

FairCLIP: Harnessing Fairness in Vision-Language Learning

Yan Luo^{1*} Min Shi^{1*} Muhammad Osama Khan^{2*} Muhammad Muneeb Afzal² Hao Huang³
Shuaihang Yuan³ Yu Tian¹ Luo Song¹ Ava Kouhana¹ Tobias Elze¹ Yi Fang^{2,3†} Mengyu Wang^{1†}

¹Harvard Ophthalmology AI Lab, Harvard University

²Tandon School of Engineering, New York University

³Multimedia and Visual Computing Lab, New York University Abu Dhabi

{yluo16, mshi6}@meei.harvard.edu, {osama.khan, muneeb.afzal, hh1811, sy2366}@nyu.edu,

{ytian11, lsong7, akouhana, tobias.elze}@meei.harvard.edu,

yfang@nyu.edu, mengyu.wang@meei.harvard.edu

Abstract

Fairness is a critical concern in deep learning, especially in healthcare, where these models influence diagnoses and treatment decisions. Although fairness has been investigated in the vision-only domain, the fairness of medical vision-language (VL) models remains unexplored due to the scarcity of medical VL datasets for studying fairness. To bridge this research gap, we introduce the first fair vision-language medical dataset (Harvard-FairVLMed) that provides detailed demographic attributes, ground-truth labels, and clinical notes to facilitate an in-depth examination of fairness within VL foundation models. Using Harvard-FairVLMed, we conduct a comprehensive fairness analysis of two widely-used VL models (CLIP and BLIP2), pre-trained on both natural and medical domains, across four different protected attributes. Our results highlight significant biases in all VL models, with Asian, Male, Non-Hispanic, and Spanish being the preferred subgroups across the protected attributes of race, gender, ethnicity, and language, respectively. In order to alleviate these biases, we propose FairCLIP, an optimal-transport-based approach that achieves a favorable trade-off between performance and fairness by reducing the Sinkhorn distance between the overall sample distribution and the distributions corresponding to each demographic group. As the first VL dataset of its kind, Harvard-FairVLMed holds the potential to catalyze advancements in the development of machine learning models that are both ethically aware and clinically effective. Our dataset and code are available at <https://ophai.hms.harvard.edu/datasets/harvard-fairvmed10k>.

*Contributed equally as co-first authors.

†Contributed equally as co-senior authors.

Table 1. Public medical fairness datasets. In contrast to existing medical VL datasets, our Harvard-FairVLMed provides detailed demographic attributes, ground-truth labels, and clinical notes with imaging as well as non-imaging clinical information.

Dataset	Data Modality	# of Images	# of Patients	Identity Attribute	V	L	L Info Type
Fitzpatrick17k [16]	Skin Photos	16,012	1,373	Skin type	✓	✗	
IAM1000 [75]	Dermatoscopy	9,948	-	Age; Gender	✓	✗	
OLSI[30]	Heart CT	8,139	1,686	Age; Gender	✓	✗	
ODIR-2019 [1]	Color Fundus	8,000	5,000	Age; Gender	✓	✗	
PAPILA [34]	Color Fundus	488	244	Age; Gender	✓	✗	
Harvard-GDP [45]	OCT	1,000	1,000	Age; Gender; Race;	✓	✗	
				Ethnicity			
COVID-CT-MD[3]	Lung CT	308	305	Age; Gender	✓	✗	
AMD-OCT[14]	OCT	384	-	Age	✓	✗	
ADNI 1.5T [83]	Brain MRI	550	-	Age; Gender	✓	✗	
CheXpert [24]	Chest X-ray and Radiology Report	222,793	65,240	Age; Gender; Race	✓	✓	Image Description
MIMIC-CXR [27]	Chest X-ray and Radiology Report	370,955	65,079	Age; Gender; Race	✓	✓	Image Description
PadChest [12]	Chest X-ray and Radiology Report (Spanish)	160,868	69,882	Age; Gender	✓	✓	Image Description
Harvard-FairVLMed	SLO Fundus and Clinical Note	Fundus: 10,000; Note: 10,000	10,000	Age; Gender; Race; Ethnicity; Preferred Language; Marital Status	✓	✓	Image Description + Non-Imaging Clinical Info

1. Introduction

Fairness has received increasing interest in deep learning in recent years. This is vital, especially in healthcare, where these deep learning models influence diagnoses and treatment decisions. Biases in these models related to factors like race, gender, or socioeconomic status can lead to healthcare disparities and adverse patient outcomes. Hence, ensuring that these models are free from bias is not only an ethical and legal requirement but also a necessity for ensuring patient safety and healthcare equity. This makes fairness in medical computer vision a critical and immediate concern, essential for providing equitable healthcare.

Previous research has identified biases in deep learning-based medical imaging models, focusing primarily on Chest X-ray diagnosis [15, 31]. In contrast to these vision-only models, there has been a recent surge of vision-language (VL) foundation models [21, 29, 30, 44, 78, 84, 85], which have set new benchmarks across a wide spectrum of tasks. However, despite these strong performances, the fairness of these VL models remains unclear. Given the existence of biases in vision-only models (which operate on machine-acquired images) and the human-written nature of clinical

notes, VL models could further aggravate fairness issues. Hence, as the field moves towards multi-modal foundation models, it becomes increasingly critical to scrutinize how the interplay of vision and text affects equity in algorithmic outcomes. The current landscape for such an investigation is, however, constrained by the scarcity of VL datasets that include comprehensive demographic information, with the existing public VL datasets primarily focused on Chest X-rays [24, 26]. Prior studies [24, 65] have highlighted the challenges of studying fairness using these datasets since their *ground-truth* labels are automatically extracted from radiology reports, potentially leading to inaccurate fairness conclusions due to the noisy labels. Moreover, since these datasets were not primarily designed for fairness, they provide only a handful of demographic characteristics, limiting the potential for a comprehensive fairness study across multiple dimensions. Furthermore, radiology reports, focusing mainly on direct observations of the imaging data with limited additional patient-specific information, are not representative of most clinical text, thus limiting their utility in fairness studies of medical VL models.

To bridge this research gap, we introduce the first fair vision-language medical dataset (*Harvard-FairVLMed* for short) that provides detailed demographic attributes, ground-truth labels, and clinical notes to facilitate an in-depth examination of fairness within VL foundation models (Table 1). *Harvard-FairVLMed* contains records for 10,000 patients, each paired with an SLO fundus image and a clinical note for diagnosing Glaucoma, along with fine-grained protected attributes such as age, gender, race, ethnicity, preferred language, and marital status. Unlike radiology reports, the clinical notes in our dataset provide much more detailed information featuring not only image descriptions but also rich non-imaging clinical information such as medication, non-imaging test results, and family history. Hence, these clinical notes are more representative of clinical textual information, making them better suited for studying the fairness of medical VL models.

Glaucoma, affecting millions globally [17, 46, 47, 67, 69, 70, 72, 74], exemplifies the need for equitable diagnostic models. Timely detection is crucial in order to avoid irreversible vision loss. However, many patients remain undiagnosed [66] due to the disease’s asymptomatic nature and barriers to ophthalmic care. Moreover, this undiagnosed issue is even more pronounced among minority groups. For example, previous studies have shown that individuals from Black communities are 4.4 times more likely to have undiagnosed and untreated Glaucoma compared to their White counterparts [66], highlighting the importance of addressing healthcare disparities. Deep learning systems hold significant potential for improving healthcare. However, it is imperative to address potential fairness issues in these deep learning systems prior to their clinical implementation to

ensure equitable healthcare delivery.

In this work, we conduct an extensive fairness analysis with two widely-used VL methods (*i.e.*, CLIP [61] and BLIP2 [39]) on *Harvard-FairVLMed*. Our experimental findings reveal significant disparities across various groups based on race, gender, ethnicity, and language. To address these fairness issues, we introduce FairCLIP, an optimal transport-based approach. FairCLIP is designed to enhance fairness by optimizing the Sinkhorn distance, thereby aligning the overall sample distribution with the distributions specific to each demographic group.

Our main contributions can be summarized as follows:

- We introduce the first fair vision-language medical dataset (*Harvard-FairVLMed*) featuring detailed demographic attributes, ground-truth labels, and clinical notes for studying the fairness of VL foundation models.
- Using *Harvard-FairVLMed*, we conduct a comprehensive fairness analysis of two widely-used VL models (*i.e.*, CLIP and BLIP2), pre-trained on both natural and medical domains, across four different protected attributes.
- Our results highlight significant biases in all VL models, with Asian, Male, Non-Hispanic, and Spanish being the preferred subgroups across the protected attributes of race, gender, ethnicity, and language, respectively.
- We propose an optimal transport-based approach named FairCLIP, which demonstrates a significant improvement over CLIP in terms of both performance and fairness.

2. Related Work

Fairness Models Fairness in machine learning is crucial for creating dependable systems that avoid discrimination. Historical instances of bias have been identified in various technologies. In the field of image processing, biases exist in facial recognition [7, 11, 18, 76, 77] and pedestrian detection algorithms [9, 32, 37]. Recent efforts in deep learning aim to rectify these biases by adjusting training processes to minimize unfair predictions across different demographic groups. These deep learning methods focus on removing or neutralizing sensitive features and fostering representations that are fair and unbiased.

Particularly in the medical domain, fairness models have emerged as a critical tool to mitigate biases and promote equity, particularly within healthcare applications. This field has developed a range of methodologies, each targeting a different stage: pre-processing, in-processing, and post-processing. Pre-processing strategies [10, 28, 55, 59, 62, 62, 65, 74, 93, 93, 97, 98] are designed to mitigate biases by enhancing the dataset prior to model training. In in-processing approaches [8, 8, 63, 63, 64, 64, 68, 89, 91, 91] fairness is incorporated directly into the model’s training stage. Post-processing methods [49, 58, 82], applied after the model has been trained, aim to adjust outputs to improve fairness.

To address the challenges of fairness in VL tasks, this work introduces FairCLIP, a method designed to promote fairness. FairCLIP aligns the distribution of all samples with those of specific demographic groups, leveraging the similarities between visual and textual features.

Vision-Language Models The integration of vision and language in machine learning, especially through models like CLIP (Contrastive Language-Image Pretraining) [61], represents a significant advancement in computer vision. CLIP, trained from a large scale of images and their related text descriptions, marks an important step in aligning visual information with text. BLIP (Bootstrapping Language-Image Pre-training) [38] represents a novel approach in vision-language models, utilizing a visual transformer as an image encoder and a transformer for text processing. BLIP-2 [39] advances the field with its scalable multimodal pre-training approach, enabling LLMs to process and comprehend images.

The success of CLIP-like models has inspired numerous adaptations in the medical field [36, 43, 78, 79, 87, 92, 96], where it confronts challenges like limited data and the need for accurate interpretations. For instance, Pathology Language-Image Pretraining (PLIP) [21] utilizes the OpenPath dataset [21], consisting of pathology images with language descriptions derived from medical Twitter posts. PMC-CLIP (PubMedCentral-CLIP) [40] joins these efforts by enhancing CLIP’s application in the biomedical domain.

Beyond these adaptations of CLIP, several other vision-language models have emerged to address various medical tasks that range from generating reports from medical images [42, 52, 53, 73, 81, 86] to diagnosing medical conditions [50, 51, 94, 95], and facilitating medical image-based question-answering (VQA) [6, 33].

Despite these advancements, a gap remains in understanding and addressing fairness concerns within vision-language models (VLMs). The complex interplay of visual and textual data in VLMs poses unique challenges for fairness, which have been largely unexplored, especially in the context of medical applications.

Vision-Language Medical Fairness Datasets Medical fairness datasets serve a critical role in the development of machine learning models [23, 27, 45, 47, 48], ensuring that they operate equitably across diverse patient populations. In this context, datasets equipped with imaging data as well as associated textual reports are particularly valuable. They not only enable the development of models but also allow for nuanced explorations of fairness by providing a richer context for each medical case.

CheXpert [23] is a prominent chest radiography dataset annotated for 14 distinct observations. CheXpert uses an automated system to extract labels from radiology reports. However, this method introduces uncertainties due to the inherent ambiguities in the reports and varying interpreta-

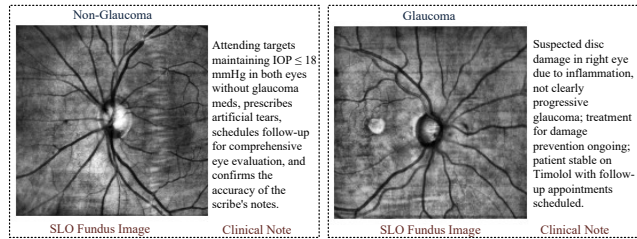


Figure 1. Examples of non-glaucomatous and glaucomatous samples with the corresponding SLO fundus image and clinical note.

tions by radiologists. This uncertainty, such as the 16% of Atelectasis labels that are uncertain, represents a significant challenge in using datasets like this in real-world scenarios [2, 23]. MIMIC-CXR [27] contributes to this area in which each study is complemented by a semi-structured free-text radiology report. Padchest [12] further expands the resources by providing chest X-ray images along with a set of radiographic findings that have been processed using radiology reports. However, the radiological reports are only available in Spanish, which may present a barrier for the global research community. Similarly, datasets pertaining to image captioning [13, 25, 56, 71] and VQA [20, 35, 41] have been developed in the medical domain.

In contrast, fundus image datasets present a different scenario. While ODIR2019 [1] and PAPILA [34] offer valuable imaging data for ocular conditions, they do not provide the same depth of textual information as CheXpert and MIMIC-CXR. Also, the Glaucoma Detection and Progression [45], Eye Fairness [47], and Glaucoma Fairness [48] datasets lack textual information required for studying the VL problem. While datasets such as CheXpert and MIMIC-CXR have advanced the field of fairness research by pairing images with radiology reports, they fall short in providing extensive demographic data and are prone to noisy labels due to their automated labeling processes. Additionally, they were not initially created with a focus on fairness, which poses limitations in studying fairness.

Our Harvard-FairVLMed dataset stands as a pioneering resource tailored for fairness studies in VL modeling. It comprises ground-truth labels, ensuring high label quality, and includes a broad range of identity attributes such as age, gender, race, ethnicity, preferred language, and marital status. Hence, this dataset not only enhances the quality and reliability of fairness studies but also broadens the horizon for research in the VL domain.

3. Dataset Analysis

This study strictly adheres to the principles outlined in the Declaration of Helsinki and has been approved by our institute’s Institutional Review Board. All data in this dataset are de-identified.

3.1. Data Collection and Quality Control

Harvard-FairVLMed contains subjects who received glaucoma care from Massachusetts Eye and Ear at Harvard Medical School between 2015 and 2022. There are three types of data to be released in this study: (1) scanning laser ophthalmoscopy (SLO) fundus images; (2) demographic identity group information; and (3) de-identified clinical notes written by ophthalmologists to provide a summary of the glaucoma diagnosis. SLO fundus images are a valuable marker for assessing retinal damage from diseases like glaucoma. Each SLO fundus image is associated with six demographic identity attributes, including age, gender, race, ethnicity, preferred language, and marital status. The accompanying clinical notes vary in length. These notes may detail assessments, treatment plans, and diagnostic strategies, and are considered to correspond with the visual semantics in SLO fundus images. Two examples of pairs of SLO fundus images and clinical notes are shown in Figure 1. The subjects are categorized into non-glaucoma (visual function measured by visual field (VF) test is normal: VF mean deviation ≥ -1 dB and normal VF glaucoma hemifield test and pattern standard deviation (PSD) results) and glaucoma categories (visual function measured by VF test is abnormal: VF mean deviation < -3 dB and abnormal VF glaucoma hemifield test and PSD results).

3.2. Protected Information De-Identification

The raw clinical notes may contain protected sensitive information, such as the date of glaucoma diagnosis, patient name, phone number, email address, physical location, institution, etc. We de-identified this sensitive information in the following three steps. First, we anonymized all the clinical notes using Microsoft’s Presidio¹, which replaced sensitive information with respective placeholders (e.g., PERSON_NAME, PHONE_NUMBER, and LOCATION) so that the original sentence structure and coherence can be maintained. Then, we used rules to match and de-identify the protected information (e.g., physical address) that has not been fully recognized by Presidio. Lastly, the de-identified clinical notes were further verified by four medical experts. Specially, every clinical note was checked by an expert and the sensitive information was manually replaced with respective placeholders when necessary.

3.3. Data Characteristics

Our Harvard-FairVLMed dataset comprises 10,000 samples from 10,000 subjects. It is divided into 7,000 training, 1,000 validation, and 2,000 test samples. The dataset’s collective age average is 60.9 ± 16.2 years. The dataset includes samples from three major groups: Asian, with 819 samples; Black, with 1,491 samples; and White, with 7,690 samples.

¹<https://github.com/microsoft/presidio>.

Gender-wise, females constitute 56.3% of the subjects, with the remainder being males. The ethnic distribution is highlighted by 90.6% Non-Hispanic, 4.0% Hispanic, and 5.4% unspecified. In terms of preferred language, 92.5% of the subjects prefer English, 1.7% prefer Spanish, 0.8% prefer other languages, and 5.0% remain unknown. From a marital status perspective, 57.4% are in a marriage or partnered, 26.4% are single, 6.6% have experienced divorce, 1.0% are legally separated, 6.1% are widowed, and 2.5% are not specified. After de-identification, clinical notes vary from 11 to 332 words, with an average word count of 147.

4. Methodology

4.1. Background

With the labeled data $D = \{(x_{Ii}, x_{Ti}, y_i, a_i)\}$, where $x_{Ii} \in \mathcal{X}_I$ represents an SLO fundus image, $x_{Ti} \in \mathcal{X}_T$ represents a clinical note, $y \in \mathcal{Y}$ denotes a glaucoma diagnosis label, and $a \in \mathcal{A}$ signifies an identity attribute associated with the patient, such as gender, race, or ethnicity. To process multimodal data, a vision-language (VL) pre-trained model (e.g., CLIP [60]) f employs a vision encoder f_I and a text encoder f_T to generate vision features z_I and text features z_T , respectively. Given a batch of N image and text pairs, CLIP is trained to generate a similarity matrix $M \in \mathbb{R}^{N \times N}$ with each entry $M_{ij} = z_{Ii}^\top z_{Tj}$ representing the cosine similarity between the i -th image feature z_{Ii} and the j -th text feature within the given batch. To achieve this, CLIP adopts a contrastive learning scheme to maximize the cosine similarity of the image and text features of the N positive pairs in the batch, while minimizing the cosine similarity of the image and text features of the remaining $N^2 - N$ negative pairs. Following [60], a symmetric cross-entropy loss is calculated over M to optimize f_I and f_T . Mathematically, the optimization goal is defined as follows:

$$\min_f - \sum_{i=1}^N \sum_{j=1}^N \delta(i-j) \log(z_{Ii}^\top z_{Tj} / (\|z_{Ii}\| \cdot \|z_{Tj}\|)),$$

where $\delta(i-j)$ is Dirac’s Delta function, $\delta(i-j) = 1$ when $i = j$. In the context of fairness learning, it requires to incorporate identity information during model training, i.e., $f \in \mathcal{F} : \mathcal{X}_I \times \mathcal{X}_T \times \mathcal{A} \xrightarrow{\theta} \mathcal{Y}$. Consequently, fairness learning aims to minimize disparities between different identity groups while also maximizing accuracy.

4.2. FairCLIP

As illustrated in Figure 2, the proposed FairCLIP framework is designed to improve fairness during the pre-training phase. This is achieved by minimizing the disparity between the probability distributions of $M_{i,i}$, which represents the correlation between visual and language features, across different racial groups (or other attribute-

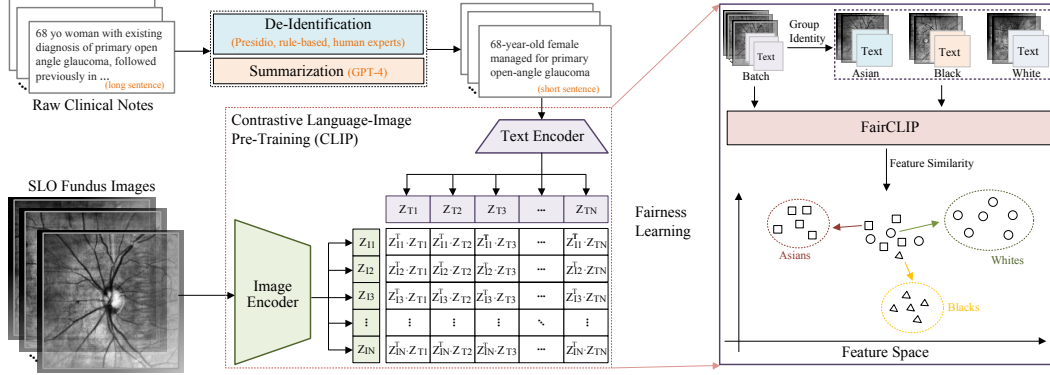


Figure 2. Schematic view of the proposed FairCLIP. Clinical notes containing PHI (e.g., names and gender) undergo de-identification and summarization to fit text encoder limitations, such as CLIP’s 77-token maximum length. FairCLIP equalizes the overall sample distribution with the distributions corresponding to each demographic group, thereby achieving a favorable trade-off between performance and fairness.

based groups). Denote $\mathcal{D}_{\{(x_I, x_T, a)\}}|f$ as a distribution of $M_{i,i}$ given the model f . If the samples are from a specific group (e.g., white), the corresponding distribution is $\mathcal{D}_{\{(x_I, x_T, a)|a=\text{white}\}}|f$. Then, the objective of enhancing fairness can be defined as:

$$\min_f \sum_{\alpha}^A d(\mathcal{D}_{\{(x_I, x_T, a)\}}|f - \mathcal{D}_{\{(x_I, x_T, a)|a=\alpha}\}}|f) \quad (1)$$

where d is a distance function, $\mathcal{D}_{\{(x_I, x_T, a)\}}|f$ and $\mathcal{D}_{\{(x_I, x_T, a)|a=\alpha}\}}|f$ are the underlying distributions that are computationally intractable. Instead, we use empirical distributions in Equation (1) that are estimated based on batch \mathcal{B} , i.e., $\mathcal{D}_{\mathcal{B}}|f$ and $\mathcal{D}_{\mathcal{B}_a}|f$. \mathcal{B}_a means the samples in the batch are from group a .

To optimize the objective (1), one straightforward way is to minimize Kullback–Leibler (KL) divergence between the two distributions. However, KL-divergence is not symmetric and does not satisfy the triangle inequality, and thus not a true metric of distance. Instead, we follow [57] to minimize the Sinkhorn distance between the two distributions. Sinkhorn distance is a type of probability metric and a variant of the Wasserstein distance.

Specifically, we have a measurable set $\mathcal{Z} = \{z_I^T, z_T\}$, denoted by $\mathcal{M}(\mathcal{Z})$ the set of measures (not necessarily probability measures) on \mathcal{Z} , and $\mathcal{P}(\mathcal{Z})$ the set of probability measures on \mathcal{Z} . Consider distributions $\mathcal{D}_{\mathcal{B}}, \mathcal{D}_{\mathcal{B}_a} \in \mathcal{P}(\mathcal{Z})$, and let $\mu, \nu \in \mathcal{M}(\mathcal{Z})$ be two reference measures. For regularization parameter $\epsilon \geq 0$, the Sinkhorn distance between two distributions $\mathcal{D}_{\mathcal{B}}$ and $\mathcal{D}_{\mathcal{B}_a}$ is defined as

$$\mathcal{W}_{\epsilon}(\mathcal{D}_{\mathcal{B}}, \mathcal{D}_{\mathcal{B}_a}) = \inf_{\gamma \in \Gamma(\mathcal{D}_{\mathcal{B}}, \mathcal{D}_{\mathcal{B}_a})} \left\{ \mathbb{E}_{(x,y) \sim \gamma} [c(p, q)] + \epsilon H(\gamma | \mu \otimes \nu) \right\}, \quad (2)$$

where $\Gamma(\mathcal{D}_{\mathcal{B}}, \mathcal{D}_{\mathcal{B}_a})$ denotes the set of joint distributions whose first and second marginal distributions are $\mathcal{D}_{\mathcal{B}}$ and

$\mathcal{D}_{\mathcal{B}_a}$, respectively, $c(p, q)$ denotes the transport cost, and $H(\gamma | \mu \otimes \nu)$ denotes the relative entropy of γ with respect to the product measure $\mu \otimes \nu$. p and q are the points in the distributions $\mathcal{D}_{\mathcal{B}}$ and $\mathcal{D}_{\mathcal{B}_a}$, respectively. The Sinkhorn loss would be added to the loss used by CLIP in the pre-training phase for optimizing each component in CLIP.

5. Experiment & Analysis

5.1. Experimental Setup

In this section, we describe the pre-training and evaluation setups as well as the metrics used to study performance and fairness of VL models on our proposed Harvard-FairVLMed dataset.

Pre-Training: We use the widely-adopted VL methods – CLIP [60] and BLIP2 [39] – for our analysis. For the natural pre-trained variants, we use the official checkpoints provided by CLIP and BLIP2. For the medical pre-trained variants, we pre-train both methods on the Harvard-FairVLMed dataset after initializing from the official checkpoints. More details are in the supplementary material.

Evaluation: We utilize two types of evaluation strategies – linear probing and zero-shot transfer. For linear probing, we follow the official MAE [19] implementation and train a linear classifier on top of the visual features from CLIP and BLIP2, respectively. Similar to MAE, we use a BatchNorm layer [22] before the linear classifier and employ a LARS optimizer [88] with a base learning rate of 0.1, weight decay of 0, and batch size of 512. We conduct linear probing for 1000 epochs on a single V100 GPU for all experiments. For zero-shot transfer, we use the same setup as CLIP and select the class corresponding to the text embedding, which has the highest similarity with the image embedding.

Metrics: To comprehensively understand the balance between model performance and fairness, we use multiple metrics for evaluation, including Demographic Parity Dif-

Table 2. VL fairness analysis of two widely-used VL methods, *i.e.*, CLIP and BLIP2, pre-trained on both natural and medical domains and evaluated across four different protected attributes, with all scores presented in percentage.

Attribute	Model	DPD ↓	DEOdds ↓	AUC ↑	ES-AUC ↑	Group-wise AUC ↑		
Race	CLIP	5.30 ± 0.63	14.00 ± 1.01	77.27 ± 0.03	72.43 ± 0.29	Asian	Black	White
	CLIP-FT	4.01 ± 0.47	9.57 ± 0.83	80.27 ± 0.08	74.70 ± 0.33	79.74 ± 0.31	73.60 ± 0.12	77.82 ± 0.03
	BLIP2	9.44 ± 0.65	10.62 ± 0.22	73.81 ± 0.02	68.88 ± 0.04	82.19 ± 0.26	75.67 ± 0.21	81.20 ± 0.08
	BLIP2-FT	8.30 ± 0.36	10.91 ± 0.32	80.10 ± 0.03	73.81 ± 0.10	76.28 ± 0.06	69.55 ± 0.09	74.22 ± 0.03
Gender	CLIP	1.08 ± 0.19	5.19 ± 0.54	77.27 ± 0.03	72.47 ± 0.10	Female	Male	
	CLIP-FT	0.39 ± 0.26	4.55 ± 0.33	80.27 ± 0.08	75.81 ± 0.12	74.25 ± 0.07	80.88 ± 0.03	
	BLIP2	1.07 ± 0.22	5.88 ± 0.24	73.81 ± 0.02	69.16 ± 0.03	77.59 ± 0.09	83.47 ± 0.04	
	BLIP2-FT	2.41 ± 0.06	6.40 ± 0.26	80.10 ± 0.03	75.08 ± 0.15	70.76 ± 0.02	77.48 ± 0.05	
Ethnicity	CLIP	15.83 ± 0.42	14.73 ± 0.54	77.27 ± 0.03	71.70 ± 0.08	Non-Hispanic	Hispanic	
	CLIP-FT	14.50 ± 0.72	22.49 ± 1.44	80.27 ± 0.08	76.31 ± 0.36	77.51 ± 0.03	69.73 ± 0.13	
	BLIP2	8.78 ± 1.35	16.56 ± 1.89	73.81 ± 0.02	68.75 ± 0.08	80.48 ± 0.07	75.30 ± 0.46	
	BLIP2-FT	16.64 ± 1.17	18.41 ± 1.98	80.10 ± 0.03	77.13 ± 0.09	74.10 ± 0.02	66.74 ± 0.14	
Language	CLIP	13.57 ± 1.35	33.11 ± 0.53	77.27 ± 0.03	70.89 ± 0.21	English	Spanish	Others
	CLIP-FT	16.75 ± 1.08	15.74 ± 0.28	80.27 ± 0.08	67.06 ± 0.46	77.25 ± 0.03	84.00 ± 0.16	75.02 ± 0.28
	BLIP2	22.40 ± 0.33	15.41 ± 1.41	73.81 ± 0.02	70.34 ± 0.64	80.77 ± 0.07	74.43 ± 0.99	66.91 ± 0.17
	BLIP2-FT	14.08 ± 3.56	37.68 ± 2.51	80.10 ± 0.03	69.98 ± 0.56	73.40 ± 0.02	75.95 ± 0.87	76.19 ± 0.12
						80.62 ± 0.04	83.14 ± 1.18	69.51 ± 0.32

ference (DPD) [4, 5], Difference in Equalized Odds (DEOdds) [4], Area Under the Receiver Operating Characteristic Curve (AUC), Equity-Scaled AUC [48], and Group-wise AUC. More details can be found in the supplementary material.

5.2. VL Fairness Analysis

In this section, we present a comprehensive fairness analysis of two widely-used VL models, benchmarked across two different pre-training domains and four different protected attributes. Table 2 presents the linear probing results, examining various performance (AUC) and fairness (DPD, DEOdds, ES-AUC) metrics, as well as reporting the group-wise AUC scores across the individual subgroups within each of the four protected attributes. We primarily focus on the ES-AUC metric in the subsequent analysis since it captures notions of both overall performance as well as fairness – both of which are important for safety-critical medical applications (Sec. 5.1). Next, we briefly discuss the disparities in VL performance across the various protected attributes, and the impact of different VL pre-training domains (natural vs. medical) and VL pre-training methods (CLIP vs. BLIP2) on model performance and fairness.

Protected Attributes: Across the four protected attributes – race, gender, ethnicity, and language – our results indicate that VL models exhibit the best performance-fairness trade-off on ethnicity and the worst on language with average ES-AUC scores (across the four VL models) of 73.47 and 69.57, respectively. A granular analysis reveals that in terms of racial subgroups, Asian patients consistently have the highest diagnostic performance, whereas Black patients have the lowest. Across genders, male patients are consistently better diagnosed than female patients. Moreover,

non-Hispanic is the highest-performing subgroup across ethnicities, whereas Spanish speakers have the best diagnostic performance across the language subgroups. Some of these performance disparities could be attributed to the imbalances in the pre-training datasets. For instance, non-Hispanic patients make up 90.6% of our dataset, potentially leading to improved performance in this subgroup. However, this is unlikely to be the only factor responsible for these performance disparities since the Asian, Male, and Spanish subgroups exhibit superior performances despite being the minority subgroups. This indicates that the pre-training of these models could potentially play a role in the biases exhibited by these models. We delve deeper into this in the subsequent analysis.

Natural vs. Medical Pre-training: Here, we compare the VL models pre-trained on two different domains – natural vs. medical VL datasets. Comparing the two pre-training domains, we observe that pre-training on medical data (*i.e.*, paired fundus images + clinical notes) improves the performance-fairness trade-off across all protected attributes except language. Concretely, we see the most benefit on ethnicity, with improvements of as much as 8.38 in terms of ES-AUC. A granular analysis of the individual subgroups shows that the White and Hispanic subgroups witness the most improvement from medical pre-training, with improvements of as much as 6.75 and 9.65 ES-AUC in the former and latter, respectively. This indicates that medical VL pre-training could be an effective strategy for improving the performances of under-performing subgroups, thereby effectively reducing the performance disparities across different protected attributes.

CLIP vs. BLIP2: Next, we investigate the impact of different VL pre-training methods on the performance dispar-

Table 3. Zero-shot transfer results of CLIP vs. FairCLIP, reporting the mean and standard deviation across three random seeds.

Attribute	Model	DPD ↓	DEOdds ↓	AUC ↑	ES-AUC ↑	Group-wise AUC ↑		
Race						Asian	Black	White
	CLIP (ViT-B/16)	15.35 ± 6.50	15.11 ± 5.01	67.84 ± 0.90	61.67 ± 0.63	73.11 ± 2.74	70.78 ± 1.84	66.02 ± 0.60
	FairCLIP (ViT-B/16)	6.07 ± 2.44	10.50 ± 2.73	70.24 ± 1.26	65.50 ± 2.60	74.83 ± 0.46	71.39 ± 0.66	69.17 ± 2.10
	CLIP (ViT-L/14)	10.10 ± 9.44	10.79 ± 10.41	67.83 ± 2.92	63.53 ± 1.83	70.65 ± 4.58	70.12 ± 3.39	66.22 ± 2.97
FairCLIP (ViT-L/14)	17.79 ± 4.86	18.30 ± 2.07	69.88 ± 2.00	66.54 ± 1.73	71.78 ± 2.18	71.79 ± 2.13	68.67 ± 1.99	
Gender						Female	Male	
	CLIP (ViT-B/16)	4.34 ± 0.66	9.95 ± 0.64	67.84 ± 0.90	63.21 ± 0.83	64.62 ± 0.83	71.96 ± 1.22	
	FairCLIP (ViT-B/16)	0.84 ± 0.25	2.97 ± 2.07	69.76 ± 2.49	65.39 ± 2.39	66.81 ± 2.50	73.50 ± 2.41	
	CLIP (ViT-L/14)	2.93 ± 3.17	4.29 ± 4.05	67.83 ± 2.92	63.86 ± 2.36	65.13 ± 2.60	71.31 ± 3.24	
FairCLIP (ViT-L/14)	5.82 ± 1.22	8.14 ± 2.62	69.74 ± 0.95	66.00 ± 1.55	67.29 ± 1.38	72.99 ± 0.83		
Ethnicity						Non-Hispanic	Hispanic	
	CLIP (ViT-B/16)	8.86 ± 5.95	15.33 ± 5.18	67.84 ± 0.90	63.09 ± 0.26	68.08 ± 0.97	60.56 ± 0.30	
	FairCLIP (ViT-B/16)	9.12 ± 3.25	14.30 ± 7.20	69.30 ± 1.87	63.33 ± 0.25	69.62 ± 2.02	60.19 ± 1.18	
	CLIP (ViT-L/14)	7.78 ± 5.50	11.56 ± 9.03	67.83 ± 2.92	62.20 ± 0.82	68.14 ± 3.07	59.13 ± 0.92	
FairCLIP (ViT-L/14)	9.44 ± 3.28	13.03 ± 5.56	69.87 ± 1.05	64.90 ± 2.05	70.16 ± 1.05	62.43 ± 2.99		
Language						English	Spanish	Others
	CLIP (ViT-B/16)	11.11 ± 2.42	16.97 ± 2.73	67.84 ± 0.90	60.19 ± 3.47	67.88 ± 1.15	61.93 ± 4.57	61.89 ± 2.90
	FairCLIP (ViT-B/16)	7.34 ± 4.63	17.15 ± 11.13	70.08 ± 1.14	62.31 ± 0.96	70.22 ± 1.27	68.47 ± 5.49	63.47 ± 2.13
	CLIP (ViT-L/14)	7.62 ± 5.39	10.84 ± 9.84	67.83 ± 2.92	61.29 ± 1.28	67.77 ± 3.18	63.83 ± 3.20	62.24 ± 2.25
FairCLIP (ViT-L/14)	11.63 ± 3.53	11.05 ± 5.66	68.68 ± 2.26	61.50 ± 2.29	68.71 ± 2.31	65.06 ± 3.96	61.18 ± 1.20	

ities of the downstream models. In order to ensure a fair comparison, we use the same ViT-L/14 architecture for the vision encoders of both models. Unlike CLIP, BLIP2 is a parameter-efficient VL method that keeps the vision encoder frozen and leverages a Q-former to learn the most suitable image representations corresponding to the text. From the natural pre-training results in Table 2, we note that CLIP consistently outperforms BLIP2 in terms of not only overall performance (77.27 vs. 73.81 AUC) but also fairness, as evidenced by the superior ES-AUC scores across all four protected attributes. However, we observe that medical pre-training largely closes this gap, with BLIP2 yielding similar results as CLIP (80.10 vs. 80.27 AUC). This large boost in BLIP2 performance indicates that (1) the parameter-efficient training in BLIP2 is effective for adapting to the medical VL data, and that (2) the clinical notes are useful for guiding BLIP2 to extract meaningful visual features from the frozen vision encoder. Among the medical pre-trained CLIP and BLIP2 models, CLIP exhibits a better performance-fairness trade-off than BLIP2 on the race and gender attributes, whereas BLIP2 yields more favorable results on the ethnicity and language attributes.

In summary, we observe that all VL models exhibit biases, with Asian, Male, Non-Hispanic, and Spanish being the preferred subgroups across the protected attributes of race, gender, ethnicity, and language, respectively. Medical pre-training enhances the performance-fairness trade-off across all attributes except language. Furthermore, different VL pre-training methods exhibit varying strengths, with CLIP outperforming on race and gender, whereas BLIP2 yields superior results on ethnicity and language.

5.3. CLIP vs. FairCLIP

Table 3 compares the zero-shot performance of CLIP against FairCLIP across two different architectures (ViT-

B/16 and ViT-L/14) and four different protected attributes. Both CLIP and FairCLIP are fine-tuned with the pairs of images and clinical notes without supervised information (*i.e.*, labels). Then, the resulting models are evaluated in the classification task. CLIP exhibits notable disparities in group-wise AUC for attributes such as race, gender, ethnicity, and language, indicating the presence of bias in glaucoma detection. Regarding the racial groups, FairCLIP (ViT-B/16) gains a significant improvement in fairness, reducing DPD and DEOdds to 6.07 and 10.50, respectively, while increasing the AUC for the White group to 69.17. In terms of gender disparity, FairCLIP (ViT-B/16) addresses the bias effectively, elevating the female AUC to 66.81 from 64.62, indicating a substantial enhancement in gender. Similar enhancements by FairCLIP are observed in ethnicity and language groups.

Overall, FairCLIP demonstrates a significant improvement over CLIP in terms of both fairness metrics (DPD, DEOdds) as well as ES-AUC and AUC scores across various demographic subgroups. The supplementary material shows more end-to-end fine-tuning results, further validating the effectiveness of FairCLIP. These empirical findings suggest that optimizing the distance between the overall sample distribution and the distribution w.r.t. specific subgroups effectively improves fairness, indicating a promising direction in addressing and mitigating inherent biases.

5.4. Ablation Studies

Clinical Note Summarization: VL models are usually trained with images and captions, which tend to be quite short. In contrast, the clinical notes in our Harvard-FairVLMed dataset are fairly lengthy, capturing a lot more nuanced information. Hence, we summarize these clinical notes in order to align our setup with the standard VL frameworks while retaining essential medical informa-

tion. We use three different types of LLMs for summarization. In addition to the SOTA GPT-4 [54] model, we also use two LLMs specialized for the medical domain – PMC-LLAMA [80] and MED42². We use the following prompt for summarization for all three LLMs: *Summarize the key details, including the presence of glaucoma, from the clinical note within 180 characters.* Figure 3a compares the performance-fairness trade-off of BLIP2 models trained via the three summarized notes against the original notes. Across all protected attributes except language, we note that all three summarization methods yield a consistent improvement in ES-AUC. Particularly, MED42 yields the best performance across race, GPT-4 yields the best performance across gender, whereas PMC-LLAMA yields the best performance across both ethnicity and language. Overall, the medical pre-trained LLMs (PMC-LLAMA and MED42) yield superior results, highlighting the efficacy of domain-specific LLMs for clinical note summarization.

Vision vs. Multimodal Features: In order to decouple the benefits of image and text features, we conduct linear probing on the BLIP2 pre-trained models using either vision-only or (vision + language) features. Table 4 presents the performance-fairness trade-off in terms of ES-AUC. We note that the multimodal features consistently improve the performance-fairness trade-off across all protected attributes except language. This highlights that the VL models make effective use of the clinical textual features, with the most appreciable gains observed on the race attribute.

Natural vs. Medical Vision Encoder: To investigate the impact of different vision encoders on model fairness in BLIP2, we utilize two different pre-trained encoders – 1) CLIP trained on the natural domain, whereas 2) PMC-CLIP trained on the medical domain. The results in Figure 3b reveal that PMC-CLIP outperforms CLIP across all four protected attributes, with the most appreciable gains on the racial subgroups. We note that medical-specific LLM summarizers and vision encoders consistently improve the performance-fairness trade-off of VL models, with the most appreciable improvements across the race attribute.

Comparison with Adversary Fairness: Beutel et al. [8] introduce a fairness approach that employs adversarial loss to prevent the model from inaccurately predicting sensitive attributes. This method aims to ensure that the model predicts the label of an image without relying on its sensitive attributes, thereby reducing bias in classification. Figure 3c shows the performance comparison among CLIP, CLIP with adversarial loss (CLIP w/ Adv), and FairCLIP. The performance of CLIP with adversarial training (CLIP w/ Adv) does not consistently surpass that of the standard CLIP across all attributes. In contrast, FairCLIP consistently outperforms CLIP. This variation in performance can be attributed to the inherent challenges of adversarial train-

Table 4. Impact of vision-only and (vision + language) features on the performance-fairness trade-off of linear probing via BLIP2.

Attribute	V	L	ES-AUC ↑	Group-wise AUC ↑		
Race	✓	✗	73.76	Asian	Black	White
	✓	✓	76.70	82.08	74.35	81.03
Gender	✓	✗	75.22	Female	Male	
	✓	✓	79.15	77.13	83.66	83.38
Ethnicity	✓	✗	77.18	Non-Hispanic	Hispanic	
	✓	✓	75.54	80.28	76.46	73.85
Language	✓	✗	69.88	English	Spanish	Others
	✓	✓	66.35	80.64	83.52	69.36
				82.32	69.89	71.01




Figure 3. Extensive analyses on Harvard-FairVLMed, including (a) Effects of various LLM summarizations on the performance-fairness trade-off of BLIP2, (b) Effects of using pre-trained vision encoders from natural (CLIP) and medical (PMC-CLIP) domains on the performance-fairness trade-off of BLIP2, and (c) Performance comparison among CLIP, CLIP with adversarial loss (CLIP w/ Adv), and FairCLIP.

ing in maintaining equivalent prediction accuracy for each attribute. On the other hand, FairCLIP utilizes Sinkhorn loss, which effectively encourages uniformity in the distribution of all samples relative to the distributions corresponding to each group.

6. Conclusion

Given the critical need for fairness in healthcare-focused settings, we introduce the first fair vision-language medical dataset (*Harvard-FairVLMed*) for studying the fairness of medical VL foundation models. Our comprehensive fairness analysis on Harvard-FairVLMed reveals significant biases in all VL models. To address these biases, we propose FairCLIP, an optimal transport-based approach that effectively balances both performance and fairness. To aid future research in this area, we make our dataset and code publicly available at <https://ophai.hms.harvard.edu/datasets/harvard-fairvmed10k>.

Acknowledgement

This work was supported by NIH R00 EY028631, NIH R21 EY035298, NIH P30 EY003790, Research To Prevent Blindness International Research Collaborators Award, and Alcon Young Investigator Grant. We also acknowledge the generous funding resources provided by NYU Abu Dhabi with code AD131.

²<https://huggingface.co/m42-health/med42-70b>

References

- [1] Peking university international competition on ocular disease intelligent recognition (odir-2019). 1, 3
- [2] Mohamed Abdalla and Benjamin Fine. Hurdles to artificial intelligence deployment: Noise in schemas and “gold” labels. *Radiology: Artificial Intelligence*, 5(2):e220056, 2023. 3
- [3] Parnian Afshar, Shahin Heidarian, Nastaran Enshaei, Farnoosh Naderkhani, Moezedin Javad Rafiee, Anastasia Oikonomou, Farnak Babaki Fard, Kaveh Samimi, Konstantinos N Plataniotis, and Arash Mohammadi. Covid-ct-md, covid-19 computed tomography scan dataset applicable in machine learning and deep learning. *Scientific Data*, 8(1): 121, 2021. 1
- [4] Alekh Agarwal, Alina Beygelzimer, Miroslav Dudík, John Langford, and Hanna Wallach. A reductions approach to fair classification. In *International Conference on Machine Learning*, pages 60–69. PMLR, 2018. 6
- [5] Alekh Agarwal, Miroslav Dudík, and Zhiwei Steven Wu. Fair regression: Quantitative definitions and reduction-based algorithms. In *International Conference on Machine Learning*, pages 120–129. PMLR, 2019. 6
- [6] Yakoub Bazi, Mohamad Mahmoud Al Rahhal, Laila Bashmal, and Mansour Zuair. Vision–language model for visual question answering in medical imagery. *Bioengineering*, 10(3):380, 2023. 3
- [7] Sebastian Benthall and Bruce D Haynes. Racial categories in machine learning. In *Proceedings of the conference on fairness, accountability, and transparency*, pages 289–298, 2019. 2
- [8] Alex Beutel, Jilin Chen, Zhe Zhao, and Ed H Chi. Data decisions and theoretical implications when adversarially learning fair representations. *arXiv preprint arXiv:1707.00075*, 2017. 2, 8
- [9] Martim Brandao. Age and gender bias in pedestrian detection algorithms. *arXiv preprint arXiv:1906.10490*, 2019. 2
- [10] Alexander Brown, Nenad Tomasev, Jan Freyberg, Yuan Liu, Alan Karthikesalingam, and Jessica Schrouff. Detecting and preventing shortcut learning for fair medical ai using shortcut testing (short). *arXiv preprint arXiv:2207.10384*, 2022. 2
- [11] Joy Buolamwini and Timnit Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*, pages 77–91. PMLR, 2018. 2
- [12] Aurelia Bustos, Antonio Pertusa, Jose-Maria Salinas, and Maria De La Iglesia-Vaya. Padchest: A large chest x-ray image dataset with multi-label annotated reports. *Medical image analysis*, 66:101797, 2020. 1, 3
- [13] Dina Demner-Fushman, Marc D Kohli, Marc B Rosenman, Sonya E Shooshan, Laritza Rodriguez, Sameer Antani, George R Thoma, and Clement J McDonald. Preparing a collection of radiology examinations for distribution and retrieval. *Journal of the American Medical Informatics Association*, 23(2):304–310, 2016. 3
- [14] Sina Farsiu, Stephanie J Chiu, Rachelle V O’Connell, Francisco A Folgar, Eric Yuan, Joseph A Izatt, Cynthia A Toth, Age-Related Eye Disease Study 2 Ancillary Spectral Domain Optical Coherence Tomography Study Group, et al. Quantitative classification of eyes with and without intermediate age-related macular degeneration using optical coherence tomography. *Ophthalmology*, 121(1):162–172, 2014. 1
- [15] Ben Glocker, Charles Jones, Mélanie Bernhardt, and Stefan Winzeck. Algorithmic encoding of protected characteristics in chest x-ray disease detection models. *Ebiomedicine*, 89, 2023. 1
- [16] Matthew Groh, Caleb Harris, Luis Soenksen, Felix Lau, Rachel Han, Aerin Kim, Arash Koochek, and Omar Badri. Evaluating deep neural networks trained on clinical images in dermatology with the fitzpatrick 17k dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1820–1828, 2021. 1
- [17] Priya Gupta, Di Zhao, Eliseo Guallar, Fang Ko, Michael V Boland, and David S Friedman. Prevalence of glaucoma in the united states: the 2005–2008 national health and nutrition examination survey. *Investigative ophthalmology & visual science*, 57(6):2905–2913, 2016. 2
- [18] Alex Hanna, Emily Denton, Andrew Smart, and Jamila Smith-Loud. Towards a critical race methodology in algorithmic fairness. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, pages 501–512, 2020. 2
- [19] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16000–16009, 2022. 5
- [20] Xuehai He, Zhuo Cai, Wenlan Wei, Yichen Zhang, Luntian Mou, Eric Xing, and Pengtao Xie. Pathological visual question answering. *arXiv preprint arXiv:2010.12435*, 2020. 3
- [21] Zhi Huang, Federico Bianchi, Mert Yuksekgonul, Thomas J Montine, and James Zou. A visual–language foundation model for pathology image analysis using medical twitter. *Nature medicine*, pages 1–10, 2023. 1, 3
- [22] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. pmlr, 2015. 5
- [23] Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silvana Ciurea-Ilcus, Chris Chute, Henrik Marklund, Behzad Haghgoo, Robyn Ball, Katie Shpanskaya, et al. Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In *Proceedings of the AAAI conference on artificial intelligence*, pages 590–597, 2019. 3
- [24] Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silvana Ciurea-Ilcus, Chris Chute, Henrik Marklund, Behzad Haghgoo, Robyn Ball, Katie Shpanskaya, et al. Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In *Proceedings of the AAAI conference on artificial intelligence*, pages 590–597, 2019. 1, 2
- [25] Baoyu Jing, Pengtao Xie, and Eric Xing. On the automatic generation of medical imaging reports. *arXiv preprint arXiv:1711.08195*, 2017. 3

- [26] Alistair EW Johnson, Tom J Pollard, Seth J Berkowitz, Nathaniel R Greenbaum, Matthew P Lungren, Chih-ying Deng, Roger G Mark, and Steven Horng. Mimic-cxr, a de-identified publicly available database of chest radiographs with free-text reports. *Scientific data*, 6(1):317, 2019. [2](#)
- [27] Alistair EW 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. Mimic-cxr-jpg, a large publicly available database of labeled chest radiographs. *arXiv preprint arXiv:1901.07042*, 2019. [1](#), [3](#)
- [28] Neil Joshi and Phil Burlina. Ai fairness via domain adaptation. *arXiv preprint arXiv:2104.01109*, 2021. [2](#)
- [29] Chris Kelly, Luhui Hu, Jiayin Hu, Yu Tian, Deshun Yang, Bang Yang, Cindy Yang, Zihao Li, Zaoshan Huang, and Yuexian Zou. Visiongpt-3d: A generalized multimodal agent for enhanced 3d vision understanding. *arXiv preprint arXiv:2403.09530*, 2024. [1](#)
- [30] Chris Kelly, Luhui Hu, Bang Yang, Yu Tian, Deshun Yang, Cindy Yang, Zaoshan Huang, Zihao Li, Jiayin Hu, and Yuexian Zou. Visiongpt: Vision-language understanding agent using generalized multimodal framework. *arXiv preprint arXiv:2403.09027*, 2024. [1](#)
- [31] Muhammad Osama Khan, Muhammad Muneeb Afzal, Shujaat Mirza, and Yi Fang. How fair are medical imaging foundation models? In *Machine Learning for Health (ML4H)*, pages 217–231. PMLR, 2023. [1](#)
- [32] Shunsuke Kogure, Kai Watabe, Ryosuke Yamada, Yoshimitsu Aoki, Akio Nakamura, and Hirokatsu Kataoka. Age should not matter: Towards more accurate pedestrian detection via self-training. In *Computer Sciences & Mathematics Forum*, page 11. MDPI, 2022. [2](#)
- [33] Shivansh Kohli, Utkarsh Verma, Vinay V Kirpalani, and Ramamoorthy Srinath. Dermatobot: an image processing enabled chatbot for diagnosis and tele-remedy of skin diseases. In *2022 3rd International Conference for Emerging Technology (INCET)*, pages 1–5. IEEE, 2022. [3](#)
- [34] Oleksandr Kovalyk, Juan Morales-Sánchez, Rafael Verdú-Monedero, Inmaculada Sellés-Navarro, Ana Palazón-Cabanes, and José-Luis Sancho-Gómez. Papila: Dataset with fundus images and clinical data of both eyes of the same patient for glaucoma assessment. *Scientific Data*, 9(1):291, 2022. [1](#), [3](#)
- [35] Jason J Lau, Soumya Gayen, Asma Ben Abacha, and Dina Demner-Fushman. A dataset of clinically generated visual questions and answers about radiology images. *Scientific data*, 5(1):1–10, 2018. [3](#)
- [36] Yiming Lei, Zilong Li, Yan Shen, Junping Zhang, and Hongming Shan. Clip-lung: Textual knowledge-guided lung nodule malignancy prediction. *arXiv preprint arXiv:2304.08013*, 2023. [3](#)
- [37] Gang Li, Jian Li, Shanshan Zhang, and Jian Yang. Learning hierarchical graph for occluded pedestrian detection. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 1597–1605, 2020. [2](#)
- [38] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*, pages 12888–12900. PMLR, 2022. [3](#)
- [39] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023. [2](#), [3](#), [5](#)
- [40] Weixiong Lin, Ziheng Zhao, Xiaoman Zhang, Chaoyi Wu, Ya Zhang, Yanfeng Wang, and Weidi Xie. Pmc-clip: Contrastive language-image pre-training using biomedical documents. *arXiv preprint arXiv:2303.07240*, 2023. [3](#)
- [41] Bo Liu, Li-Ming Zhan, Li Xu, Lin Ma, Yan Yang, and Xiao-Ming Wu. Slake: A semantically-labeled knowledge-enhanced dataset for medical visual question answering. In *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, pages 1650–1654. IEEE, 2021. [3](#)
- [42] Guanxiong Liu, Tzu-Ming Harry Hsu, Matthew McDermott, Willie Boag, Wei-Hung Weng, Peter Szolovits, and Marzyeh Ghassemi. Clinically accurate chest x-ray report generation. In *Machine Learning for Healthcare Conference*, pages 249–269. PMLR, 2019. [3](#)
- [43] Jie Liu, Yixiao Zhang, Jie-Neng Chen, Junfei Xiao, Yongyi Lu, Bennett A Landman, Yixuan Yuan, Alan Yuille, Yucheng Tang, and Zongwei Zhou. Clip-driven universal model for organ segmentation and tumor detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 21152–21164, 2023. [3](#)
- [44] Renqian Luo, Liai Sun, Yingce Xia, Tao Qin, Sheng Zhang, Hoifung Poon, and Tie-Yan Liu. Biogpt: generative pre-trained transformer for biomedical text generation and mining. *Briefings in Bioinformatics*, 23(6):bbac409, 2022. [1](#)
- [45] Yan Luo, Min Shi, Yu Tian, Tobias Elze, and Mengyu Wang. Harvard glaucoma detection and progression: A multimodal multitask dataset and generalization-reinforced semi-supervised learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 20471–20482, 2023. [1](#), [3](#)
- [46] Yan Luo, Min Shi, Yu Tian, Tobias Elze, and Mengyu Wang. Harvard glaucoma detection and progression: A multimodal multitask dataset and generalization-reinforced semi-supervised learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20471–20482, 2023. [2](#)
- [47] Yan Luo, Yu Tian, Min Shi, Tobias Elze, and Mengyu Wang. Eye fairness: A large-scale 3d imaging dataset for equitable eye diseases screening and fair identity scaling. *arXiv preprint arXiv:2310.02492*, 2023. [2](#), [3](#)
- [48] Yan Luo, Yu Tian, Min Shi, Louis R. Pasquale, Lucy Q. Shen, Nazlee Zebardast, Tobias Elze, and Mengyu Wang. Harvard glaucoma fairness: A retinal nerve disease dataset for fairness learning and fair identity normalization. *IEEE Transactions on Medical Imaging*, pages 1–1, 2024. [3](#), [6](#)
- [49] Ricards Marcinkevics, Ece Ozkan, and Julia E Vogt. Debiasing deep chest x-ray classifiers using intra-and post-processing methods. In *Machine Learning for Healthcare Conference*, pages 504–536. PMLR, 2022. [2](#)
- [50] Masoud Monajatipoor, Mozhddeh Rouhsedaghat, Liunan Harold Li, C-C Jay Kuo, Aichi Chien, and Kai-Wei

- Chang, Berthop: An effective vision-and-language model for chest x-ray disease diagnosis. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 725–734. Springer, 2022. 3
- [51] Usman Naseem, Matloob Khushi, and Jinman Kim. Vision-language transformer for interpretable pathology visual question answering. *IEEE Journal of Biomedical and Health Informatics*, 27(4):1681–1690, 2022. 3
- [52] Hoang TN Nguyen, Dong Nie, Taivanbat Badamdorj, Yujie Liu, Lingzi Hong, Jason Truong, and Li Cheng. Eddie-transformer: Enriched disease embedding transformer for x-ray report generation. In *2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI)*, pages 1–5. IEEE, 2022. 3
- [53] Saeed Niksaz and Fahimeh Ghasemian. Improving chest x-ray report generation by leveraging text of similar images. *Available at SSRN 4211036*, 2023. 3
- [54] OpenAI. Gpt-4 technical report, 2023. 8
- [55] Sungho Park, Jewook Lee, Pilhyeon Lee, Sunhee Hwang, Dohyung Kim, and Hyeran Byun. Fair contrastive learning for facial attribute classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10389–10398, 2022. 2
- [56] Obioma Pelka, Sven Koitka, Johannes Rückert, Felix Nensa, and Christoph M Friedrich. Radiology objects in context (roco): a multimodal image dataset. In *Intravascular Imaging and Computer Assisted Stenting and Large-Scale Annotation of Biomedical Data and Expert Label Synthesis: 7th Joint International Workshop, CVII-STENT 2018 and Third International Workshop, LABELS 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 16, 2018, Proceedings 3*, pages 180–189. Springer, 2018. 3
- [57] Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019. 5
- [58] Geoff Pleiss, Manish Raghavan, Felix Wu, Jon Kleinberg, and Kilian Q Weinberger. On fairness and calibration. *Advances in neural information processing systems*, 30, 2017. 2
- [59] Esther Puyol-Antón, Bram Ruijsink, Stefan K Piechnik, Stefan Neubauer, Steffen E Petersen, Reza Razavi, and Andrew P King. Fairness in cardiac mr image analysis: an investigation of bias due to data imbalance in deep learning based segmentation. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part III 24*, pages 413–423. Springer, 2021. 2
- [60] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 4, 5
- [61] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 2, 3
- [62] Vikram V Ramaswamy, Sunnie SY Kim, and Olga Russakovsky. Fair attribute classification through latent space de-biasing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9301–9310, 2021. 2
- [63] Yuji Roh, Kangwook Lee, Steven Whang, and Changho Suh. Fr-train: A mutual information-based approach to fair and robust training. In *International Conference on Machine Learning*, pages 8147–8157. PMLR, 2020. 2
- [64] Mhd Hasan Sarhan, Nassir Navab, Abouzar Eslami, and Shadi Albarqouni. Fairness by learning orthogonal disentangled representations. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIX 16*, pages 746–761. Springer, 2020. 2
- [65] Laleh Seyyed-Kalantari, Guanxiong Liu, Matthew McDermott, Irene Y Chen, and Marzyeh Ghassemi. Chexclusion: Fairness gaps in deep chest x-ray classifiers. In *BIOCOMPUTING 2021: proceedings of the Pacific symposium*, pages 232–243. World Scientific, 2020. 2
- [66] Yahya Shaikh, Fei Yu, and Anne L Coleman. Burden of undetected and untreated glaucoma in the united states. *American journal of ophthalmology*, 158(6):1121–1129, 2014. 2
- [67] Min Shi, Anagha Lokhande, Mojtaba S Fazli, Vishal Sharma, Yu Tian, Yan Luo, Louis R Pasquale, Tobias Elze, Michael V Boland, Nazlee Zebardast, et al. Artifact-tolerant clustering-guided contrastive embedding learning for ophthalmic images in glaucoma. *IEEE Journal of Biomedical and Health Informatics*, 2023. 2
- [68] Min Shi, Yan Luo, Yu Tian, Lucy Q Shen, Tobias Elze, Nazlee Zebardast, Mohammad Eslami, Saber Kazeminasab, Michael V Boland, David S Friedman, et al. Equitable artificial intelligence for glaucoma screening with fair identity normalization. *medRxiv*, pages 2023–12, 2023. 2
- [69] Min Shi, Jessica A Sun, Anagha Lokhande, Yu Tian, Yan Luo, Tobias Elze, Lucy Q Shen, and Mengyu Wang. Artifact correction in retinal nerve fiber layer thickness maps using deep learning and its clinical utility in glaucoma. *Translational Vision Science & Technology*, 12(11):12–12, 2023. 2
- [70] Min Shi, Yu Tian, Yan Luo, Tobias Elze, and Mengyu Wang. Rnfft2vec: Artifact-corrected representation learning for retinal nerve fiber layer thickness maps. *Medical Image Analysis*, page 103110, 2024. 2
- [71] Sanjay Subramanian, Lucy Lu Wang, Sachin Mehta, Ben Bogin, Madeleine van Zuynen, Sravanthi Parasa, Sameer Singh, Matt Gardner, and Hannaneh Hajishirzi. Medicat: A dataset of medical images, captions, and textual references. *arXiv preprint arXiv:2010.06000*, 2020. 3
- [72] Yih-Chung Tham, Xiang Li, Tien Y Wong, Harry A Quigley, Tin Aung, and Ching-Yu Cheng. Global prevalence of glaucoma and projections of glaucoma burden through 2040: a systematic review and meta-analysis. *Ophthalmology*, 121(11):2081–2090, 2014. 2

- [73] Jiang Tian, Cong Li, Zhongchao Shi, and Feiyu Xu. A diagnostic report generator from ct volumes on liver tumor with semi-supervised attention mechanism. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2018: 21st International Conference, Granada, Spain, September 16–20, 2018, Proceedings, Part II 11*, pages 702–710. Springer, 2018. 3
- [74] Yu Tian, Min Shi, Yan Luo, Ava Kouhana, Tobias Elze, and Mengyu Wang. Fairseg: A large-scale medical image segmentation dataset for fairness learning with fair error-bound scaling. *arXiv preprint arXiv:2311.02189*, 2023. 2
- [75] Philipp Tschandl, Cliff Rosendahl, and Harald Kittler. The ham10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Scientific data*, 5(1):1–9, 2018. 1
- [76] Mei Wang and Weihong Deng. Mitigating bias in face recognition using skewness-aware reinforcement learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9322–9331, 2020. 2
- [77] Zeyu Wang, Klint Qinami, Ioannis Christos Karakozis, Kyle Genova, Prem Nair, Kenji Hata, and Olga Russakovsky. Towards fairness in visual recognition: Effective strategies for bias mitigation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8919–8928, 2020. 2
- [78] Zifeng Wang, Zhenbang Wu, Dinesh Agarwal, and Jimeng Sun. Medclip: Contrastive learning from unpaired medical images and text. *arXiv preprint arXiv:2210.10163*, 2022. 1, 3
- [79] Chaoyi Wu, Xiaoman Zhang, Ya Zhang, Yanfeng Wang, and Weidi Xie. Medklip: Medical knowledge enhanced language-image pre-training. *medRxiv*, pages 2023–01, 2023. 3
- [80] Chaoyi Wu, Xiaoman Zhang, Ya Zhang, Yanfeng Wang, and Weidi Xie. Pmc-llama: Further finetuning llama on medical papers. *arXiv preprint arXiv:2304.14454*, 2023. 8
- [81] Fan Wu, Haiqiong Yang, Linlin Peng, Zongkai Lian, Mingxin Li, Gang Qu, Shancheng Jiang, and Yu Han. Agnet: Automatic generation network for skin imaging reports. *Computers in biology and medicine*, 141:105037, 2022. 3
- [82] Yawen Wu, Dewen Zeng, Xiaowei Xu, Yiyu Shi, and Jingtong Hu. Fairprune: Achieving fairness through pruning for dermatological disease diagnosis. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 743–753. Springer, 2022. 2
- [83] Bradley T Wyman, Danielle J Harvey, Karen Crawford, Matt A Bernstein, Owen Carmichael, Patricia E Cole, Paul K Crane, Charles DeCarli, Nick C Fox, Jeffrey L Gunter, et al. Standardization of analysis sets for reporting results from adni mri data. *Alzheimer’s & Dementia*, 9(3):332–337, 2013. 1
- [84] Shawn Xu, Lin Yang, Christopher Kelly, Marcin Sieniek, Timo Kohlberger, Martin Ma, Wei-Hung Weng, Attila Kiraly, Sahar Kazemzadeh, Zakkai Melamed, et al. Elixr: Towards a general purpose x-ray artificial intelligence system through alignment of large language models and radiology vision encoders. *arXiv preprint arXiv:2308.01317*, 2023. 1
- [85] Deshun Yang, Luhui Hu, Yu Tian, Zihao Li, Chris Kelly, Bang Yang, Cindy Yang, and Yuexian Zou. Worldgpt: A sora-inspired video ai agent as rich world models from text and image inputs. *arXiv preprint arXiv:2403.07944*, 2024. 1
- [86] Shaokang Yang, Jianwei Niu, Jiyan Wu, Yong Wang, Xuefeng Liu, and Qingfeng Li. Automatic ultrasound image report generation with adaptive multimodal attention mechanism. *Neurocomputing*, 427:40–49, 2021. 3
- [87] Kihyun You, Jawook Gu, Jiyeon Ham, Beomhee Park, Jiho Kim, Eun K Hong, Woonhyuk Baek, and Byungseok Roh. Cxr-clip: Toward large scale chest x-ray language-image pre-training. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 101–111. Springer, 2023. 3
- [88] Yang You, Igor Gitman, and Boris Ginsburg. Large batch training of convolutional networks. *arXiv preprint arXiv:1708.03888*, 2017. 5
- [89] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rogriguez, and Krishna P Gummadi. Fairness constraints: Mechanisms for fair classification. In *Artificial intelligence and statistics*, pages 962–970. PMLR, 2017. 2
- [90] Juan M Zambrano Chaves, Akshay S Chaudhari, Andrew L Wentland, Arjun D Desai, Imon Banerjee, Robert D Boutin, David J Maron, Fatima Rodriguez, Alexander T Sandhu, R Brooke Jeffrey, et al. Opportunistic assessment of ischemic heart disease risk using abdominopelvic computed tomography and medical record data: a multimodal explainable artificial intelligence approach. *medRxiv*, pages 2021–01, 2021. 1
- [91] Brian Hu Zhang, Blake Lemoine, and Margaret Mitchell. Mitigating unwanted biases with adversarial learning. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pages 335–340, 2018. 2
- [92] Sheng Zhang, Yanbo Xu, Naoto Usuyama, Jaspreet Bagga, Robert Tinn, Sam Preston, Rajesh Rao, Mu Wei, Naveen Valluri, Cliff Wong, et al. Large-scale domain-specific pre-training for biomedical vision-language processing. *arXiv preprint arXiv:2303.00915*, 2023. 3
- [93] Yi Zhang and Jitao Sang. Towards accuracy-fairness paradox: Adversarial example-based data augmentation for visual debiasing. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 4346–4354, 2020. 2
- [94] Zizhao Zhang, Pingjun Chen, Mason McGough, Fuyong Xing, Chunbao Wang, Marilyn Bui, Yuanpu Xie, Manish Sapkota, Lei Cui, Jasreman Dhillon, et al. Pathologist-level interpretable whole-slide cancer diagnosis with deep learning. *Nature Machine Intelligence*, 1(5):236–245, 2019. 3
- [95] Yue Zhao, Ming Tian, Jing Jin, Qiucheng Wang, Jian Song, and Yi Shen. An automatically thyroid nodules feature extraction and description network for ultrasound images. In *2021 IEEE International Ultrasonics Symposium (IUS)*, pages 1–4. IEEE, 2021. 3
- [96] Qihang Zhou, Guansong Pang, Yu Tian, Shibo He, and Jiming Chen. Anomalyclip: Object-agnostic prompt learning for zero-shot anomaly detection. *arXiv preprint arXiv:2310.18961*, 2023. 3

- [97] Yuyin Zhou, Shih-Cheng Huang, Jason Alan Fries, Alaa Youssef, Timothy J Amrhein, Marcello Chang, Imon Banerjee, Daniel Rubin, Lei Xing, Nigam Shah, et al. Radfusion: Benchmarking performance and fairness for multimodal pulmonary embolism detection from ct and ehr. *arXiv preprint arXiv:2111.11665*, 2021. [2](#)
- [98] Dominik Zietlow, Michael Lohaus, Guha Balakrishnan, Matthäus Kleindessner, Francesco Locatello, Bernhard Schölkopf, and Chris Russell. Leveling down in computer vision: Pareto inefficiencies in fair deep classifiers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10410–10421, 2022. [2](#)