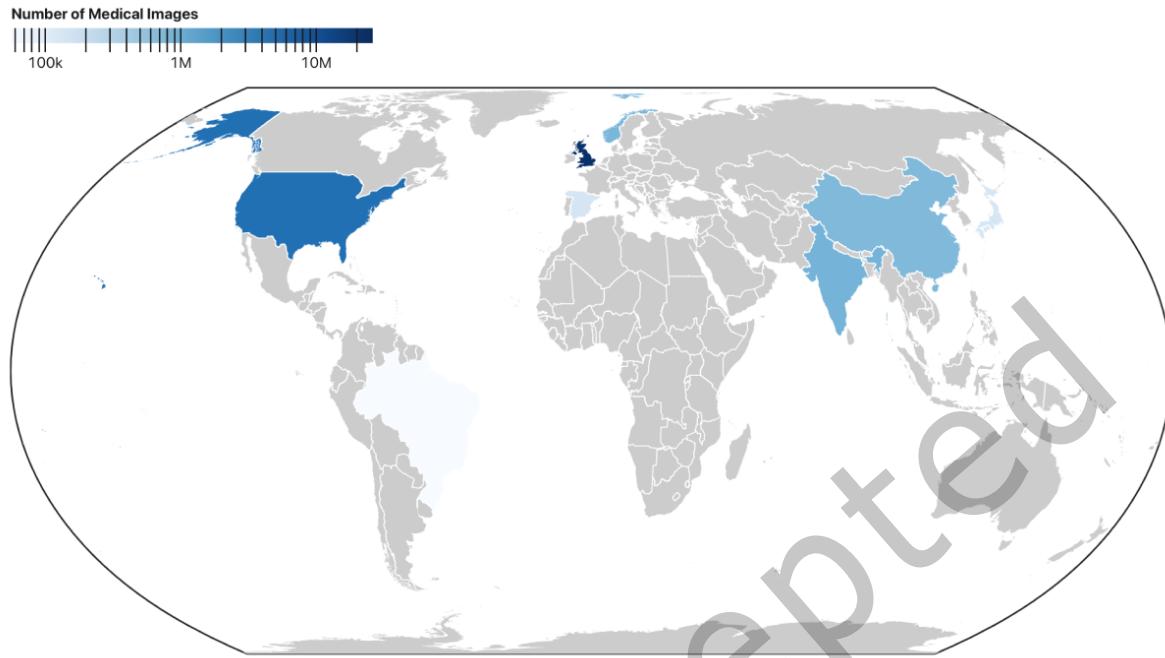


Fig. 2. Global distribution of medical imaging data: This geographic visualization depicts the volume of medical imaging datasets by country, excluding multi-country datasets listed in Table 2. The figure highlights pronounced disparities in data representation, underscoring its critical role in the development of equitable FMs. Notably, a limited number of countries account for the majority of available medical imaging data.

3.1 Data Creation

The first step in FMs development involves comprehensive data collection. Healthcare facilities generate medical images through specialized imaging equipment that captures various anatomical structures. The pre-training phase utilizes extensive unlabeled datasets, enabling FMs to acquire general representations and discover underlying patterns without annotation requirements. For subsequent fine-tuning processes, labeled data plays an essential role in enabling models to develop task-specific capabilities through the refinement of learned representations. Both unlabeled and labeled datasets must maintain high quality and diversity standards, as these characteristics fundamentally influence FMs performance and generalization capabilities [40, 118].

Inclusion Criteria. Datasets were identified for inclusion in Table 2 through a structured, multi-step process. The initial step involved identifying the pre-training datasets for the principal FMs listed in a recent survey [134]. To enhance geographic representation, a supplementary search was conducted on the Hugging Face and PhysioNet platforms to identify datasets from underrepresented regions. A minimum size threshold of 10,000 images was applied to all datasets to ensure they represent large-scale collections. Finally, metadata for each selected dataset were manually extracted from the corresponding publication and official website to populate the table.

While not exhaustive, this methodological approach enabled the systematic curation of 74,184,415 medical images, representing a comprehensive cross-section of diverse imaging modalities and anatomical regions. The comprehensive review encompassed unimodal imaging and multimodal image-text paired datasets, spanning diverse imaging modalities and anatomical regions. This methodological documentation framework facilitates

Table 2. Medical imaging datasets and their characteristics. Overview of publicly available medical imaging datasets, ordered by size and annotated with key attributes including computational framework, anatomical region, and data accessibility parameters.

| Dataset | Model | Language | Region | Modality | Images | Demographics | Origin | License |
|---------------------------------|------------|----------|-----------|------------|------------|--------------|----------|-------------|
| Moorfields BioResource 001 [50] | Vision | – | Ocular | Multimodal | 26,548,820 | A | UK | Restricted |
| MedTrinity-25M[159] | Multimodal | English | Multiple | Multimodal | 25,000,000 | – | Multiple | Group |
| PMC-15M[171] | Multimodal | English | Multiple | Multimodal | 15,000,000 | – | Multiple | CC BY-SA |
| BRATS24-MICCAI[38] | Vision | – | Cerebral | MRI | 2,535,132 | – | USA | CC BY 4.0 |
| PMC-OA[90] | Multimodal | English | Multiple | Multimodal | 1,600,000 | A | Multiple | OpenRAIL |
| RadImageNet[107] | Vision | – | Multiple | Multimodal | 1,350,000 | – | USA | CC BY 4.0 |
| ISIC[165] | Vision | – | Dermal | Dermoscopy | 1,162,456 | A, S | Multiple | CC BY-NC-SA |
| TCGA[76] | Vision | – | Dermal | Histology | 1,142,221 | A, R, S | USA | CC BY-NC-SA |
| HyperKvasir[21] | Vision | – | Colon | Endoscopy | 1,000,000 | – | Norway | CC BY 4.0 |
| BRATS-ISBI[75] | Vision | – | Cerebral | MRI | 987,340 | – | Multiple | CC BY 4.0 |
| BHX[129] | Vision | – | Cerebral | MRI | 973,908 | – | India | CC BY 4.0 |
| LDPolypVideo[101] | Vision | – | Colon | Endoscopy | 901,666 | – | China | – |
| MIMIC-CXR-JPG[73] | Multimodal | English | Pulmonary | Radiograph | 370,955 | A, R, S | USA | PHDL |
| CheXpert[68] | Multimodal | English | Pulmonary | Radiograph | 222,793 | A, R, S | USA | RUA |
| PadChest[23] | Multimodal | Spanish | Pulmonary | Radiograph | 160,868 | A, S | Spain | RUA |
| SUN-SEG[69] | Vision | – | Colon | Endoscopy | 158,690 | A, S | Japan | MIT |
| NIH-CXR14[149] | Vision | – | Pulmonary | Radiograph | 112,120 | A, S | USA | CC0 |
| Kermany et al.[77] | Vision | – | Ocular | OCT | 108,312 | – | USA | CC BY 4.0 |
| EyePACS[59] | Vision | – | Ocular | Fundus | 92,501 | – | India | CC BY 4.0 |
| BRAX[128] | Vision | – | Pulmonary | Radiograph | 40,967 | A, S | Brazil | PHDL |
| MURA[127] | Vision | – | Multiple | Radiograph | 40,561 | – | USA | CC BY 4.0 |
| ASU-Mayo[138] | Vision | – | Colon | Endoscopy | 19,400 | – | USA | – |
| BRSET[114] | Vision | – | Ocular | Fundus | 16,266 | A, S, N | Brazil | PHDL |
| Harvard-FairVLMed[98] | Multimodal | English | Ocular | SLO | 10,000 | A, R, S, E | USA | CC BY-NC-SA |

Demographics: A, Age; R, Race; S, Sex; N, Nationality; E, Ethnicity; –, None reported

Licenses: CC, Creative Commons (BY, Attribution; NC, NonCommercial; SA, ShareAlike); PHDL, PhysioNet Health Data License; RUA, Research Use Agreement

a systematic assessment of the current medical imaging data ecosystem while elucidating opportunities for enhanced demographic representation in FMs development. The geographic distribution of dataset origins appears in Figure 2.

Pre-training. The utilization of large unlabeled datasets for FMs training significantly reduces dependence on costly and time-intensive labeled data collection. Advanced learning methodologies, particularly self-supervised learning approaches [9, 33, 61], facilitate effective utilization of unlabeled data while minimizing labeling bias. This methodology proves especially valuable in medical contexts, where individual clinician annotation preferences, influenced by patient attributes, may introduce systematic biases [16]. Using unlabeled data enables FMs to identify generalizable patterns while simultaneously reducing both cost constraints and annotation-induced biases.

The development of robust datasets for FMs presents significant challenges in ensuring comprehensive representation across populations, imaging modalities, equipment specifications, and disease classifications. The accumulation of large-scale data frequently amplifies existing biases and imbalances that reflect underlying disparities in healthcare access and infrastructure. As illustrated in Figure 2, the predominant source of available datasets resides in developed nations, which substantially constrains the diversity of documented diseases and patient demographics. This geographic concentration of data resources is particularly evident in regions such as Africa, where large-scale medical imaging datasets remain notably absent (see Section 4.1 for a detailed discussion of the implications of such underrepresentation).

Fine-tuning. Smaller, more accurate, and more representative datasets are ideal for specializing FMs for specific tasks. This specialization is achieved through supervised learning on labeled datasets. Due to their data efficiency and generalization capabilities [10, 143, 156, 175], FMs require less data than training supervised models from scratch. Consequently, fine-tuning FMs offers significant advantages, as it requires fewer computational resources and enables countries and institutions to adapt these models to their specific tasks with greater accessibility.

Precise labels and comprehensive patient information, including gender, age, and race, facilitate the evaluation of model performance and enable systematic bias identification [57, 68, 114, 128]. This demographic information proves essential for model evaluation, as it enables more rigorous identification and mitigation of potential biases (discussed in the evaluation section 4.2). The availability of labeled data supports dataset balancing across specific attributes, thereby enhancing model fairness and performance. This represents a significant advantage over unlabeled data, which presents greater challenges in achieving these objectives [8].

Generative Models. Generative models constitute a driving force behind recent advancements in AI and its applications. These models learn to approximate the probability distribution of output features conditioned on input features, enabling the generation of novel data instances that closely approximate the characteristics of the training data. Advanced techniques, including generative adversarial networks (GANs) [94], variational autoencoders (VAEs) [121], diffusion models [122], have significantly enhanced the capabilities of generative AI systems.

Data augmentation through generative approaches facilitates both pre-training and fine-tuning stages. These models enable dataset balancing across protected attributes through synthetic data generation for underrepresented patient groups and enhance model robustness by generating challenging cases, reducing the fairness gap in chest radiograph classifiers trained on synthetic and real images by 44.6% [37]. Empirical studies demonstrate that fine-tuning models with augmented data improve both fairness and robustness, resulting in enhanced performance across diverse populations [82]. However, generative models may inadvertently amplify biases present in training data, such as generating synthetic medical images that systematically underrepresent darker skin tones [169], which necessitates robust bias mitigation strategies when utilizing synthetic data for augmentation.

Hallucinations. Understood as model-generated content that is not aligned with the real world, constitute a major source of misinformation that undermines the reliability of text and image generation models. They are typically associated with high uncertainty and incomplete or inaccurate knowledge [66, 70, 92, 102], particularly in instances where the model lacks relevant training data and exhibits low confidence [80, 92, 158]. As illustrated in Figure 2, large regions of the world are underrepresented in current datasets, which implies that models are more prone to hallucinate in these contexts because their underlying knowledge is partial or missing [32]. Nevertheless, it is important to note that hallucinations also occur in settings where the model has the knowledge and model confidence is high [135], indicating that they remain a persistent concern even under apparently well-specified conditions.

3.2 Data Curation

Although increasing data volume in parallel with neural network size can enhance model performance [55, 56], the relationship between data quantity and model improvement is not consistently linear. Recent research demonstrates that systematic data curation plays a critical role in both natural language processing and computer vision tasks [40, 118]. Ensuring dataset fairness requires careful attention to multiple dimensions: diversity, global representativeness, equipment variability, disease representation, gender balance, and age distribution. Furthermore, data deduplication serves as an essential process for eliminating near-duplicate images, thereby reducing redundancy and enhancing dataset diversity and representativeness [85, 118]. In parallel, removing noisy samples, incomplete entries, and ambiguous abbreviations is crucial for mitigating hallucinations and achieve fair representations [62, 80].

Protected attributes. The acquisition of metadata and labels in medical imaging presents substantial challenges, primarily due to the resource-intensive nature of manual curation given data volume and complexity. While fairness research frequently presumes the availability of demographic data, such information often remains inaccessible due to legal, ethical, and practical constraints [7]. This limited access to demographic data creates significant obstacles for both data curation processes and the development of unbiased models [124]. Current methodologies attempting to address fairness without demographic metadata face fundamental limitations, including systematic biases, accuracy limitations, technical barriers, and compromised transparency, thus emphasizing the critical need for rigorous evaluation protocols and explicit methodological guidelines [8].

Text and image. Recent research on VLMs highlights that data curation plays a more critical role in CLIP model performance than the architecture itself [161]. While enhancements such as the addition of sigmoid loss contribute to performance [167], the success of VLMs primarily depends on abundant high-quality data and substantial compute resources for scaling [142]. The process of curation leverages text metadata to construct more relevant datasets by selecting key terms, applying substring matching, and most importantly, balancing the dataset through sub-sampling based on chosen metadata [161].

Scalable methods without protected attributes. FMs serve as highly effective feature extractors, enabling clustering-based approaches for automatic data curation that do not require protected attributes. This capability facilitates the enhancement of diversity and balance in large datasets [124, 147]. When compared to uncurated data, features derived from these automatically curated datasets demonstrate superior performance [147], for example by reducing a 4.44% gender imbalance to achieve an equal distribution of male and female samples in medical imaging [124]. The clustering techniques developed through this approach serve a dual purpose: they not only improve data quality but also provide a systematic method for identifying potential biases in models [124]. These findings establish FMs as versatile tools that excel not only in feature extraction but also in crucial data management tasks, including bias assessment and automated curation.

Data curation alone is insufficient. While data balancing contributes to model fairness, it does not fully eliminate inherent biases in AI systems. The integration of multiple data domains, particularly the combination of textual and visual data, can potentially introduce or intensify biases beyond those present in unimodal models [20, 60, 67]. Consequently, achieving fair downstream behavior requires comprehensive mitigation strategies beyond data balancing alone [4]. In the context of unimodal models, performance disparities in downstream tasks frequently correlate with class difficulty, where more challenging classes exhibit higher misclassification rates and decreased performance metrics [37].

4 Environmental Impact

This section examines the environmental considerations throughout FMs development, encompassing training requirements, model evaluation protocols, and deployment strategies.

4.1 Training

Fairness. Self-supervised learning models leveraging diverse, unlabeled datasets demonstrate enhanced fairness and inclusivity compared to supervised approaches, yielding outcomes characterized by increased robustness and reduced bias [56]. However, FMs retain inherent biases despite these advantages [54, 98, 145]. For example, in a chest radiograph classification task, female patients experienced a 6.8–7.8% decrease in performance on the “no finding” label [53]. The mitigation of such biases necessitates targeted interventions within the training loop. This optimization process faces dual challenges: the substantial computational requirements and the inherent difficulty of acquiring large-scale datasets that exclude protected attributes.

Computational Challenges. The substantial computational requirements of FMs arise from two primary factors: the necessity of extensive datasets during pre-training and the complexity of architectures involving

numerous parameters. These resource demands create significant entry barriers, constraining many institutions to fine-tuning existing pre-trained models rather than developing their own. Additionally, as established in the Data Creation section, this computational divide disproportionately affects certain countries, potentially exacerbating existing biases and societal inequities.

Bias. The lack of data representativeness, as illustrated in Figure 2, can cause models to perform disproportionately better on more prevalent or overrepresented data subsets, thereby limiting generalizability [37]. Models are known to internalize diverse biases related to complexity, shape, and other intrinsic features present in training data, which compromises their robustness [83, 137]. Hence, assembling datasets that comprehensively reflect the true population distribution is paramount to enhancing model robustness and mitigating bias-related performance degradation during training.

Bias Amplification in VLMs. Multimodal models not only perpetuate existing societal biases but also amplify them relative to single-modality systems [20, 60, 67]. The pairing of images and text reintroduces the well-documented problem of label bias [27, 148], and generating non-factual medical diagnoses and overconfidence in generated diagnoses [157]. Moreover, as shown in the Table 2, most multimodal datasets remain English-centric, which restricts their applicability in global contexts [13]. This limitation is reinforced by the fact that most multilingual solutions rely heavily on translation pipelines, which could introduce semantic drift and cultural misalignment [41], in the medical domain this could lead to a misdiagnosis of local disease.

Synthetic data. Integrating synthetically generated data during training offers a promising approach to enhancing fairness in FMs. Evidence demonstrates that models trained exclusively on synthetic data can match or exceed the accuracy of models trained with real data [79, 82, 112]. For instance, in chest radiography applications, models trained with synthetic data achieved a 6.5% accuracy increase in downstream classification performance compared to models trained with real data [112]. Models that integrate medical findings with patient demographics prove essential for generating well-distributed synthetic datasets, as exemplified by RoentGen-v2 [112]. Model performance depends critically on synthetic data quality. Traditional high-quality generation methods require substantial computational resources. Emerging techniques such as flow matching and consistency models address this limitation through end-to-end architectures that reduce computational demands while maintaining generation quality [52, 136].

4.1.1 Pre-training. Protected attributes. The computational challenges inherent in vision and language models stem from power-law scaling relationships, wherein incremental performance improvements necessitate exponential increases in computational resources [55, 56]. As established in the Data Creation section, the prevalent absence of metadata in large-scale datasets introduces significant complexities for training loop implementation in pre-trained FMs. These constraints fundamentally limit the application of fairness techniques during the training process. Furthermore, fairness methodologies developed for pre-trained FMs must demonstrate efficient scalability across expansive datasets and sophisticated architectural frameworks.

Data selection. The incorporation of systematic data curation into the training loop can be achieved through active selection strategies, wherein computational priority is assigned to data elements that maximize task performance contributions [131]. Research demonstrates that small curated models can facilitate the training of larger models through systematic identification of both straightforward and challenging image cases [48, 49]. Implementation of these selection methodologies yields dual benefits: reduced training duration for pre-trained models and enhanced quality of resultant outputs.

Loss. The Reducible Holdout Loss (RHO) employs a secondary model to identify three categories of data points: those that are learnable, those worth dedicating computational resources to learn, and those not yet acquired by the model [108]. One established approach for enhancing model robustness and fairness involves the integration of protected attributes into loss functions, thereby promoting equitable outcomes across demographic groups [103]. However, as examined in the Data Documentation section, the acquisition of protected attributes presents

significant practical challenges. To address this limitation, alternative methodologies utilizing clustering-based data curation [147] offer potential proxy measures for protected attributes [124].

World Models. Contemporary research in robust and unbiased model development prioritizes the prediction of representations within embedding spaces over traditional token-level forecasting approaches (e.g., word or pixel prediction) [84]. Within this framework, a novel class of FMs, termed World Models enables simultaneous training across multiple languages and modalities while maintaining scalability and minimizing bias [139]. In the visual domain, the I-JEPA implements this embedding-centric approach, yielding substantial improvements in robustness, scalability, and computational efficiency relative to MAE [9, 91].

Vision Language World Models. Current VLMs, such as LLaVA-Med [87] and MedViNT [172], exhibit notable challenges including overconfidence in diagnostic outputs, privacy breaches, and the perpetuation of health disparities, as mentioned in Section 4.1 [157]. A promising approach to mitigate these issues involves integrating language and vision within world models equipped with a planning mechanism [28]. This mechanism operates by evaluating the consequences of an action, generating multiple potential outcomes, and predicting the optimal course that minimizes a cost function as assessed by a critic model [28, 84]. Crucially, fairness constraints can be embedded into this critic model. As this component is a self-supervised language model, establishing robust fairness principles in the language domain becomes fundamental to mitigating systemic bias throughout the entire vision-language architecture.

Knowledge Agglomeration. Knowledge agglomeration has recently emerged as a promising approach for pre-training. This approach leverages knowledge from existing FMs with diverse perspectives, such as CLIP [4], DINO [119], and SAM [81], to train improved models. Recent examples, like RADIO [63], demonstrate that models trained via knowledge agglomeration can surpass their teacher's performance. In the medical imaging domain, models such as MedSAM [99] and MedCLIP [150] can be combined to enhance performance. Crucially, this approach allows integrating models with fairness-aware representations alongside those optimized for general performance, jointly improving subgroup representation while balancing the fairness-utility trade-off (see Section 4.2 for a detailed discussion of utility-fairness trade-offs).

4.1.2 Fine-tuning. Efficacy of Fine-tuning for Fairness Enhancement. Fine-tuning serves as a critical mechanism for aligning pre-trained FMs with targeted objectives. These models demonstrate exceptional capability in adapting to novel data distributions with minimal sample requirements [10]. Fine-tuning effectively mitigates bias through increased model sensitivity to training distribution, particularly when employing balanced and curated datasets [4].

Resource Optimization in Fine-tuning. In contrast to the computationally intensive pre-training phase, fine-tuning procedures can be executed with substantially reduced resource requirements, rendering this approach particularly advantageous for resource-constrained initiatives. The utilization of domain-specific pre-trained FMs, especially those optimized for medical imaging applications, further enhances computational efficiency in settings with limited infrastructure.

Parameter-efficient Fine-tuning. Methodologies such as LoRa [65] and QLoRa [42] enhance the efficiency of low-resource fine-tuning through selective parameter modification, targeting only specific subsets within the base model architecture. The impact of these optimization techniques on bias mitigation remains an active area of investigation, with current research providing inconclusive evidence regarding their effects on model fairness [43, 72]. Notably, alternative computational optimization strategies, including model pruning and differentially private training approaches, demonstrate increased bias manifestation within specific demographic subgroups [11, 141].

Mitigating Bias. Despite extensive research into bias mitigation strategies within deep learning frameworks, conventional fairness interventions demonstrate variable efficacy when applied to FMs [72]. The limitations of these approaches extend beyond FMs applications, with traditional deep learning implementations often

achieving only modest improvements in fairness metrics [177]. Advanced data augmentation strategies, including AutoAug, Mixup, and CutMix methodologies, demonstrate promise in addressing challenging FMs scenarios [37]. Furthermore, the integration of synthetically generated data during the fine-tuning process presents compelling evidence for enhanced fairness outcomes [82].

Blackbox FMs. The increasing commercialization of FMs and proliferation of their APIs has resulted in a distribution model where access is frequently restricted to embedding outputs, limiting direct interaction with model architectures. This constrained accessibility presents distinct challenges for bias and hallucinations mitigations in medical imaging applications. Recent research demonstrates, however, that effective bias and hallucinations eliminations techniques can be implemented without requiring access to internal model parameters, offering viable solutions for both open-source and proprietary model frameworks [71, 102].

4.2 Model Evaluation

Fairness Considerations and Metrics. Within deep learning research, fairness has become a critical consideration, with particular significance in medical imaging applications [31, 130, 134, 145, 162]. The evaluation of fairness encompasses two fundamental methodological approaches: individual fairness, which requires consistent model outputs for similar inputs, and group fairness, which assesses model performance across demographic categories defined by protected attributes such as race and gender. Despite the widespread adoption of group fairness metrics in practical applications, a significant number of influential FMs studies in medical imaging have conducted evaluations without incorporating explicit fairness metrics [10, 97, 174, 175], thus leaving critical questions about potential biases and their clinical implications unexplored.

Utility-Fairness Trade-offs. In real-world scenarios, models may use protected attributes as shortcuts for specific tasks, introducing bias [163]. However, bias does not always stem from such shortcuts; other factors like intensity-based and morphology-based effects can also influence model behavior [137]. This bias complexity creates a utility-fairness trade-off, where achieving statistical parity can reduce utility for some groups [151, 173]. Identifying the optimal trade-off point is essential to balance both metrics effectively [39, 120].

Hallucinations. Medical hallucinations arise in specialized tasks such as diagnostic reasoning, employing domain-specific terminology that masks critical inaccuracies [80]. Expert scrutiny remains essential for detecting these plausible yet incorrect outputs, which can delay appropriate interventions and compromise patient safety [80]. Large language model taxonomies typically categorize hallucinations into factual errors, outdated references, spurious correlations, incomplete reasoning chains, and fabricated sources across text and multimodal contexts [80]. Detecting and quantifying hallucinations is critical for ensuring model safety and reliability [30, 58, 80]. Exploring the relationship between protected attributes and hallucinations represents an emerging research area that advances fairness evaluations and addresses bias in generative models [80, 112]. In this context, rigorous quality checks of generated images are essential to remove hallucinations and mitigate bias; for example, in a dataset of 623,712 prompts, 58,558 (9.4%) failed quality control, predominantly due to violations of race-related criteria [112].

Robustness. A fundamental challenge in medical imaging applications lies in ensuring that fairness-aware models maintain their equitable performance when transitioning between different data distributions (A to B), particularly given the dynamic nature of population characteristics and deployment contexts [132]. Although adversarial training techniques enhance overall model robustness, the improvements typically manifest unevenly across different classes [100, 160]. Research indicates that fairness enhancement strategies can simultaneously strengthen model robustness [113]. However, contemporary deep learning systems frequently demonstrate inadequate robustness to protected attributes, specifically sex and gender variables [177]. Within FMs frameworks, the direct deployment of models without fine-tuning procedures further compromises robustness to these demographic attributes, thus constraining their efficacy across diverse application scenarios [124].

Data-Efficient Generalization. FMs demonstrate exceptional capabilities in data-efficient generalization, facilitating fine-tuning processes for medical imaging tasks with minimal labeled data requirements [10]. This characteristic proves particularly advantageous for resource-constrained institutions implementing model deployment within their specific data environments. Nevertheless, substantial uncertainty persists regarding the implications of limited labeled data usage on fairness metrics in FMs applications. The development of ethical frameworks and bias mitigation strategies assumes critical importance in these contexts, especially given that institutions operating under resource constraints typically depend on restricted labeled data availability.

Benchmark. Benchmarks are essential for assessing fairness in FMs development, providing teams with systematic evaluation tools. While fairness benchmarks exist for deep learning in medical imaging [46, 168, 176, 177], comprehensive benchmarks and libraries specifically designed for FMs remain lacking [72, 78]. Developing such resources is crucial for holistic fairness evaluation, encompassing metrics, utility, utility-fairness trade-offs, robustness, and data-efficient generalization. Furthermore, FMs should undergo testing on unbiased data and pipelines, with evaluations conducted on distributions that reflect real-world applications to ensure valid and applicable fairness metrics [95]. Benchmarks such as Med-HallMark, which introduces detection methods including MediHallDetector, provide systematic frameworks for evaluating and mitigating hallucinations in medical applications [30].

The Fairness Benchmarking for Medical Imaging Foundation Models (FairMedFM) [72] provides a comprehensive framework for comparing models, techniques, and datasets, covering most of the aspects discussed here. Given the importance of fairness and the lack of protected attributes in many datasets (see Table 2), we encourage both new and existing datasets to collect this information. Even if such metadata is limited to the test set or a reduced patient cohort, it should accurately reflect the dataset distribution to support fairness evaluation and model development, as demonstrated by datasets like Harvard-FairVLMed [98]. Legacy datasets such as PMC-15M [171], which are significant for VLMs research, could be enhanced by extracting demographic metadata such as age and sex from captions. Methods for dataset curation like MetaCLIP [161] can facilitate this process, improving both evaluation and model training. In cases where protected attributes cannot be obtained, we recommend techniques that create groups approximating demographic attributes to facilitate fairness analysis [124, 147].

4.3 Deployment

Documentation. Comprehensive documentation constitutes a fundamental requirement for the ethical implementation of FMs [19, 95]. The documentation must encompass detailed specifications of training data sources, testing benchmark methodologies, and explicit deployment guidelines that minimize risk and bias [110]. Moreover, the documentation should present information in a stratified manner, ensuring accessibility across varying levels of technical expertise while maintaining rigorous detail. This multi-level documentation approach enables stakeholders to comprehend the model's inherent risks, biases, and operational constraints. Such transparency serves as a crucial mechanism for fostering trust and promoting responsible model deployment in clinical settings.

License. Open FMs represent a strategic approach to addressing global bias mitigation, enhancing transparency, and facilitating equitable power distribution [18, 104]. By providing unrestricted access to data, code, and model weights, this approach enables resource-constrained institutions to engage in model development and adaptation, thereby fostering innovation through reduced computational barriers. As established in Section 4.1.2, while closed models present substantial impediments to achieving fairness and impartial outcomes, open-weight architectures facilitate superior downstream task adaptation. Nevertheless, ensuring adherence to intended model applications remains crucial for preventing bias propagation and maintaining equity. The implementation of a Responsible AI License framework provides a structured mechanism for guiding ethical and equitable model utilization.

Monitoring. In medical settings, the prevalence of the disease and the distribution of patients accessing a specific hospital can change over time. While assessing fairness and bias during pretraining and fine-tuning is

critical for FMs, continuous monitoring of deployed models in real-world scenarios is equally important. Despite its importance, such monitoring is not widely practiced; in the United States, only 44% of institutions reported conducting local evaluations for bias [115].

5 Policymakers

This section examines two critical aspects of FMs development: governance frameworks for ensuring reliability and the strategic allocation of essential resources. Figure 3 summarizes the key points discussed in the text, providing a comprehensive overview of foundation model development along with the cost and importance of each step.

5.1 Governance

Reliability. The implementation of FMs in healthcare faces critical reliability challenges due to insufficient standardization in bias and hallucinations detection and mitigation practices [54, 72, 78, 89, 177]. Despite a growing emphasis on AI ethics, healthcare institutions lack both established evaluation frameworks and qualified experts needed to assess these systems effectively. This deficiency is particularly concerning because bias continues to emerge in deployment scenarios, while current mitigation methods demonstrate limited reliability [72, 177]. The challenge of bias mitigation becomes particularly critical in regions with limited data representation, notably in Global South nations (Figure 2), where healthcare systems encounter substantial obstacles in implementing and adapting these models for their populations. This disparity underscores a critical gap in current FMs deployment strategies, as the absence of representative data further compounds existing healthcare inequities.

FMs as institutions. Research demonstrates that algorithmic systems function as institutional frameworks that organize complex machine-human interactions in decision-making processes [5]. FMs represent a significant evolution in this infrastructure, enabling complex decision-making capabilities across medical domains, offering pre-trained architectures adaptable to diverse medical tasks. However, these models systematically perpetuate societal power imbalances and biases through their operational frameworks [36]. Conceptualizing FMs as institutions, rather than purely technical implementations, facilitates the establishment of comprehensive evaluation protocols and governance frameworks that enhance model reliability through systematic oversight, standardized assessment criteria, and robust accountability measures.

Regulation. Legislative frameworks are crucial in enhancing AI system reliability in healthcare. The EU AI Act exemplifies a pioneering approach by classifying medical AI systems as "high-risk" and establishing provisions for General Purpose AI Models (GPAI), which encompass FMs. These models present distinct challenges due to their scale, adaptability, and cross-domain impact [109]. The EU AI Act mandates requirements to detect, prevent, and mitigate bias in applications that could result in discrimination. These regulatory requirements transform FMs development pipelines by necessitating the adoption of fairness-enhancing techniques discussed in this work. Although these changes impose extensive technical documentation requirements and associated costs [2, 25], the adoption of such pipelines enhances the reliability of companies' and government products and enables the scalable deployment of AI from local settings to national and global levels.

5.2 Resource Allocation

Workforce. Fair FMs demand interdisciplinary collaboration among technical, legal, and social experts to ensure comprehensive bias mitigation [15, 86]. These multidisciplinary teams integrate diverse societal values and ethical perspectives into model development and evaluation protocols. Essential to this process is the active participation of marginalized communities, whose insights prove crucial for identifying systemic biases in FMs. The global distribution of FM expertise, however, remains heavily concentrated in the United States and China [6, 179], creating significant barriers to diverse team formation. Another challenge arises from data sources and computing

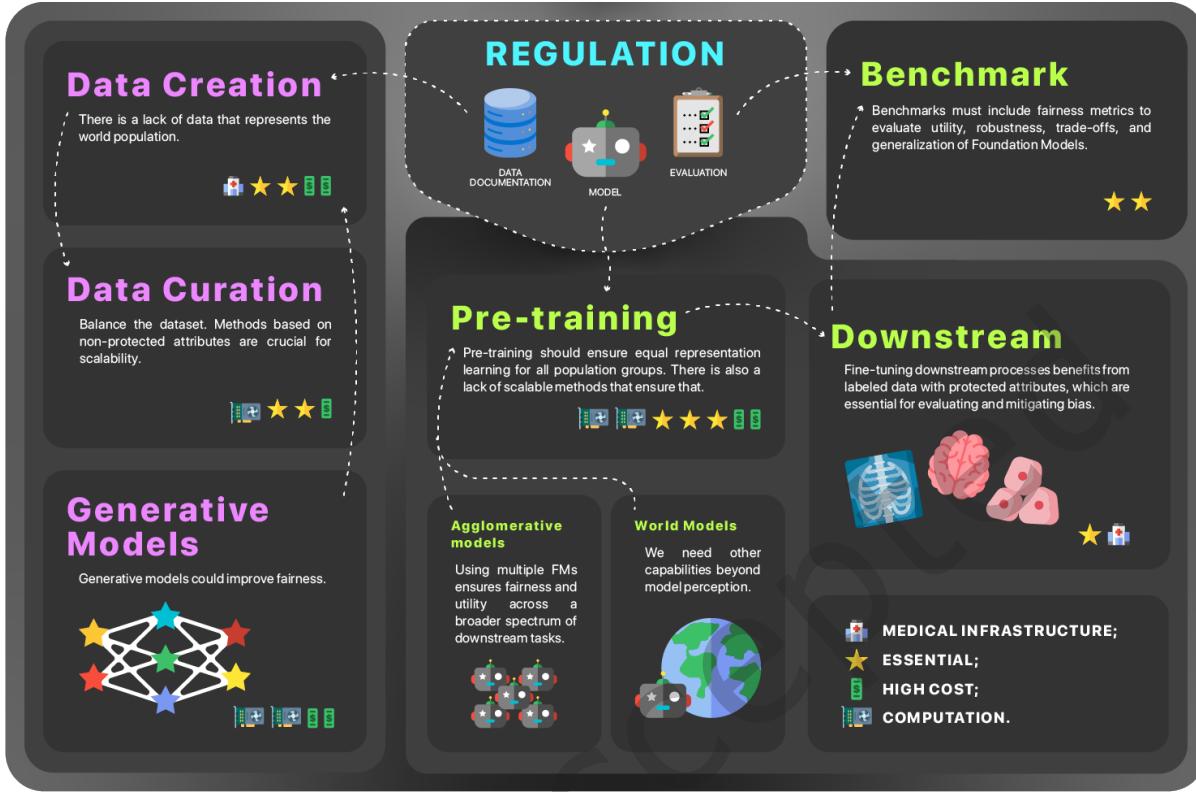


Fig. 3. Overview recommendation: The figure summarizes recommendations for achieving fair foundation models, indicating for each stage its medical infrastructure, importance for fairness (essential), high cost, and computation. The recommendations focus on where governments and industry should prioritize technical interventions to advance fairness. Medical infrastructure is critical for both downstream applications and data creation, but their cost profiles differ: data creation depends primarily on expensive imaging equipment and large-scale data storage, whereas downstream tasks require substantial clinical expertise for data annotation. With respect to fairness mitigation, the most critical and challenging stage is pre-training, which demands extensive computational infrastructure and incurs high costs. The figure also highlights evaluation and data curation as comparatively low-cost yet important strategies for improving fairness. Generative models offer a promising avenue to balance data creation and alleviate infrastructure constraints by better representing the underlying population. At the same time, real-data collection remains more critical for fairness, because equitable access to medical infrastructure across all population groups is essential, whereas generative models can only complement these efforts by improving population representation in the training data. Finally, agglomerative models and world models are identified as future directions with potential to enhance fairness, although they are tightly coupled to pre-training and therefore are not quantitatively assessed in the figure.

resources; many countries face difficulties in providing high-performance computing (HPC) access to institutions. A valuable contribution in this regard is the establishment of training centers that offer education, data, and computing resources.

Data source. AI research disparities manifest critically in data infrastructure and computational resources across global regions. Figure 2 illustrates the significant underrepresentation of Global South populations in major medical imaging datasets. Quantitative analyses reveal systematic exclusion across modalities, with Global

South regions constituting less than 0.7% of text-domain datasets [96]. This limitation extends to linguistic diversity: among medical imaging repositories, only PadChest provides non-English (Spanish) annotations (Table 2). Developing equitable FMs requires datasets that reflect global population diversity across both geographic and demographic dimensions. Achieving representative data collection necessitates integrated technical and governance frameworks that actively incorporate marginalized populations into healthcare data systems [12, 29, 153, 155].

Computing. FMs development requires extensive computational infrastructure, limiting access to organizations with advanced resources [36]. Analysis of the November 2024 Top500 supercomputer rankings demonstrates pronounced regional disparities: North America holds 55.6% of capacity, Europe 27.4%, and Asia 15.9%, while South America and Oceania represent just 0.6% and 0.5% respectively [1]. The combined scarcity of computational power and large-scale datasets creates a fundamental barrier in the Global South, where only Brazil, Australia, and Argentina maintain Top500-listed facilities. Africa's absence from these rankings underscores how infrastructure limitations constrain both FMs development and deployment.

6 Conclusion

FMs represent a transformative advancement in medical imaging analysis, yet their implementation presents both opportunities and challenges for achieving equitable healthcare delivery. Our comprehensive review demonstrates that effective bias mitigation in FMs requires systematic interventions throughout the development pipeline, from data curation to deployment protocols. While technical innovations in training methodologies show promise for enhancing fairness without relying on protected attributes, the substantial computational and data demands of these models risk exacerbating global inequalities. The emergence of regulatory frameworks such as the EU AI Act reflects growing recognition of FMs societal impact and the need for governance structures that ensure responsible development. However, significant disparities persist, particularly in Global South nations where limited access to essential resources including specialized workforce, datasets, and computational infrastructure hinders both development and implementation of fair FMs. Moving forward, addressing these challenges requires coordinated action between technologists, healthcare providers, and policymakers to develop accessible solutions and appropriate frameworks for low-resource countries and institutions. As FMs continue to evolve, their successful implementation in healthcare will depend on our ability to balance technical innovation with ethical principles, ultimately working toward reducing rather than amplifying existing healthcare disparities.

References

- [1] [n. d.]. November 2024 | TOP500. <https://top500.org/lists/top500/2024/11/>.
- [2] Mateo Aboy, Timo Minssen, and Effy Vayena. 2024. Navigating the EU AI Act: Implications for Regulated Digital Medical Products. *npj Digital Medicine* (2024). doi:10.1038/s41746-024-01232-3
- [3] Abejide Ade-Ibijola and Chinedu Okonkwo. 2023. Artificial Intelligence in Africa: Emerging Challenges. In *Responsible AI in Africa: Challenges and Opportunities*, Damian Okaibedi Eke, Kutoma Wakunuma, and Simisola Akintoye (Eds.). Springer International Publishing, Cham. doi:10.1007/978-3-031-08215-3_5
- [4] Ibrahim Alabdulmohsin, Xiao Wang, Andreas Steiner, Priya Goyal, Alexander D'Amour, and Xiaohua Zhai. 2024. CLIP the Bias: How Useful Is Balancing Data in Multimodal Learning? *arXiv:2403.04547* doi:10.48550/arXiv.2403.04547
- [5] Virgilio Almeida, Fernando Filgueiras, and Ricardo Fabrino Mendonca. 2022. Algorithms and Institutions: How Social Sciences Can Contribute to Governance of Algorithms. *IEEE Internet Computing* (2022). doi:10.1109/MIC.2022.3147923
- [6] Bedoor AlShebli, Shahzad Ali Memon, James A. Evans, and Talal Rahwan. 2024. China and the U.S. Produce More Impactful AI Research When Collaborating Together. *Scientific Reports* (2024). doi:10.1038/s41598-024-79863-5
- [7] McKane Andrus and Sarah Villeneuve. 2022. Demographic-Reliant Algorithmic Fairness: Characterizing the Risks of Demographic Data Collection in the Pursuit of Fairness. *arXiv:2205.01038* doi:10.48550/arXiv.2205.01038
- [8] Carolyn Ashurst and Adrian Weller. 2023. Fairness Without Demographic Data: A Survey of Approaches. In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO '23)*. Association for Computing Machinery, New York, NY, USA. doi:10.1145/3617694.3623234

- [9] Mahmoud Assran, Quentin Duval, Ishan Misra, Piotr Bojanowski, Pascal Vincent, Michael Rabbat, Yann LeCun, and Nicolas Ballas. 2023. Self-Supervised Learning from Images with a Joint-Embedding Predictive Architecture. arXiv:2301.08243 [cs] doi:10.48550/arXiv.2301.08243
- [10] Shekoofeh Azizi, Laura Culp, Jan Freyberg, Basil Mustafa, Sebastien Baur, Simon Kornblith, Ting Chen, Nenad Tomasev, Jovana Mitrović, Patricia Strachan, S. Sara Mahdavi, Ellery Wulczyn, Boris Babenko, Megan Walker, Aaron Loh, Po-Hsuan Cameron Chen, Yuan Liu, Pinal Bavishi, Scott Mayer McKinney, Jim Winkens, Abhijit Guha Roy, Zach Beaver, Fiona Ryan, Justin Krogue, Mozziyar Etemadi, Umesh Telang, Yun Liu, Lily Peng, Greg S. Corrado, Dale R. Webster, David Fleet, Geoffrey Hinton, Neil Houlsby, Alan Karthikesalingam, Mohammad Norouzi, and Vivek Natarajan. 2023. Robust and Data-Efficient Generalization of Self-Supervised Machine Learning for Diagnostic Imaging. *Nature Biomedical Engineering* (2023). doi:10.1038/s41551-023-01049-7
- [11] Eugene Bagdasaryan and Vitaly Shmatikov. 2019. Differential Privacy Has Disparate Impact on Model Accuracy. arXiv:1905.12101 [cs] doi:10.48550/arXiv.1905.12101
- [12] Zinzi D Bailey, Nancy Krieger, Madina Agénor, Jasmine Graves, Natalia Linos, and Mary T Bassett. 2017. Structural Racism and Health Inequities in the USA: Evidence and Interventions. *The Lancet* (2017). doi:10.1016/S0140-6736(17)30569-X
- [13] Elisa Bassignana, Amanda Cercas Curry, and Dirk Hovy. 2025. The AI Gap: How Socioeconomic Status Affects Language Technology Interactions. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Vienna, Austria. doi:10.18653/v1/2025.acl-long.914
- [14] Tom L. Beauchamp and James F. Childress. 1994. *Principles of Biomedical Ethics*. Edicoes Loyola.
- [15] Y Bengio. [n. d.]. International AI Safety Report. ([n. d.]).
- [16] Adugnaw Berhane and Fikre Enquselassie. 2015. Patients' Preferences for Attributes Related to Health Care Services at Hospitals in Amhara Region, Northern Ethiopia: A Discrete Choice Experiment. *Patient Preference and Adherence* (2015). doi:10.2147/PPA.S87928
- [17] Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, Erik Brynjolfsson, Shyamal Buch, Dallas Card, Rodrigo Castellon, Niladri Chatterji, Annie Chen, Kathleen Creel, Jared Quincy Davis, Dora Demszky, Chris Donahue, Moussa Doumbouya, Esin Durmus, Stefano Ermon, John Etchemendy, Kawin Ethayarajh, Li Fei-Fei, Chelsea Finn, Trevor Gale, Lauren Gillespie, Karan Goel, Noah Goodman, Shelby Grossman, Neel Guha, Tatsunori Hashimoto, Peter Henderson, John Hewitt, Daniel E. Ho, Jenny Hong, Kyle Hsu, Jing Huang, Thomas Icard, Saahil Jain, Dan Jurafsky, Pratyusha Kalluri, Siddharth Karamcheti, Geoff Keeling, Fereshte Khani, Omar Khattab, Pang Wei Koh, Mark Krass, Ranjay Krishna, Rohith Kuditipudi, Ananya Kumar, Faisal Ladhak, Mina Lee, Tony Lee, Jure Leskovec, Isabelle Levent, Xiang Lisa Li, Xuechen Li, Tengyu Ma, Ali Malik, Christopher D. Manning, Suvin Mirchandani, Eric Mitchell, Zanele Munyikwa, Suraj Nair, Avanika Narayan, Deepak Narayanan, Ben Newman, Allen Nie, Juan Carlos Niebles, Hamed Nilforoshan, Julian Nyarko, Giray O gut, Laurel Orr, Isabel Papadimitriou, Joon Sung Park, Chris Piech, Eva Portelance, Christopher Potts, Aditi Raghuhanathan, Rob Reich, Hongyu Ren, Frieda Rong, Yusuf Roohani, Camilo Ruiz, Jack Ryan, Christopher Ré, Dorsa Sadigh, Shiori Sagawa, Keshav Santhanam, Andy Shih, Krishnan Srinivasan, Alex Tamkin, Rohan Taori, Armin W. Thomas, Florian Tramèr, Rose E. Wang, William Wang, Bohan Wu, Jiajun Wu, Yuhuai Wu, Sang Michael Xie, Michihiro Yasunaga, Jiaxuan You, Matei Zaharia, Michael Zhang, Tianyi Zhang, Xikun Zhang, Yuhui Zhang, Lucia Zheng, Kaitlyn Zhou, and Percy Liang. 2022. On the Opportunities and Risks of Foundation Models. arXiv:2108.07258 [cs] doi:10.48550/arXiv.2108.07258
- [18] Rishi Bommasani, Sayash Kapoor, Kevin Klyman, Shayne Longpre, Ashwin Ramaswami, Daniel Zhang, Marietje Schaake, Daniel E. Ho, Arvind Narayanan, and Percy Liang. 2024. Considerations for Governing Open Foundation Models. *Science* (2024). doi:10.1126/science.adp1848
- [19] Rishi Bommasani, Kevin Klyman, Shayne Longpre, Sayash Kapoor, Nestor Maslej, Betty Xiong, Daniel Zhang, and Percy Liang. 2023. The Foundation Model Transparency Index. arXiv:2310.12941 [cs] doi:10.48550/arXiv.2310.12941
- [20] Brandon M. Booth, Louis Hickman, Shree Krishna Subburaj, Louis Tay, Sang Eun Woo, and Sidney K. D'Mello. 2021. Bias and Fairness in Multimodal Machine Learning: A Case Study of Automated Video Interviews. In *Proceedings of the 2021 International Conference on Multimodal Interaction (ICMI '21)*. Association for Computing Machinery, New York, NY, USA. doi:10.1145/3462244.3479897
- [21] Hanna Borgli, Vajira Thambawita, Pia H. Smedsrød, Steven Hicks, Debesh Jha, Sigrun L. Eskeland, Kristin Ranheim Randel, Konstantin Pogorelov, Mathias Lux, Duc Tien Dang Nguyen, Dag Johansen, Carsten Griwodz, Håkon K. Stensland, Enrique Garcia-Ceja, Peter T. Schmidt, Hugo L. Hammer, Michael A. Riegler, Pål Halvorsen, and Thomas De Lange. 2020. HyperKvasir, a Comprehensive Multi-Class Image and Video Dataset for Gastrointestinal Endoscopy. *Scientific Data* (2020). doi:10.1038/s41597-020-00622-y
- [22] Philippe Burlina, Neil Joshi, William Paul, Katia D. Pacheco, and Neil M. Bressler. 2021. Addressing Artificial Intelligence Bias in Retinal Diagnostics. *Translational Vision Science & Technology* (2021). doi:10.1167/tvst.10.2.13
- [23] Aurelia Bustos, Antonio Pertusa, Jose-Maria Salinas, and Maria de la Iglesia-Vayá. 2020. PadChest: A Large Chest x-Ray Image Dataset with Multi-Label Annotated Reports. *Medical Image Analysis* (2020). doi:10.1016/j.media.2020.101797
- [24] Flavio Calmon, Dennis Wei, Bhanukiran Vinzamuri, Karthikeyan Natesan Ramamurthy, and Kush R Varshney. 2017. Optimized Pre-Processing for Discrimination Prevention. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc.
- [25] Ann-Kathrin Carl and David Hochmann. 2024. Impact of the New European Medical Device Regulation: A Two-Year Comparison. *Biomedical Engineering / Biomedizinische Technik* (2024). doi:10.1515/bmt-2023-0325

[26] L. Elisa Celis and Vijay Keswani. 2019. Improved Adversarial Learning for Fair Classification. *arXiv:1901.10443* [cs] doi:10.48550/arXiv.1901.10443

[27] Delong Chen, Samuel Cahyawijaya, Etsuko Ishii, Ho Shu Chan, Yejin Bang, and Pascale Fung. 2024. What Makes for Good Image Captions? doi:10.48550/ARXIV.2405.00485

[28] Delong Chen, Theo Moutakanni, Willy Chung, Yejin Bang, Ziwei Ji, Allen Bolourchi, and Pascale Fung. 2025. Planning with Reasoning Using Vision Language World Model. doi:10.48550/ARXIV.2509.02722

[29] Irene Y. Chen, Emma Pierson, Sherri Rose, Shalmali Joshi, Kadja Ferryman, and Marzyeh Ghassemi. 2021. Ethical Machine Learning in Healthcare. *Annual Review of Biomedical Data Science* (2021). doi:10.1146/annurev-biodatasci-092820-114757

[30] Jiawei Chen, Dingkang Yang, Tong Wu, Yue Jiang, Xiaolu Hou, Mingcheng Li, Shunli Wang, Dongling Xiao, Ke Li, and Lihua Zhang. 2024. Detecting and Evaluating Medical Hallucinations in Large Vision Language Models. *arXiv:2406.10185* [cs] doi:10.48550/arXiv.2406.10185

[31] Richard J. Chen, Judy J. Wang, Drew F. K. Williamson, Tiffany Y. Chen, Jana Lipkova, Ming Y. Lu, Sharifa Sahai, and Faisal Mahmood. 2023. Algorithmic Fairness in Artificial Intelligence for Medicine and Healthcare. *Nature Biomedical Engineering* (2023). doi:10.1038/s41551-023-01056-8

[32] Shan Chen, Jack Gallifant, Mingye Gao, Pedro Moreira, Nikolaj Munch, Ajay Muthukumar, Arvind Rajan, Jaya Kolluri, Amelia Fiske, Janna Hastings, Hugo Aerts, Brian Anthony, Leo Anthony Celi, William G. La Cava, and Danielle S. Bitterman. 2024. Cross-Care: Assessing the Healthcare Implications of Pre-training Data on Language Model Bias. *arXiv:2405.05506* [cs] doi:10.48550/arXiv.2405.05506

[33] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A Simple Framework for Contrastive Learning of Visual Representations. *arXiv:2002.05709* [cs, stat] doi:10.48550/arXiv.2002.05709

[34] Alexandra Chouldechova, Diana Benavides-Prado, Oleksandr Fialko, and Rhema Vaithianathan. 2018. A Case Study of Algorithm-Assisted Decision Making in Child Maltreatment Hotline Screening Decisions. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*. PMLR.

[35] Gary S Collins, Karel G M Moons, Paula Dhiman, Richard D Riley, Andrew L Beam, Ben Van Calster, Marzyeh Ghassemi, Xiaoxuan Liu, Johannes B Reitsma, Maarten Van Smeden, Anne-Laure Boulesteix, Jennifer Catherine Camaradou, Leo Anthony Celi, Spiros Denaxas, Alastair K Denniston, Ben Glocker, Robert M Golub, Hugh Harvey, Georg Heinze, Michael M Hoffman, André Pascal Kengne, Emily Lam, Naomi Lee, Elizabeth W Loder, Lena Maier-Hein, Bilal A Mateen, Melissa D McCradden, Lauren Oakden-Rayner, Johan Ordish, Richard Parnell, Sherri Rose, Karandeep Singh, Laure Wynants, and Patricia Logullo. 2024. TRIPOD+AI Statement: Updated Guidance for Reporting Clinical Prediction Models That Use Regression or Machine Learning Methods. *BMJ* (2024). doi:10.1136/bmj-2023-078378

[36] Ben Cottier, Robi Rahman, Loredana Fattorini, Nestor Maslej, and David Owen. 2024. The Rising Costs of Training Frontier AI Models. *arXiv:2405.21015* [cs] doi:10.48550/arXiv.2405.21015

[37] Jiequan Cui, Beier Zhu, Xin Wen, Xiaojuan Qi, Bei Yu, and Hanwang Zhang. 2024. Classes Are Not Equal: An Empirical Study on Image Recognition Fairness. *arXiv:2402.18133* [cs] doi:10.48550/arXiv.2402.18133

[38] Maria Correia de Verdier, Rachit Saluja, Louis Gagnon, Dominic LaBella, Ujjwall Baid, Nourel Hoda Tahon, Martha Foltyn-Dumitru, Jikai Zhang, Maram Alafif, Saif Baig, Ken Chang, Gennaro D'Anna, Lisa Deptula, Diviya Gupta, Muhammad Ammar Haider, Ali Hussain, Michael Iv, Marinos Kontzialis, Paul Manning, Farzan Moodi, Teresa Nunes, Aaron Simon, Nico Sollmann, David Vu, Maruf Adewole, Jake Albrecht, Uduanna Anazodo, Rongrong Chai, Verena Chung, Shahriar Faghani, Keyvan Farahani, Anahita Fathi Kazerooni, Eugenio Iglesias, Florian Kofler, Hongwei Li, Marius George Linguraru, Bjoern Menze, Ahmed W. Moawad, Yury Velichko, Benedikt Wiestler, Talissa Altes, Patil Basavasagar, Martin Bendszus, Gianluca Brugnara, Jaeyoung Cho, Yaseen Dhemes, Brandon K. K. Fields, Filip Garrett, Jaime Gass, Lubomir Hadjiiski, Jona Hattangadi-Gluth, Christopher Hess, Jessica L. Houk, Edvin Isufi, Lester J. Layfield, George Mastorakos, John Mongan, Pierre Nedelec, Uyen Nguyen, Sebastian Oliva, Matthew W. Pease, Aditya Rastogi, Jason Sinclair, Robert X. Smith, Leo P. Sugrue, Jonathan Thacker, Igor Vidic, Javier Villanueva-Meyer, Nathan S. White, Mariam Aboian, Gian Marco Conte, Anders Dale, Mert R. Sabuncu, Tyler M. Seibert, Brent Weinberg, Aly Abayazeed, Raymond Huang, Sevcan Turk, Andreas M. Rauschecker, Nikdokht Farid, Philipp Vollmuth, Ayman Nada, Spyridon Bakas, Evan Calabrese, and Jeffrey D. Rudie. 2024. The 2024 Brain Tumor Segmentation (BraTS) Challenge: Glioma Segmentation on Post-treatment MRI. *arXiv:2405.18368* [cs] doi:10.48550/arXiv.2405.18368

[39] Sepehr Dehdashtian, Bashir Sadeghi, and Vishnu Naresh Boddeti. 2024. Utility-Fairness Trade-Offs and How to Find Them. *arXiv:2404.09454* [cs] doi:10.48550/arXiv.2404.09454

[40] Matt Deitke, Christopher Clark, Sangho Lee, Rohun Tripathi, Yue Yang, Jae Sung Park, Mohammadreza Salehi, Niklas Muennighoff, Kyle Lo, Luca Soldaini, Jiasen Lu, Taira Anderson, Erin Bransom, Kiana Ehsani, Huong Ngo, YenSung Chen, Ajay Patel, Mark Yatskar, Chris Callison-Burch, Andrew Head, Rose Hendrix, Favyen Bastani, Eli VanderBilt, Nathan Lambert, Yvonne Chou, Arnavi Chheda, Jenna Sparks, Sam Skjonsberg, Michael Schmitz, Aaron Sarnat, Byron Bischoff, Pete Walsh, Chris Newell, Piper Wolters, Tanmay Gupta, Kuo-Hao Zeng, Jon Borchardt, Dirk Groeneveld, Jen Dumas, Crystal Nam, Sophie Lebrecht, Caitlin Wittif, Carissa Schoenick, Oscar Michel, Ranjay Krishna, Luca Weihs, Noah A. Smith, Hannaneh Hajishirzi, Ross Girshick, Ali Farhadi, and Aniruddha Kembhavi. 2024. Molmo and PixMo: Open Weights and Open Data for State-of-the-Art Multimodal Models. *arXiv:2409.17146* doi:10.48550/arXiv.2409.17146

[41] Mohammad Mahdi Derakhshani, Dheeraj Varghese, Marzieh Fadaee, and Cees G. M. Snoek. 2025. NeoBabel: A Multilingual Open Tower for Visual Generation. doi:10.48550/ARXIV.2507.06137

[42] Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. QLoRA: Efficient Finetuning of Quantized LLMs. arXiv:2305.14314 [cs] doi:10.48550/arXiv.2305.14314

[43] Zhoujie Ding, Ken Ziyu Liu, Pura Peetathawatchai, Berivan Isik, and Sanmi Koyejo. 2024. On Fairness of Low-Rank Adaptation of Large Models. arXiv:2405.17512 [cs] doi:10.48550/arXiv.2405.17512

[44] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. An Image Is Worth 16x16 Words: Transformers for Image Recognition at Scale. arXiv:2010.11929 doi:10.48550/arXiv.2010.11929

[45] Mengnan Du, Fan Yang, Na Zou, and Xia Hu. 2021. Fairness in Deep Learning: A Computational Perspective. *IEEE Intelligent Systems* (2021). doi:10.1109/MIS.2020.3000681

[46] Raman Dutt, Ondrej Bohdal, Sotirios A. Tsaftaris, and Timothy Hospedales. 2024. FairTune: Optimizing Parameter Efficient Fine Tuning for Fairness in Medical Image Analysis. arXiv:2310.05055 [cs] doi:10.48550/arXiv.2310.05055

[47] European Parliament and Council. 2024. Regulation of the European Parliament and of the Council Laying down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act).

[48] Talfan Evans, Nikhil Parthasarathy, Hamza Merzic, and Olivier J. Henaff. 2024. Data Curation via Joint Example Selection Further Accelerates Multimodal Learning. doi:10.48550/ARXIV.2406.17711

[49] Talfan Evans, Shreya Pathak, Hamza Merzic, Jonathan Schwarz, Ryutaro Tanno, and Olivier J. Henaff. 2024. Bad Students Make Great Teachers: Active Learning Accelerates Large-Scale Visual Understanding. arXiv:2312.05328 [cs] doi:10.48550/arXiv.2312.05328

[50] Health Data Research Innovation Gateway. 2024. Moorfields Eye Image BioResource 001.

[51] Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A. Wichmann. 2020. Shortcut Learning in Deep Neural Networks. *Nature Machine Intelligence* (2020). doi:10.1038/s42256-020-00257-z

[52] Zhengyang Geng, Mingyang Deng, Xingjian Bai, J. Zico Kolter, and Kaiming He. 2025. Mean Flows for One-step Generative Modeling. arXiv:2505.13447 [cs] doi:10.48550/arXiv.2505.13447

[53] Ben Glocker, Charles Jones, Mélanie Bernhardt, and Stefan Winzeck. 2023. Algorithmic Encoding of Protected Characteristics in Chest X-ray Disease Detection Models. *eBioMedicine* (2023). doi:10.1016/j.ebiom.2023.104467

[54] Ben Glocker, Charles Jones, Mélanie Roschewitz, and Stefan Winzeck. 2023. Risk of Bias in Chest Radiography Deep Learning Foundation Models. *Radiology: Artificial Intelligence* (2023). doi:10.1148/ryai.2300060

[55] Priya Goyal, Mathilde Caron, Benjamin Lefauveux, Min Xu, Pengchao Wang, Vivek Pai, Mannat Singh, Vitaliy Liptchinsky, Ishan Misra, Armand Joulin, and Piotr Bojanowski. 2021. Self-Supervised Pretraining of Visual Features in the Wild. arXiv:2103.01988 doi:10.48550/arXiv.2103.01988

[56] Priya Goyal, Quentin Duval, Isaac Seessel, Mathilde Caron, Ishan Misra, Levent Sagun, Armand Joulin, and Piotr Bojanowski. 2022. Vision Models Are More Robust And Fair When Pretrained On Uncurated Images Without Supervision. arXiv:2202.08360 doi:10.48550/arXiv.2202.08360

[57] Matthew Groh, Caleb Harris, Luis Soenksen, Felix Lau, Rachel Han, Aerin Kim, Arash Koochek, and Omar Badri. 2021. Evaluating Deep Neural Networks Trained on Clinical Images in Dermatology with the Fitzpatrick 17k Dataset. arXiv:2104.09957 doi:10.48550/arXiv.2104.09957

[58] Zishan Gu, Changchang Yin, Fenglin Liu, and Ping Zhang. 2024. MedVH: Towards Systematic Evaluation of Hallucination for Large Vision Language Models in the Medical Context. arXiv:2407.02730 [cs] doi:10.48550/arXiv.2407.02730

[59] Varun Gulshan, Lily Peng, Marc Coram, Martin C. Stumpe, Derek Wu, Arunachalam Narayanaswamy, Subhashini Venugopalan, Kasumi Widner, Tom Madams, Jorge Cuadros, Ramasamy Kim, Rajiv Raman, Philip C. Nelson, Jessica L. Mega, and Dale R. Webster. 2016. Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs. *JAMA* (2016). doi:10.1001/jama.2016.17216

[60] Melissa Hall, Laura Gustafson, Aaron Adcock, Ishan Misra, and Candace Ross. 2023. Vision-Language Models Performing Zero-Shot Tasks Exhibit Gender-based Disparities. arXiv:2301.11100 doi:10.48550/arXiv.2301.11100

[61] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. 2021. Masked Autoencoders Are Scalable Vision Learners. arXiv:2111.06377 [cs] doi:10.48550/arXiv.2111.06377

[62] Stefan Hegselmann, Shannon Zejiang Shen, Florian Gierse, Monica Agrawal, David Sontag, and Xiaoyi Jiang. 2024. A Data-Centric Approach To Generate Faithful and High Quality Patient Summaries with Large Language Models. arXiv:2402.15422 [cs] doi:10.48550/arXiv.2402.15422

[63] Greg Heinrich, Mike Ranzinger, Hongxu Yin, Yao Lu, Jan Kautz, Andrew Tao, Bryan Catanzaro, and Pavlo Molchanov. 2025. RADIOv2.5: Improved Baselines for Agglomerative Vision Foundation Models. arXiv:2412.07679 [cs] doi:10.48550/arXiv.2412.07679

[64] High-Level Expert Group on Artificial Intelligence. 2019. Ethics Guidelines for Trustworthy AI. <https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>.

[65] Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. LoRA: Low-Rank Adaptation of Large Language Models. arXiv:2106.09685 [cs] doi:10.48550/arXiv.2106.09685

[66] Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. 2025. A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions. *ACM Trans. Inf. Syst.* (2025). doi:10.1145/3703155

[67] Ben Hutchinson, Jason Baldridge, and Vinodkumar Prabhakaran. 2022. Underspecification in Scene Description-to-Depiction Tasks. arXiv:2210.05815 doi:10.48550/arXiv.2210.05815

[68] Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silviana Ciurea-Ilcus, Chris Chute, Henrik Marklund, Behzad Haghgoo, Robyn Ball, Katie Shpanskaya, Jayne Seekins, David A. Mong, Safwan S. Halabi, Jesse K. Sandberg, Ricky Jones, David B. Larson, Curtis P. Langlotz, Bhavik N. Patel, Matthew P. Lungren, and Andrew Y. Ng. 2019. CheXpert: A Large Chest Radiograph Dataset with Uncertainty Labels and Expert Comparison. arXiv:1901.07031 doi:10.48550/arXiv.1901.07031

[69] Ge-Peng Ji, Guobao Xiao, Yu-Cheng Chou, Deng-Ping Fan, Kai Zhao, Geng Chen, and Luc Van Gool. 2022. Video Polyp Segmentation: A Deep Learning Perspective. *Machine Intelligence Research* (2022). doi:10.1007/s11633-022-1371-y

[70] Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of Hallucination in Natural Language Generation. *ACM Comput. Surv.* (2023). doi:10.1145/3571730

[71] Ruinan Jin, Wenlong Deng, Minghui Chen, and Xiaoxiao Li. 2024. Universal Debiased Editing on Foundation Models for Fair Medical Image Classification. *CoRR* (2024).

[72] Ruinan Jin, Zikang Xu, Yuan Zhong, Qingsong Yao, Qi Dou, S. K. Zhou, and Xiaoxiao Li. 2024. FairMedFM: Fairness Benchmarking for Medical Imaging Foundation Models. *Advances in Neural Information Processing Systems* (2024).

[73] Alistair E. W. Johnson, Tom J. Pollard, Nathaniel R. Greenbaum, Matthew P. Lungren, Chih-ying Deng, Yifan Peng, Zhiyong Lu, Roger G. Mark, Seth J. Berkowitz, and Steven Horng. 2019. MIMIC-CXR-JPG, a Large Publicly Available Database of Labeled Chest Radiographs. arXiv:1901.07042 [cs] doi:10.48550/arXiv.1901.07042

[74] Toshihiro Kamishima, Shotaro Akaho, Hideki Asoh, and Jun Sakuma. 2012. Fairness-Aware Classifier with Prejudice Remover Regularizer. In *Machine Learning and Knowledge Discovery in Databases*, Peter A. Flach, Tijl De Bie, and Nello Cristianini (Eds.). Springer, Berlin, Heidelberg. doi:10.1007/978-3-642-33486-3_3

[75] Alexandros Karargyris, Renato Umeton, Micah J. Sheller, Alejandro Aristizabal, John George, Anna Wuest, Sarthak Pati, Hasan Kassem, Maximilian Zenk, Ujjwal Baid, Prakash Narayana Moorthy, Alexander Chowdhury, Junyi Guo, Sahil Nalawade, Jacob Rosenthal, David Kanter, Maria Xenocristou, Daniel J. Beutel, Verena Chung, Timothy Bergquist, James Eddy, Abubakar Abid, Lewis Tunstall, Omar Sanseviero, Dimitrios Dimitriadis, Yiming Qian, Xinxing Xu, Yong Liu, Rick Siow Mong Goh, Srinivas Bala, Victor Bittorf, Sreekar Reddy Puchala, Biagio Ricciuti, Soujanya Samineni, Eshna Sengupta, Akshay Chaudhari, Cody Coleman, Bala Desinghu, Gregory Diamos, Debo Dutta, Diane Feddema, Grigori Fursin, Xinyuan Huang, Satyananda Kashyap, Nicholas Lane, Indranil Mallick, Pietro Mascagni, Virendra Mehta, Cassiano Ferro Moraes, Vivek Natarajan, Nikola Nikolov, Nicolas Padov, Gennady Pekhimenko, Vijay Janapa Reddi, G. Anthony Reina, Pablo Ribalta, Abhishek Singh, Jayaraman J. Thiagarajan, Jacob Albrecht, Thomas Wolf, Geralyn Miller, Huazhu Fu, Prashant Shah, Daguang Xu, Poonam Yadav, David Talby, Mark M. Awad, Jeremy P. Howard, Michael Rosenthal, Luigi Marchionni, Massimo Loda, Jason M. Johnson, Spyridon Bakas, and Peter Mattson. 2023. Federated Benchmarking of Medical Artificial Intelligence with MedPerf. *Nature Machine Intelligence* (2023). doi:10.1038/s42256-023-00652-2

[76] Masakata Kawai, Noriaki Ota, and Shinsuke Yamaoka. 2023. Large-Scale Pretraining on Pathological Images for Fine-Tuning of Small Pathological Benchmarks. In *Medical Image Learning with Limited and Noisy Data*, Zhiyun Xue, Sameer Antani, Ghada Zamzmi, Feng Yang, Sivaramakrishnan Rajaraman, Sharon Xiaolei Huang, Marius George Linguraru, and Zhaohui Liang (Eds.). Springer Nature Switzerland, Cham. doi:10.1007/978-3-031-44917-8_25

[77] Daniel S. Kermany, Michael Goldbaum, Wenjia Cai, Carolina C. S. Valentim, Huiying Liang, Sally L. Baxter, Alex McKeown, Ge Yang, Xiaokang Wu, Fangbing Yan, Justin Dong, Made K. Prasadha, Jacqueline Pei, Magdalene Y. L. Ting, Jie Zhu, Christina Li, Sierra Hewett, Jason Dong, Ian Ziyar, Alexander Shi, Runze Zhang, Lianghong Zheng, Rui Hou, William Shi, Xin Fu, Yaou Duan, Viet A. N. Huu, Cindy Wen, Edward D. Zhang, Charlotte L. Zhang, Oulan Li, Xiaobo Wang, Michael A. Singer, Xiaodong Sun, Jie Xu, Ali Tafreshi, M. Anthony Lewis, Huimin Xia, and Kang Zhang. 2018. Identifying Medical Diagnoses and Treatable Diseases by Image-Based Deep Learning. *Cell* (2018). doi:10.1016/j.cell.2018.02.010

[78] Muhammad Osama Khan, Muhammad Muneeb Afzal, Shujaat Mirza, and Yi Fang. 2023. How Fair Are Medical Imaging Foundation Models?. In *Proceedings of the 3rd Machine Learning for Health Symposium*. PMLR.

[79] Bardia Khosravi, Frank Li, Theo Dapamede, Pouria Rouzrok, Cooper U. Gamble, Hari M. Trivedi, Cody C. Wyles, Andrew B. Sellergren, Saptarshi Purkayastha, Bradley J. Erickson, and Judy W. Gichoya. 2024. Synthetically Enhanced: Unveiling Synthetic Data's Potential in Medical Imaging Research. *eBioMedicine* (2024). doi:10.1016/j.ebiom.2024.105174

[80] Yubin Kim, Hyewon Jeong, Shan Chen, Shuyue Stella Li, Chanwoo Park, Mingyu Lu, Kumail Alhamoud, Jimin Mun, Cristina Grau, Minseok Jung, Rodrigo Gameiro, Lizhou Fan, Eugene Park, Tristan Lin, Joonsil Yoon, Wonjin Yoon, Maarten Sap, Yulia Tsvetkov, Paul Liang, Xuhai Xu, Xin Liu, Chunjong Park, Hyeonhoon Lee, Hae Won Park, Daniel McDuff, Samir Tulebaev, and Cynthia Breazeal. 2025. Medical Hallucinations in Foundation Models and Their Impact on Healthcare. arXiv:2503.05777 [cs] doi:10.48550/arXiv.2503.05777

[81] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. 2023. Segment Anything. arXiv:2304.02643 [cs] doi:10.48550/arXiv.2304.02643

[82] Ira Ktena, Olivia Wiles, Isabela Albuquerque, Sylvestre-Alvise Rebuffi, Ryutaro Tanno, Abhijit Guha Roy, Shekoofeh Azizi, Danielle Belgrave, Pushmeet Kohli, Alan Karthikesalingam, Taylan Cemgil, and Sven Gowal. 2023. Generative Models Improve Fairness of Medical Classifiers under Distribution Shifts. arXiv:2304.09218 [cs] doi:10.48550/arXiv.2304.09218

[83] Andrew Kyle Lampinen, Stephanie C. Y. Chan, and Katherine Hermann. 2024. Learned Feature Representations Are Biased by Complexity, Learning Order, Position, and More. doi:10.48550/ARXIV.2405.05847

[84] Yann LeCun. [n. d.]. A Path Towards Autonomous Machine Intelligence. ([n. d.].)

[85] Katherine Lee, Daphne Ippolito, Andrew Nystrom, Chiyuan Zhang, Douglas Eck, Chris Callison-Burch, and Nicholas Carlini. 2022. Deduplicating Training Data Makes Language Models Better. arXiv:2107.06499 doi:10.48550/arXiv.2107.06499

[86] Karim Lekadir, Alejandro F. Frangi, Antonio R. Porras, Ben Glocker, Celia Cintas, Curtis P. Langlotz, Eva Weicken, Folkert W. Asselbergs, Fred Prior, Gary S. Collins, Georgios Kaassis, Gianna Tsakou, Irène Buvat, Jayashree Kalpathy-Cramer, John Mongan, Julia A. Schnabel, Kaisar Kushiabar, Katrine Riklund, Kostas Marias, Lameck M. Amugongo, Lauren A. Fromont, Lena Maier-Hein, Leonor Cerdá-Alberich, Luis Martí-Bonmatí, M. Jorge Cardoso, Maciej Bobowicz, Mahsa Shabani, Manolis Tsiknakis, Maria A. Zuluaga, Marie-Christine Fritzsche, Marina Camacho, Marius George Linguraru, Markus Wenzel, Marleen De Bruijne, Martin G. Tolsgaard, Melanie Goisauf, Mónica Cano Abadía, Nikolaos Papanikolaou, Noussair Lazrak, Oriol Pujol, Richard Osuala, Sandy Napel, Sara Colantonio, Smriti Joshi, Stefan Klein, Susanna Aussó, Wendy A. Rogers, Zohaib Salahuddin, and Martijn P. A. Starmans. 2025. FUTURE-AI: International Consensus Guideline for Trustworthy and Deployable Artificial Intelligence in Healthcare. *BMJ* (2025). doi:10.1136/bmj-2024-081554

[87] Chunyuan Li, Cliff Wong, Sheng Zhang, Naoto Usuyama, Haotian Liu, Jianwei Yang, Tristan Naumann, Hoifung Poon, and Jianfeng Gao. 2023. LLaVA-Med: Training a Large Language-and-Vision Assistant for Biomedicine in One Day. doi:10.48550/ARXIV.2306.00890

[88] Qing Li, Zhihang Hu, Yixuan Wang, Lei Li, Yimin Fan, Irwin King, Gengjie Jia, Sheng Wang, Le Song, and Yu Li. 2024. Progress and Opportunities of Foundation Models in Bioinformatics. *Briefings in Bioinformatics* (2024). doi:10.1093/bib/bbae548

[89] Qing Li, Yizhe Zhang, Yan Li, Jun Lyu, Meng Liu, Longyu Sun, Mengting Sun, Qirong Li, Wenyue Mao, Xinran Wu, Yajing Zhang, Yinghua Chu, Shuo Wang, and Chengyan Wang. 2024. An Empirical Study on the Fairness of Foundation Models for Multi-Organ Image Segmentation. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2024*, Marius George Linguraru, Qi Dou, Aasa Feragen, Stamatia Giannarou, Ben Glocker, Karim Lekadir, and Julia A. Schnabel (Eds.). Springer Nature Switzerland, Cham. doi:10.1007/978-3-031-72390-2_41

[90] Weixiong Lin, Ziheng Zhao, Xiaoman Zhang, Chaoyi Wu, Ya Zhang, Yanfeng Wang, and Weidi Xie. 2023. PMC-CLIP: Contrastive Language-Image Pre-training Using Biomedical Documents. arXiv:2303.07240 [cs] doi:10.48550/arXiv.2303.07240

[91] Etaı Littwin, Omid Saremi, Madhu Advani, Vimal Thilak, Preetum Nakkirian, Chen Huang, and Joshua M. Susskind. 2024. How JEPAs Avoids Noisy Features: The Implicit Bias of Deep Linear Self Distillation Networks. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.

[92] Hanchao Liu, Wenyuan Xue, Yifei Chen, Dapeng Chen, Xiutian Zhao, Ke Wang, Liping Hou, Rongjun Li, and Wei Peng. 2024. A Survey on Hallucination in Large Vision-Language Models. arXiv:2402.00253 [cs] doi:10.48550/arXiv.2402.00253

[93] Xiaoxuan Liu, Samantha Cruz Rivera, David Moher, Melanie J Calvert, and Alastair K Denniston. 2020. Reporting Guidelines for Clinical Trial Reports for Interventions Involving Artificial Intelligence: The CONSORT-AI Extension. *BMJ* (2020). doi:10.1136/bmj.m3164

[94] Zelong Liu, Alexander Zhou, Arnold Yang, Alara Yilmaz, Maxwell Yoo, Mikey Sullivan, Catherine Zhang, James Grant, Daiqing Li, Zahi A. Fayad, Sean Huver, Timothy Deyer, and Xueyan Mei. 2023. RadImageGAN – A Multi-modal Dataset-Scale Generative AI for Medical Imaging. arXiv:2312.05953 doi:10.48550/arXiv.2312.05953

[95] Shayne Longpre, Stella Biderman, Alon Albalak, Hailey Schoelkopf, Daniel McDuff, Sayash Kapoor, Kevin Klyman, Kyle Lo, Gabriel Ilharco, Nay San, Maribeth Rauh, Aviya Skowron, Bertie Vidgen, Laura Weidinger, Arvind Narayanan, Victor Sanh, David Adelani, Percy Liang, Rishi Bommasani, Peter Henderson, Sasha Luccioni, Yacine Jernite, and Luca Soldaini. 2024. The Responsible Foundation Model Development Cheatsheet: A Review of Tools & Resources. arXiv:2406.16746 [cs]

[96] Shayne Longpre, Nikhil Singh, Manuel Cherep, Kushagra Tiwary, Joanna Materzynska, William Brannon, Robert Mahari, Manan Dey, Mohammed Hamdy, Nayan Saxena, Ahmad Mustafa Anis, Emad A. Alghamdi, Vu Minh Chien, Naana Obeng-Marnu, Da Yin, Kun Qian, Yizhi Li, Minnie Liang, An Dinh, Shrestha Mohanty, Deividas Mataciunas, Tobin South, Jianguo Zhang, Ariel N. Lee, Campbell S. Lund, Christopher Klamm, Damien Sileo, Dignanta Misra, Enrico Shippole, Kevin Klyman, Lester JV Miranda, Niklas Muenninghoff, Seonghyeon Ye, Seungone Kim, Vipul Gupta, Vivek Sharma, Xuhui Zhou, Caiming Xiong, Luis Villa, Stella Biderman, Alex Pentland, Sara Hooker, and Jad Kabbara. 2024. Bridging the Data Provenance Gap Across Text, Speech and Video. arXiv:2412.17847 [cs] doi:10.48550/arXiv.2412.17847

[97] Ming Y. Lu, Bowen Chen, Drew F. K. Williamson, Richard J. Chen, Ivy Liang, Tong Ding, Guillaume Jaume, Igor Odintsov, Long Phi Le, Georg Gerber, Anil V. Parwani, Andrew Zhang, and Faisal Mahmood. 2024. A Visual-Language Foundation Model for Computational Pathology. *Nature Medicine* (2024). doi:10.1038/s41591-024-02856-4

[98] Yan Luo, Min Shi, Muhammad Osama Khan, Muhammad Muneeb Afzal, Hao Huang, Shuaihang Yuan, Yu Tian, Luo Song, Ava Kouhana, Tobias Elze, Yi Fang, and Mengyu Wang. 2024. FairCLIP: Harnessing Fairness in Vision-Language Learning. arXiv:2403.19949 [cs] doi:10.48550/arXiv.2403.19949

[99] Jun Ma, Yuting He, Feifei Li, Lin Han, Chenyu You, and Bo Wang. 2024. Segment Anything in Medical Images. *Nature Communications* (2024). arXiv:2304.12306 [eess] doi:10.1038/s41467-024-44824-z

[100] Xinsong Ma, Zekai Wang, and Weiwei Liu. 2022. On the Tradeoff Between Robustness and Fairness. *Advances in Neural Information Processing Systems* (2022).

[101] Yiting Ma, Xuejin Chen, Kai Cheng, Yang Li, and Bin Sun. 2021. LDPolypVideo Benchmark: A Large-Scale Colonoscopy Video Dataset of Diverse Polyps. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2021*, Marleen de Bruijne, Philippe C. Cattin, Stéphanie Cotin, Nicolas Padoy, Stefanie Speidel, Yefeng Zheng, and Caroline Essert (Eds.). Springer International Publishing, Cham. doi:10.1007/978-3-030-87240-3_37

[102] Potsawee Manakul, Adian Liusie, and Mark J. F. Gales. 2023. SelfCheckGPT: Zero-Resource Black-Box Hallucination Detection for Generative Large Language Models. arXiv:2303.08896 [cs] doi:10.48550/arXiv.2303.08896

[103] Debmalya Mandal, Samuel Deng, Suman Jana, Jeannette Wing, and Daniel J Hsu. 2020. Ensuring Fairness Beyond the Training Data. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc.

[104] Michael E. Matheny, Jennifer C. Goldsack, Suchi Saria, Nigam H. Shah, Jacqueline Gerhart, I. Glenn Cohen, W. Nicholson Price, Bakul Patel, Philip R. O. Payne, Peter J. Embi, Brian Anderson, and Eric Horvitz. 2025. Artificial Intelligence In Health And Health Care: Priorities For Action. *Health Affairs* (2025). doi:10.1377/hlthaff.2024.01003

[105] Melissa D. McCradden, Shalmali Joshi, Mjaye Mazwi, and James A. Anderson. 2020. Ethical Limitations of Algorithmic Fairness Solutions in Health Care Machine Learning. *The Lancet Digital Health* (2020). doi:10.1016/S2589-7500(20)30065-0

[106] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2021. A Survey on Bias and Fairness in Machine Learning. *Comput. Surveys* (2021). doi:10.1145/3457607

[107] Xueyan Mei, Zelong Liu, Philip M. Robson, Brett Marinelli, Mingqian Huang, Amish Doshi, Adam Jacobi, Chendi Cao, Katherine E. Link, Thomas Yang, Ying Wang, Hayit Greenspan, Timothy Deyer, Zahi A. Fayad, and Yang Yang. 2022. RadImageNet: An Open Radiologic Deep Learning Research Dataset for Effective Transfer Learning. *Radiology: Artificial Intelligence* (2022). doi:10.1148/ryai.210315

[108] Sören Mindermann, Jan M. Brauner, Muhammed T. Razzak, Mrinank Sharma, Andreas Kirsch, Winnie Xu, Benedikt Höltgen, Aidan N. Gomez, Adrien Morisot, Sebastian Farquhar, and Yarin Gal. 2022. Prioritized Training on Points That Are Learnable, Worth Learning, and Not yet Learnt. In *Proceedings of the 39th International Conference on Machine Learning*. PMLR.

[109] Timo Minssen, Effy Vayena, and I. Glenn Cohen. 2023. The Challenges for Regulating Medical Use of ChatGPT and Other Large Language Models. *JAMA* (2023). doi:10.1001/jama.2023.9651

[110] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. Model Cards for Model Reporting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT* '19)*. Association for Computing Machinery, New York, NY, USA. doi:10.1145/3287560.3287596

[111] Michael Moor, Oishi Banerjee, Zahra Shakeri Hossein Abad, Harlan M. Krumholz, Jure Leskovec, Eric J. Topol, and Pranav Rajpurkar. 2023. Foundation Models for Generalist Medical Artificial Intelligence. *Nature* (2023). doi:10.1038/s41586-023-05881-4

[112] Stefania L. Moroianu, Christian Blauthgen, Pierre Chambon, Mehdi Cherti, Jean-Benoit Delbrouck, Magdalini Paschali, Brandon Price, Judy Gichoya, Jenia Jitsev, Curtis P. Langlotz, and Akshay S. Chaudhari. 2025. Improving Performance, Robustness, and Fairness of Radiographic AI Models with Finely-Controllable Synthetic Data. arXiv:2508.16783 [cs] doi:10.48550/arXiv.2508.16783

[113] Ningping Mou, Xinli Yue, Lingchen Zhao, and Qian Wang. 2024. Fairness Is Essential for Robustness: Fair Adversarial Training by Identifying and Augmenting Hard Examples. *Frontiers of Computer Science* (2024). doi:10.1007/s11704-024-3587-1

[114] Luis Filipe Nakayama, David Restrepo, João Matos, Lucas Zago Ribeiro, Fernando Korn Malerbi, Leo Anthony Celi, and Caio Saito Regatieri. 2024. BRSET: A Brazilian Multilabel Ophthalmological Dataset of Retina Fundus Photos. *PLOS Digital Health* (2024). doi:10.1371/journal.pdig.0000454

[115] Paige Nong, Julia Adler-Milstein, Nate C. Apathy, A. Jay Holmgren, and Jordan Everson. 2025. Current Use And Evaluation Of Artificial Intelligence And Predictive Models In US Hospitals. *Health Affairs* (2025). doi:10.1377/hlthaff.2024.00842

[116] Peter A. Noseworthy, Zachi I. Attia, LaPrincess C. Brewer, Sharonne N. Hayes, Xiaoxi Yao, Suraj Kapa, Paul A. Friedman, and Francisco Lopez-Jimenez. 2020. Assessing and Mitigating Bias in Medical Artificial Intelligence. *Circulation: Arrhythmia and Electrophysiology* (2020). doi:10.1161/CIRCEP.119.007988

[117] Ziad Obermeyer, Rebecca Nissan, Michael Stern, Stephanie Eanoff, Emily Joy Bembeneck, and Sendhil Mullainathan. [n. d.]. Algorithmic Bias Playbook. ([n. d.].)

[118] Maxime Oquab, Timothée Dariset, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Mahmoud Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Hervé Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. 2024. DINov2: Learning Robust Visual Features without Supervision. arXiv:2304.07193 doi:10.48550/arXiv.2304.07193

[119] Maxime Oquab, Timothée Dariset, Théo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Mido Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Hervé Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. 2023. DINov2: Learning Robust Visual Features without Supervision. *Transactions on Machine Learning Research*

(2023).

[120] Gökhan Özbulak, Oscar Jimenez-del-Toro, Maíra Fatoretto, Lilian Berton, and André Anjos. 2025. A Multi-Objective Evaluation Framework for Analyzing Utility-Fairness Trade-Offs in Machine Learning Systems. arXiv:2503.11120 [cs] doi:10.48550/arXiv.2503.11120

[121] Mehran Pesteie, Purang Abolmaesumi, and Robert N. Rohling. 2019. Adaptive Augmentation of Medical Data Using Independently Conditional Variational Auto-Encoders. *IEEE Transactions on Medical Imaging* (2019). doi:10.1109/TMI.2019.2914656

[122] Walter H. L. Pinaya, Petru-Daniel Tudosi, Jessica Dafflon, Pedro F. da Costa, Virginia Fernandez, Parashkev Nachev, Sébastien Ourselin, and M. Jorge Cardoso. 2022. Brain Imaging Generation with Latent Diffusion Models. arXiv:2209.07162 doi:10.48550/arXiv.2209.07162

[123] Geoff Pleiss, Manish Raghavan, Felix Wu, Jon Kleinberg, and Kilian Q Weinberger. 2017. On Fairness and Calibration. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc.

[124] Dilermando Queiroz, André Anjos, and Lilian Berton. 2025. Using Backbone Foundation Model for Evaluating Fairness in Chest Radiography Without Demographic Data. In *Ethics and Fairness in Medical Imaging*, Esther Puyol-Antón, Ghada Zamzmi, Aasa Feragen, Andrew P. King, Veronika Cheplygina, Melanie Ganz-Benjaminsen, Enzo Ferrante, Ben Glocker, Eike Petersen, John S. H. Baxter, Islem Rekik, and Roy Eagleson (Eds.). Springer Nature Switzerland, Cham. doi:10.1007/978-3-031-72787-0_11

[125] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning Transferable Visual Models From Natural Language Supervision. arXiv:2103.00020

[126] Alvin Rajkomar, Michaela Hardt, Michael D. Howell, Greg Corrado, and Marshall H. Chin. 2018. Ensuring Fairness in Machine Learning to Advance Health Equity. *Annals of Internal Medicine* (2018). doi:10.7326/M18-1990

[127] Pranav Rajpurkar, Jeremy Irvin, Aarti Bagul, Daisy Ding, Tony Duan, Hershel Mehta, Brandon Yang, Kaylie Zhu, Dillon Laird, Robyn L. Ball, Curtis Langlotz, Katie Shpanskaya, Matthew P. Lungren, and Andrew Y. Ng. 2018. MURA: Large Dataset for Abnormality Detection in Musculoskeletal Radiographs. arXiv:1712.06957 [physics] doi:10.48550/arXiv.1712.06957

[128] Eduardo P. Reis, Joselisa P. Q. de Paiva, Maria C. B. da Silva, Guilherme A. S. Ribeiro, Victor F. Paiva, Lucas Bulgarelli, Henrique M. H. Lee, Paulo V. Santos, Vanessa M. Brito, Lucas T. W. Amaral, Gabriel L. Beraldo, Jorge N. Haidar Filho, Gustavo B. S. Teles, Gilberto Szarf, Tom Pollard, Alistair E. W. Johnson, Leo A. Celi, and Edson Amaro. 2022. BRAX, Brazilian Labeled Chest x-Ray Dataset. *Scientific Data* (2022). doi:10.1038/s41597-022-01608-8

[129] Eduardo Pontes Reis, Felipe Nascimento, Mateus Aranha, Fernando Mainetti Secol, Birajara Machado, Marcelo Felix, Anouk Stein, and Edson Amaro. [n. d.]. Brain Hemorrhage Extended (BHX): Bounding Box Extrapolation from Thick to Thin Slice CT Images. doi:10.13026/9CFT-HG92

[130] María Agustina Ricci Lara, Rodrigo Echeveste, and Enzo Ferrante. 2022. Addressing Fairness in Artificial Intelligence for Medical Imaging. *Nature Communications* (2022). doi:10.1038/s41467-022-32186-3

[131] Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. 2015. Prioritized Experience Replay. doi:10.48550/ARXIV.1511.05952

[132] Jessica Schrouff, Natalie Harris, Oluwasanmi Koyejo, Ibrahim Alabdulmohsin, Eva Schnider, Krista Opsahl-Ong, Alex Brown, Subhrajit Roy, Diana Mincu, Christina Chen, Awa Dieng, Yuan Liu, Vivek Natarajan, Alan Karthikesalingam, Katherine Heller, Silvia Chiappa, and Alexander D'Amour. 2023. Diagnosing Failures of Fairness Transfer across Distribution Shift in Real-World Medical Settings. arXiv:2202.01034 [cs, stat] doi:10.48550/arXiv.2202.01034

[133] Andrew Sellergren, Sahar Kazemzadeh, Tiam Jaroensri, Atilla Kiraly, Madeleine Traverse, Timo Kohlberger, Shawn Xu, Fayaz Jamil, Cian Hughes, Charles Lau, Justin Chen, Fereshteh Mahvar, Liron Yatziv, Tiffany Chen, Bram Sterling, Stefanie Anna Baby, Susanna Maria Baby, Jeremy Lai, Samuel Schmidgall, Lu Yang, Kejia Chen, Per Björnsson, Shashir Reddy, Ryan Brush, Kenneth Philbrick, Mercy Asiedu, Ines Mezherreg, Howard Hu, Howard Yang, Richa Tiwari, Sunny Jansen, Preeti Singh, Yun Liu, Shekoofeh Azizi, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Riviere, Louis Rouillard, Thomas Mesnard, Geoffrey Cideron, Jean-bastien Grill, Sabela Ramos, Edouard Yvinec, Michelle Casbon, Elena Buchatskaya, Jean-Baptiste Alayrac, Dmitry Lepikhin, Vlad Feinberg, Sébastien Borgeaud, Alek Andreev, Cassidy Hardin, Robert Dadashi, Léonard Hussenot, Armand Joulin, Olivier Bachem, Yossi Matias, Katherine Chou, Avinatan Hassidim, Kavi Goel, Clement Farabet, Joelle Barral, Tris Warkentin, Jonathon Shlens, David Fleet, Victor Cotrata, Omar Sanseviero, Gus Martins, Phoebe Kirk, Anand Rao, Shravya Shetty, David F. Steiner, Can Kirmizibayrak, Rory Pilgrim, Daniel Golden, and Lin Yang. 2025. MedGemma Technical Report. arXiv:2507.05201 [cs] doi:10.48550/arXiv.2507.05201

[134] Congzhen Shi, Ryan Rezai, Jiaxi Yang, Qi Dou, and Xiaoxiao Li. 2024. A Survey on Trustworthiness in Foundation Models for Medical Image Analysis. arXiv:2407.15851 [cs] doi:10.48550/arXiv.2407.15851

[135] Adi Simhi, Itay Itzhak, Fazl Barez, Gabriel Stanovsky, and Yonatan Belinkov. 2025. Trust Me, I'm Wrong: LLMs Hallucinate with Certainty Despite Knowing the Answer. arXiv:2502.12964 [cs] doi:10.48550/arXiv.2502.12964

[136] Yang Song and Prafulla Dhariwal. 2023. Improved Techniques for Training Consistency Models. arXiv:2310.14189 [cs] doi:10.48550/arXiv.2310.14189

[137] Emma A.M. Stanley, Raissa Souza, Matthias Wilms, and Nils D. Forkert. 2025. Where, Why, and How Is Bias Learned in Medical Image Analysis Models? A Study of Bias Encoding within Convolutional Networks Using Synthetic Data. *eBioMedicine* (2025). doi:10.1016/j.ebiom.2024.105501

[138] Nima Tajbakhsh, Suryakanth R. Gurudu, and Jianming Liang. 2016. Automated Polyp Detection in Colonoscopy Videos Using Shape and Context Information. *IEEE Transactions on Medical Imaging* (2016). doi:10.1109/TMI.2015.2487997

[139] L. C. M. team, Loic Barrault, Paul-Ambroise Duquenne, Maha Elbayad, Artyom Kozhevnikov, Belen Alastrauey, Pierre Andrews, Mariano Coria, Guillaume Couairon, Marta R. Costa-jussà, David Dale, Hady Elsahar, Kevin Heffernan, João Maria Janeiro, Tuan Tran, Christophe Ropers, Eduardo Sánchez, Robin San Roman, Alexandre Mourachko, Safiyyah Saleem, and Holger Schwenk. 2024. Large Concept Models: Language Modeling in a Sentence Representation Space. arXiv:2412.08821 [cs] doi:10.48550/arXiv.2412.08821

[140] Ali S. Tejani, Michail E. Klontzas, Anthony A. Gatti, John T. Mongan, Linda Moy, Seong Ho Park, Charles E. Kahn, for the CLAIM 2024 Update Panel, Sunhy Abbara, Saif Afat, Uduanna C. Anazodo, Anna Andreychenko, Folkert W. Asselbergs, Aldo Badano, Bettina Baessler, Bayarbaatar Bold, Sotirios Bisdas, Torkel B. Brismar, Giovanni E. Cacciamani, John A. Carrino, Julius Chapiro, Michael F. Chiang, Tessa S. Cook, Renato Cuocolo, John Damilakis, Roxana Daneshjou, Carlo N. De Cecco, Hesham Elhalawani, Guillermo Elizondo-Riojas, Andrey Fedorov, Benjamin Fine, Adam E. Flanders, Judy Wawira Gichoya, Maryellen L. Giger, Sufwan S. Halabi, Sven Haller, William Hsu, Krishna Juluru, Jayashree Kalpathy-Cramer, Apostolos H. Karantanis, Felipe C. Kitamura, Burak Kocak, Dow-Mu Koh, Elmar Kotter, Elizabeth A. Krupinski, Curtis P. Langlotz, Cecilia S. Lee, Mario Maas, Anant Madabhushi, Lena Maier-Hein, Kostas Marias, Luis Martí-Bonmatí, Jaishree Naidoo, Emanuele Neri, Robert Ochs, Nikolaos Papanikolaou, Thomas Papathomas, Katja Pinker-Domenig, Daniel Pinto Dos Santos, Fred Prior, Alexandros Protonotarios, Mauricio Reyes, Veronica Rotemberg, Jeffrey D. Rudie, Emmanuel Salinas-Miranda, Francesco Sardanelli, Mark E. Schweitzer, Luca Maria Sconfienza, Ronnie Sebro, Prateek Sharma, An Tang, Antonios Tzortzakakis, Jeroen Van Der Laak, Peter M. A. Van Ooijen, Vasantha K. Venugopal, Jacob J. Visser, Bradford J. Wood, Carol C. Wu, Greg Zaharchuk, and Marc Zins. 2024. Checklist for Artificial Intelligence in Medical Imaging (CLAIM): 2024 Update. *Radiology: Artificial Intelligence* (2024). doi:10.1148/ryai.240300

[141] Cuong Tran, Ferdinandio Fioretto, Jung-Eun Kim, and Rakshit Naidu. 2022. Pruning Has a Disparate Impact on Model Accuracy. arXiv:2205.13574 [cs] doi:10.48550/arXiv.2205.13574

[142] Michael Tschannen, Alexey Gritsenko, Xiao Wang, Muhammad Ferjad Naeem, Ibrahim Alabdulmohsin, Nikhil Parthasarathy, Talfan Evans, Lucas Beyer, Ye Xia, Basil Mustafa, Olivier Hénaff, Jeremiah Harmsen, Andreas Steiner, and Xiaohua Zhai. 2025. SigLIP 2: Multilingual Vision-Language Encoders with Improved Semantic Understanding, Localization, and Dense Features. arXiv:2502.14786 [cs] doi:10.48550/arXiv.2502.14786

[143] Tao Tu, Shekoofeh Azizi, Danny Driess, Mike Schaeckermann, Mohamed Amin, Pi-Chuan Chang, Andrew Carroll, Chuck Lau, Ryutaro Tanno, Ira Ktena, Basil Mustafa, Aakanksha Chowdhery, Yun Liu, Simon Kornblith, David Fleet, Philip Mansfield, Sushant Prakash, Renee Wong, Sunny Virmani, Christopher Semturs, S. Sara Mahdavi, Bradley Green, Ewa Dominowska, Blaise Aguera y Arcas, Joelle Barral, Dale Webster, Greg S. Corrado, Yossi Matias, Karan Singh, Pete Florence, Alan Karthikesalingam, and Vivek Natarajan. 2023. Towards Generalist Biomedical AI. arXiv:2307.14334 [cs] doi:10.48550/arXiv.2307.14334

[144] United States: National Archives and Records Administration: Office of the Federal Register. 2010. An Act Entitled The Patient Protection and Affordable Care Act.

[145] Anurag Vaidya, Richard J. Chen, Drew F. K. Williamson, Andrew H. Song, Guillaume Jaume, Yuzhe Yang, Thomas Hartvigsen, Emma C. Dyer, Ming Y. Lu, Jana Lipkova, Muhammad Shaban, Tiffany Y. Chen, and Faisal Mahmood. 2024. Demographic Bias in Misdiagnosis by Computational Pathology Models. *Nature Medicine* (2024). doi:10.1038/s41591-024-02885-z

[146] Baptiste Vasey, Myura Nagendran, Bruce Campbell, David A Clifton, Gary S Collins, Spiros Denaxas, Alastair K Denniston, Livia Faes, Bart Geerts, Mudathir Ibrahim, Xiaoxuan Liu, Bilal A Mateen, Piyush Mathur, Melissa D McCradden, Lauren Morgan, Johan Ordish, Campbell Rogers, Suchi Saria, Daniel S W Ting, Peter Watkinson, Wim Weber, Peter Wheatstone, and Peter McCulloch. 2022. Reporting Guideline for the Early Stage Clinical Evaluation of Decision Support Systems Driven by Artificial Intelligence: DECIDE-AI. *BMJ* (2022). doi:10.1136/bmj-2022-070904

[147] Huy V. Vo, Vasil Khalidov, Timothée Darcet, Théo Moutakanni, Nikita Smetanin, Marc Szafraniec, Hugo Touvron, Camille Couprie, Maxime Oquab, Armand Joulin, Hervé Jégou, Patrick Labatut, and Piotr Bojanowski. 2024. Automatic Data Curation for Self-Supervised Learning: A Clustering-Based Approach. arXiv:2405.15613 doi:10.48550/arXiv.2405.15613

[148] Xiya Wang, Jiahui Chen, Zhaoyang Wang, Yuhang Zhou, Yiyang Zhou, Huaxiu Yao, Tianyi Zhou, Tom Goldstein, Parminder Bhatia, Furong Huang, and Cao Xiao. 2024. Enhancing Visual-Language Modality Alignment in Large Vision Language Models via Self-Improvement. doi:10.48550/ARXIV.2405.15973

[149] Xiaosong Wang, Yifan Peng, Le Lu, Zhiyong Lu, Mohammadhdadi Bagheri, and Ronald M. Summers. 2017. ChestX-Ray8: Hospital-Scale Chest X-Ray Database and Benchmarks on Weakly-Supervised Classification and Localization of Common Thorax Diseases. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. doi:10.1109/CVPR.2017.369

[150] Zifeng Wang, Zhenbang Wu, Dinesh Agarwal, and Jimeng Sun. 2022. MedCLIP: Contrastive Learning from Unpaired Medical Images and Text. arXiv:2210.10163 [cs] doi:10.48550/arXiv.2210.10163

[151] Susan Wei and Marc Niethammer. 2021. The Fairness-Accuracy Pareto Front. arXiv:2008.10797 [cs] doi:10.48550/arXiv.2008.10797

[152] Kasumi Widner, Sunny Virmani, Jonathan Krause, Jay Nayar, Richa Tiwari, Elin Rønby Pedersen, Divleen Jeji, Naama Hammel, Yossi Matias, Greg S. Corrado, Yun Liu, Lily Peng, and Dale R. Webster. 2023. Lessons Learned from Translating AI from Development to Deployment in Healthcare. *Nature Medicine* (2023). doi:10.1038/s41591-023-02293-9

[153] David R. Williams, Jourdyn A. Lawrence, Brigitte A. Davis, and Cecilia Vu. 2019. Understanding How Discrimination Can Affect Health. *Health Services Research* (2019). doi:10.1111/1475-6773.13222

[154] Benjamin D. Wissel, Hansel M. Greiner, Tracy A. Glauser, Francesco T. Mangano, Daniel Santel, John P. Pestian, Rhonda D. Szczesniak, and Judith W. Dexheimer. 2019. Investigation of Bias in an Epilepsy Machine Learning Algorithm Trained on Physician Notes. *Epilepsia* (2019). doi:10.1111/epi.16320

[155] World Health Organization. 2010. A Conceptual Framework for Action on the Social Determinants of Health. (2010).

[156] Chaoyi Wu, Xiaoman Zhang, Ya Zhang, Yanfeng Wang, and Weidi Xie. 2023. Towards Generalist Foundation Model for Radiology by Leveraging Web-scale 2D&3D Medical Data. arXiv:2308.02463 [cs] doi:10.48550/arXiv.2308.02463

[157] Peng Xia, Ze Chen, Juanxi Tian, Yangrui Gong, Ruibo Hou, Yue Xu, Zhenbang Wu, Zhiyuan Fan, Yiyang Zhou, Kangyu Zhu, Wenhao Zheng, Zhaoyang Wang, Xiao Wang, Xuchao Zhang, Chetan Bansal, Marc Niethammer, Junzhou Huang, Hongtu Zhu, Yun Li, Jimeng Sun, Zongyuan Ge, Gang Li, James Zou, and Huaxiu Yao. 2024. CARES: A Comprehensive Benchmark of Trustworthiness in Medical Vision Language Models. In *Advances in Neural Information Processing Systems*, A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (Eds.). Curran Associates, Inc.

[158] Yijun Xiao and William Yang Wang. 2021. On Hallucination and Predictive Uncertainty in Conditional Language Generation. arXiv:2103.15025 [cs] doi:10.48550/arXiv.2103.15025

[159] Yunfei Xie, Ce Zhou, Lang Gao, Juncheng Wu, Xianhang Li, Hong-Yu Zhou, Sheng Liu, Lei Xing, James Zou, Cihang Xie, and Yuyin Zhou. 2024. MedTrinity-25M: A Large-scale Multimodal Dataset with Multigranular Annotations for Medicine. arXiv:2408.02900 [cs] doi:10.48550/arXiv.2408.02900

[160] Han Xu, Xiaorui Liu, Yixin Li, Anil K. Jain, and Jiliang Tang. 2021. To Be Robust or to Be Fair: Towards Fairness in Adversarial Training. arXiv:2010.06121 [cs] doi:10.48550/arXiv.2010.06121

[161] Hu Xu, Saining Xie, Xiaoqing Tan, Po-Yao Huang, Russell Howes, Vasu Sharma, Shang-Wen Li, Gargi Ghosh, Luke Zettlemoyer, and Christoph Feichtenhofer. 2023. Demystifying CLIP Data. In *The Twelfth International Conference on Learning Representations*.

[162] Zikang Xu, Jun Li, Qingsong Yao, Han Li, Mingyue Zhao, and S. Kevin Zhou. 2024. Addressing Fairness Issues in Deep Learning-Based Medical Image Analysis: A Systematic Review. *npj Digital Medicine* (2024). doi:10.1038/s41746-024-01276-5

[163] Yuzhe Yang, Haoran Zhang, Judy W. Gichoya, Dina Katabi, and Marzyeh Ghassemi. 2024. The Limits of Fair Medical Imaging AI in Real-World Generalization. *Nature Medicine* (2024). doi:10.1038/s41591-024-03113-4

[164] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rodriguez, and Krishna P. Gummadi. 2017. Fairness Constraints: Mechanisms for Fair Classification. arXiv:1507.05259 [stat] doi:10.48550/arXiv.1507.05259

[165] Anna Zawacki. 2020. SIIM-ISIC Melanoma Classification. <https://kaggle.com/siim-isic-melanoma-classification>.

[166] Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. 2023. Sigmoid Loss for Language Image Pre-Training. arXiv:2303.15343 [cs] doi:10.48550/arXiv.2303.15343

[167] Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. 2023. Sigmoid Loss for Language Image Pre-Training. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*. IEEE, Paris, France. doi:10.1109/ICCV51070.2023.01100

[168] Haoran Zhang, Natalie Dullerud, Karsten Roth, Lauren Oakden-Rayner, Stephen Robert Pfahl, and Marzyeh Ghassemi. 2022. Improving the Fairness of Chest X-ray Classifiers. arXiv:2203.12609 [cs] doi:10.48550/arXiv.2203.12609

[169] Ruichen Zhang, Yuguang Yao, Zhen Tan, Zhiming Li, Pan Wang, Huan Liu, Jingtong Hu, Sijia Liu, and Tianlong Chen. 2024. FairSkin: Fair Diffusion for Skin Disease Image Generation. arXiv:2410.22551

[170] Shaoting Zhang and Dimitris Metaxas. 2023. On the Challenges and Perspectives of Foundation Models for Medical Image Analysis. *Medical Image Analysis* (2023). doi:10.1016/j.media.2023.102996

[171] Sheng Zhang, Yanbo Xu, Naoto Usuyama, Hanwen Xu, Jaspreet Bagga, Robert Tinn, Sam Preston, Rajesh Rao, Mu Wei, Naveen Valluri, Cliff Wong, Andrea Tupini, Yu Wang, Matt Mazzola, Swadheen Shukla, Lars Liden, Jianfeng Gao, Angela Crabtree, Brian Piening, Carlo Bifulco, Matthew P. Lungren, Tristan Naumann, Sheng Wang, and Hoifung Poon. 2025. BiomedCLIP: A Multimodal Biomedical Foundation Model Pretrained from Fifteen Million Scientific Image-Text Pairs. arXiv:2303.00915 [cs] doi:10.48550/arXiv.2303.00915

[172] Xiaoman Zhang, Chaoyi Wu, Ziheng Zhao, Weixiong Lin, Ya Zhang, Yanfeng Wang, and Weidi Xie. 2023. PMC-VQA: Visual Instruction Tuning for Medical Visual Question Answering. doi:10.48550/ARXIV.2305.10415

[173] Han Zhao and Geoffrey J. Gordon. 2022. Inherent Tradeoffs in Learning Fair Representations. arXiv:1906.08386 [cs] doi:10.48550/arXiv.1906.08386

[174] Theodore Zhao, Yu Gu, Jianwei Yang, Naoto Usuyama, Ho Hin Lee, Sid Kiblawi, Tristan Naumann, Jianfeng Gao, Angela Crabtree, Jacob Abel, Christine Moung-Wen, Brian Piening, Carlo Bifulco, Mu Wei, Hoifung Poon, and Sheng Wang. 2024. A Foundation Model for Joint Segmentation, Detection and Recognition of Biomedical Objects across Nine Modalities. *Nature Methods* (2024). doi:10.1038/s41592-024-02499-w

[175] Yukun Zhou, Mark A. Chia, Siegfried K. Wagner, Murat S. Ayhan, Dominic J. Williamson, Robbert R. Struyven, Timing Liu, Moucheng Xu, Mateo G. Lozano, Peter Woodward-Court, Yuka Kihara, Andre Altmann, Aaron Y. Lee, Eric J. Topol, Alastair K. Denniston, Daniel C. Alexander, and Pearse A. Keane. 2023. A Foundation Model for Generalizable Disease Detection from Retinal Images. *Nature* (2023). doi:10.1038/s41586-023-06555-x

- [176] Yuyin Zhou, Shih-Cheng Huang, Jason Alan Fries, Alaa Youssef, Timothy J. Amrhein, Marcello Chang, Imon Banerjee, Daniel Rubin, Lei Xing, Nigam Shah, and Matthew P. Lungren. 2021. RadFusion: Benchmarking Performance and Fairness for Multimodal Pulmonary Embolism Detection from CT and EHR. arXiv:2111.11665 [eess] doi:10.48550/arXiv.2111.11665
- [177] Yongshuo Zong, Yongxin Yang, and Timothy Hospedales. 2023. MEDFAIR: Benchmarking Fairness for Medical Imaging. arXiv:2210.01725 [cs, eess] doi:10.48550/arXiv.2210.01725
- [178] James Zou, Judy Wawira Gichoya, Daniel E. Ho, and Ziad Obermeyer. 2023. Implications of Predicting Race Variables from Medical Images. *Science* (2023). doi:10.1126/science.adh4260
- [179] Remco Zwetsloot, Baobao Zhang, Noemi Dreksler, Lauren Kahn, Markus Anderljung, Allan Dafoe, and Michael C. Horowitz. 2021. Skilled and Mobile: Survey Evidence of AI Researchers' Immigration Preferences. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society (AIES '21)*. Association for Computing Machinery, New York, NY, USA. doi:10.1145/3461702.3462617

Received 28 March 2025; revised 30 November 2025; accepted 13 January 2026

Just Accepted