

# Markedness in Visual Semantic AI

Robert Wolfe  
University of Washington  
Seattle, WA, USA  
rwolfe3@uw.edu

Aylin Caliskan  
University of Washington  
Seattle, WA, USA  
aylin@uw.edu

## ABSTRACT

We evaluate the state-of-the-art multimodal "visual semantic" model CLIP ("Contrastive Language Image Pretraining") for biases related to the marking of age, gender, and race or ethnicity. Given the option to label an image as "a photo of a person" or to select a label denoting race or ethnicity, CLIP chooses the "person" label 47.9% of the time for White individuals, compared with 5.0% or less for individuals who are Black, East Asian, Southeast Asian, Indian, or Latino or Hispanic. The model is also more likely to rank the unmarked "person" label higher than labels denoting gender for Male individuals (26.7% of the time) vs. Female individuals (15.2% of the time). Age also affects whether an individual is marked by the model: Female individuals under the age of 20 are more likely than Male individuals to be marked with a gender label, but less likely to be marked with an age label, while Female individuals over the age of 40 are more likely to be marked based on age than Male individuals. We trace our results back to the CLIP embedding space by examining the self-similarity (mean pairwise cosine similarity) for each social group, where higher self-similarity denotes greater attention directed by CLIP to the shared characteristics (*i.e.*, age, race, or gender) of the social group. The results indicate that, as age increases, the self-similarity of representations of Female individuals increases at a higher rate than for Male individuals, with the disparity most pronounced at the "more than 70" age range. Six of the ten least self-similar social groups are individuals who are White and Male, while all ten of the most self-similar social groups are individuals under the age of 10 or over the age of 70, and six of the ten are Female individuals. Our results yield evidence that bias in CLIP is intersectional: existing biases of self-similarity and markedness between Male and Female gender groups are further exacerbated when the groups compared are individuals who are White and Male and individuals who are Black and Female. CLIP is an English-language model trained on internet content gathered based on a query list generated from an American website (Wikipedia), and results indicate that CLIP reflects the biases of the language and society which produced this training data.

## CCS CONCEPTS

• **Computing methodologies** → **Artificial intelligence.**



This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike International 4.0 License.

FACCT '22, June 21–24, 2022, Seoul, Republic of Korea  
© 2022 Copyright held by the owner/author(s).  
ACM ISBN 978-1-4503-9352-2/22/06.  
<https://doi.org/10.1145/3531146.3533183>

## KEYWORDS

multimodal, bias in AI, visual semantics, language-and-vision AI, markedness, age bias

### ACM Reference Format:

Robert Wolfe and Aylin Caliskan. 2022. Markedness in Visual Semantic AI. In *2022 ACM Conference on Fairness, Accountability, and Transparency (FACCT '22)*, June 21–24, 2022, Seoul, Republic of Korea. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3531146.3533183>

## 1 INTRODUCTION

Recent progress in multimodal artificial intelligence (AI) has produced the first "zero-shot" language-and-vision model, known as CLIP ("Contrastive Language Image Pretraining"), which allows for the definition of image classes in natural language and achieves performance competitive with the state of the art on datasets on which it was not trained [40]. The advances achieved by CLIP have spawned a new wave of interest and innovation in "visual semantic" AI, which combines language and image representations in the same embedding space. The past year has seen the development of numerous proprietary models similar to CLIP [28, 53], which are likely to serve as important components in the future foundation of the internet [38]. Like word embeddings and language models, visual semantic AI is trained using human-authored language on an internet-scale dataset: CLIP's training data is composed of 400 million pairs of images and associated text scraped from the English-language internet [40]. Prior work has uncovered semantic biases in CLIP associating immigrants and religious minorities with unpleasant stereotypes [22], and gender biases related to underrepresentation of women when using the model for image retrieval [55]. In this research, we systematically evaluate CLIP for a previously unobserved perceptual bias: the bias of who is marked, and based on what socially defined characteristics (gender, race, and age).

Our contributions are twofold. First, we find disparities in the self-similarity (mean pairwise cosine similarity) of embedded images along the axes of gender, age, and race or ethnicity. Concretely, embedded images of White individuals in the FairFace dataset have self-similarity of 0.573, compared to self-similarity of  $> .62$  for Indian, East Asian, or Southeast Asian individuals; images of Male individuals have mean self-similarity of 0.566, compared to 0.592 for images of Female individuals; and images of individuals between 40 and 49 have mean cosine similarity of 0.573, compared to 0.623 for individuals more than 70 and 0.692 for individuals between 0 and 2. These results are also intersectional: of 126 social groups formed based on the intersection of gender, age, and race or ethnicity, the ten least self-similar social groups are all Male, and six of the ten are White. The ten most self-similar groups all reflect people either under the age of 10 or over the age of 70, and six of the ten are Female.

Second, we show that CLIP prefers to label images of members of some social groups based on demographic characteristics such as race, gender, and age, and to omit those labels when describing other social groups. We prompt CLIP to rank the text label "a photo of a person" against sets of text labels related to gender, age, and race or ethnicity based on the FairFace dataset [29]. If the model ranks "a photo of a person" higher than any of the FairFace gender labels, we consider that image to be unmarked based on gender. We repeat the same experiment using the set of age labels and the set of race or ethnicity labels from FairFace. CLIP ranks the "a photo of a person" label highest 47.9% of the time for images of White individuals, and less than 16.9% of the time for all other racial or ethnic groups; CLIP ranks "a photo of a person" highest 26.6% of the time for images of Male individuals, and 15.2% of the time for Female individuals; and CLIP ranks "a photo of a person" highest 59.0% of the time for images of individuals between 30 and 39, but only 6.7% of the time for images of individuals older than 70. Differences are also intersectional, as Female individuals are marked according to their gender at younger ages than Male individuals, and Female individuals who are more than 40 years old are more likely to be marked based on age than Male individuals at the same age. The disparity is greater in almost all cases between individuals who are White and Male and individuals who are Black and Female.

This form of AI bias is likely to become increasingly important as multimodal AI models use natural language supervision [40] to unify different forms of information, such as vision, language, speech, and video, into joint conceptual embedding spaces [38]. Our findings indicate that combining modalities results in previously unobserved forms of representational bias, which may amplify biases existing in single-modality embedding spaces. We make our code available at [https://github.com/wolferobert3/visual\\_semantic\\_markedness](https://github.com/wolferobert3/visual_semantic_markedness).

## 2 RELATED WORK

This research draws on multiple strands of prior work on cultural markedness, CLIP and multimodal AI, and bias in AI.

**Markedness.** Visual semantic models like CLIP are designed to detect patterns in unstructured visual information, and to categorize those patterns in ways that are interpretable to humans in natural language [40]. However, how different groups of people form categories, and what categories have utility, varies widely between communities [5]. While the biases of an advantaged social group may manifest explicitly in the expression of cultural stereotypes, they can also become embedded in the structure of language [8]. One way in which this occurs is in which words and contexts are "unmarked" in language, and more natural, and which are "marked" as less natural to the speaker [17].

The concept of "markedness" originates in structural linguistics. Trubetzkoy [54] use the term to refer to phonological distinctions in language, wherein one in a pair of opposing phonemes has a mark, and the other lacks the mark. Jakobson [27] expands this to semantics, and asserts that linguistic opposites can be characterized by the presence or absence of some attribute, which is marked in a secondary form (*i.e.*, a form lower in a linguistic hierarchy), but not in the dominant form (higher in the hierarchy). For example,

the names of male animals (such as lion) are unmarked, while the names of female animals (lioness) are marked. Greenberg [23] describes markedness as a consequence of the frequency with which a linguistic construction occurs. Dressler [17] describes markedness as linguistic unnaturalness, or as morphological difficulty, while Givón [21] finds that a marked category is cognitively more complex for the speaker.

Markedness also has sociological implications and has been studied as a sociological phenomenon. Waugh [58] proposes a semiotic theory of markedness which suggests that marginalized social categories such as homosexuality are marked, while categories such as heterosexuality are unmarked. Battistella et al. [2] generalizes this and shows that more culturally dominant terms are more broadly defined and are thus "unmarked," while "marked" terms are more narrowly defined. Mayerthaler [35] notes that that which is unmarked agrees with "the typical attributes of the speaker." Finally, Tannen [51] notes that that which is culturally unmarked "carries the meaning that goes without saying," while that which is marked is unusual, and must be explained. Our work evaluates whether CLIP has learned, as Battistella et al. [2] puts it, an "evaluative superstructure" of language which influences how the model embeds unstructured visual information.

**Intersectionality.** Intersectionality refers to the ways in which a person's demographic and socioeconomic characteristics combine to multiply marginalize or multiply advantage the person [10]. Crenshaw [11] introduced intersectionality to describe the multiple marginalization of black women, and the theory is used to understand the effects of marginalization along many sociocultural axes, including race, gender, sexual orientation, age, disability status, and class [10]. Intersectional biases associated with members of multiply marginalized individuals may not be associated with all people who share only one of these constituent characteristics [20]. Intersectional marginalization also compounds the effects of biases associated with an individual's constituent identities [12]. Intersectional biases have been discovered in both computer vision [6, 49] and in language models [25, 34, 47, 50]. We examine biases related to race or ethnicity, gender, and age in this research, and in some cases characterize results in terms of intersectional effects.

**CLIP and Multimodal AI.** CLIP is the first multimodal AI model to learn "transferable" visual representations, meaning that it performs at near state-of-the-art for image ranking, retrieval, and classification across many computer vision evaluation datasets, without explicitly training on the training data for those datasets. Before CLIP, the state-of-the-art for zero-shot image classification on ImageNet [13] was 11.6% accuracy; CLIP improved this to 76.5% [40], realizing a significant advance for semi-supervised computer vision. CLIP jointly pretrains a language model (a smaller version of the GPT-2 architecture [41]) and a computer vision model (either a ResNet [26] or a Vision Transformer [16]), and projects the representations formed by each model independently into a joint language-and-vision embedding space. The model's training objective is to maximize the cosine similarity between a projected image and its projected natural language caption, while minimizing the similarity between the image and all of the other projected captions in the batch [40], a training objective referred to as contrastive learning [40, 52]. In addition to learning transferable visual features, CLIP has also been shown to learn highly semantic word and

sentence representations which set or match state-of-the-art on some intrinsic evaluations [62]. We report results on the CLIP-ViT-Base-Patch32 model, which was downloaded more than one million times in the month prior to April 26, 2022 from the Transformers library, and constituted more than 98% of CLIP downloads from the library during that month [60].

Visual semantic embedding spaces have their origin in the research of Socher et al. [48], who designed an approach for zero-shot transfer between image and language representations. Frome et al. [19] subsequently embed images and text in a joint embedding space in DeViSe, allowing limited generalization to unseen classes. CLIP builds on recent work in multimodal transformer models, including that of Lu et al. [33], who introduce ViLBERT, a multimodal model which extended the BERT masked language model of Devlin et al. [14] to train joint "visiolinguistic" representations on the Conceptual Captions dataset [46] using co-attentional transformer layers. Li et al. [32] introduce VisualBERT, a multimodal model which implicitly aligns language and images using the attention weights of a transformer and which is capable of grounding elements of language to specific regions of an image. More recently, Zhang et al. [64] employ contrastive learning with an image encoder (ResNet 50) and a contextualizing language model (BERT) to train a medical image classifier known as ConVIRT, and Jia et al. [28] employ a similar design with ALIGN, an image classifier trained on contrastive loss (normalized softmax) between a BERT language model and an EfficientNet-L2 [63]. Wang et al. [57] introduce SimVLM, a multimodal encoder-decoder transformer trained only on language modeling loss, rather than contrastive loss. Tiwary [53] introduce Turing Bletchley, a multimodal model trained on contrastive loss and capable of performing language-vision tasks in 94 languages. Most recently, Mu et al. [36] introduce SLIP, a model which adds view-based self-supervision of images to the contrastive learning training objective.

**Bias in AI.** Unsupervised and semi-supervised pretraining is known to encode social biases inherent in the training data [3, 7]. While CLIP is multimodal in that it combines image and text representations into a joint embedding space, the transformers used to form this space are a language model and a computer vision model [40]. Accordingly, we review biases observed in computer vision, language models, and multimodal AI.

**Bias in Computer Vision.** Buolamwini and Gebru [6] show that three state-of-the-art AI facial recognition systems fail to detect the faces of women with darker skin, and that common facial recognition benchmark datasets contain primarily images of people with lighter skin, where skin tone is defined based on the Fitzpatrick dermatological skin type classification system [18]. Wilson et al. [59] show that state-of-the-art object detection systems exhibit better performance for people with lighter skin than for people with darker skin. More recently, Kim et al. [30] find that commercial emotion recognition systems detect emotion least accurately in older adults,<sup>1</sup> while Park et al. [39] examine 92 face datasets used to train facial analysis systems, and show that less than half include at

least one image of a person older than 65, and that only one dataset includes an image of a person older than 85. Steed and Caliskan [49] show that Image GPT [9] generates images which reflect harmful human social biases, such as sexualized pictures of women, and that it associates members of minority racial groups with weapons. Wang et al. [56] find that even balancing the training dataset based on gender does not necessarily prevent deep learning models from learning gender biases.

**Bias in Language Models.** Sheng et al. [47] show that the text output of language models such as GPT-2 [41] demonstrates low "regard" for women and sexual minorities. In developing a benchmark to quantify both bias and language modeling performance based on model output, Nadeem et al. [37] find that larger language models (*i.e.*, those with more parameters) exhibit better language performance and prefer stereotypical human biases more than do smaller models. Guo and Caliskan [25] adapt the Word Embedding Association Test (WEAT) of Caliskan et al. [7] to contextualized word embeddings by treating contextualization as a random effect. They find that word embeddings formed by language models such as BERT [14] encode a range of social biases, including biases related to gender and race, and that the magnitude of intersectional bias is greater than for bias based solely on gender or race. May et al. [34] measure bias in sentence vectors formed by language models and find evidence of sentence-level racial and gender biases. Recent research examines biases of representational similarity in language models. Wolfe and Caliskan [61] find that underrepresentation in the text training corpora of language models results in contextualized word embeddings which are more self-similar across identical contexts, yet undergo more change in the model, suggesting that language models generalize poorly to less frequently observed groups, and overfit to pretraining contexts. Dodge et al. [15] show that blocklist filtering used in the construction of the C4 language modeling dataset removes information authored by and about people belonging to minority populations.

**Bias in CLIP.** Radford et al. [40] and Agarwal et al. [1] find that CLIP disproportionately associates text descriptions related to appearance with images of Female individuals; that images of Black individuals are the most likely to be misclassified as animals; and that images of individuals under 20 years old are most likely to be classified into crime-related categories [40]. Goh et al. [22] identify multimodal neurons in CLIP which associate stereotypes (such as immigration) with people or regions (such as Latin America). Wang et al. [55] use CLIP as an image retrieval system, and find evidence of pervasive gender imbalances in retrieved images. Birhane et al. [4] show that LAION-400m [45], a dataset of more than 400 million image-text pairs designed to be similar to the one on which CLIP trains, contains pornographic, misogynistic, and stereotypical images and accompanying text captions. It is unclear to what extent CLIP's training data (which has not been released to the public) was filtered for such content.

### 3 DATA

In analyzing the social impact of CLIP, Radford et al. [40] use the FairFace dataset of Karkkainen and Joo [29], designed to address the lack of racial diversity in computer vision datasets by creating nearly balanced classes for seven racial or ethnic groups. In order

<sup>1</sup>There are fundamental concerns with the use of facial recognition and analysis tools, which may be used for purposes of surveillance and social control [24, 31, 44]. However, the failure of such technologies to perform according to their intended use for underrepresented social groups reflects a bias of underrepresentation in the training data.

to compare more directly with the results reported by Radford et al. [41], and because FairFace is among the largest relatively balanced datasets of human images, we use the 86,744 images in the FairFace training dataset. The FairFace dataset labels images according to seven races or ethnicities, two genders, and nine age ranges:

- Race or Ethnicity: Black, White, Latino or Hispanic, East Asian, Southeast Asian, Indian<sup>2</sup>, Middle Eastern
- Gender: Female, Male
- Age: 0-2, 3-9, 10-19, 20-29, 30-39, 40-49, 50-59, 60-69, more than 70

For our experiments, we balance the dataset on the attribute or attributes under consideration. For example, when examining gender bias, images of Male individuals (the more frequently occurring category in FairFace) are removed without replacement until we have the same number of Male images and Female images. When examining intersectional bias operating over race, gender, and age, we obtain the number of images associated with the least frequent intersectional group and randomly remove images from every other group until they contain the same number. To mitigate the effects of downsampling, we repeat this process 10,000 times and report the mean.

We do not endorse the categorical or binary division of human beings based on race, gender, or age. Nevertheless, we recognize that the categories defined by FairFace are broadly operational in our current social context, and thus we operationalize them here to undertake research related to the biases which attend them. Because FairFace labels are assigned to images by Amazon Mechanical Turkers, the FairFace dataset reflects how images of humans are perceived by other potentially biased humans, and does not necessarily reflect self-identification.

**CLIP Training Data.** The patterns learned by unsupervised and semi-supervised machine learning models are dependent on the language and society which produces the training data. Radford et al. [40] train CLIP on the WebImageText corpus (WIT), a web scrape composed of 400 million images and associated captions. Radford et al. [40] produce the query list using every word which occurs at least 100 times in English Wikipedia, plus bigrams from Wikipedia with high pointwise mutual information, the names of Wikipedia articles, and all WordNet synsets [40]. The frequency-based heuristic for inclusion in WIT suggests that biases related to underrepresentation may occur in CLIP.

## 4 APPROACH AND EXPERIMENTS

We describe an approach for measuring self-similarity and testing whether CLIP exhibits cultural markedness.

**Representational Similarity.** We obtain a projected image embedding from every image in the Fairface dataset, and measure the self-similarity (*i.e.*, the mean pairwise cosine similarity) of projected image embeddings for each race or ethnicity, gender, and age range, as well as for the intersectional groups in the dataset. Specifically, for a set of images, the self-similarity  $s$  for a group of images  $G$  consisting of  $n$  images  $i \in G$  is given by:

$$s(G) = \frac{1}{n^2 - n} \sum_i \sum_{j \neq i} \cos(\vec{l}_i, \vec{l}_j) \quad (1)$$

$\cos$  in Equation 1 refers to cosine similarity, or the angular similarity of two vectors after normalization to unit length. Note that image representations projected to the visual semantic space in CLIP are multimodal, meaning that they can be meaningfully compared not only to other images, but to text encoded by the model as well. Prior work indicates that, as embedded images reach the top layers of the CLIP image encoder before projection, they acquire abstract conceptual features, which activate in response to textual, photographic, and symbolic representations of the depicted concept and other similar concepts [22]. For example, Goh et al. [22] find that the same multimodal neuron activates in the upper layers of the CLIP image encoder for cartoon depictions of Spider-Man, text of the word "spider," and photographs of humans dressed as Spider-Man. What this suggests for the present research is that, in a multimodal embedding space, self-similarity is a meaningful measurement for evaluating the extent to which a group is culturally marked by a shared conceptual characteristic: if, for example, people over the age of 70 are marked by CLIP based on age, we would expect that a group of images of people over 70 embedded in a visual semantic embedding space would be more similar to each other (by virtue of the shared, marked characteristic) than would a group of images of people in their 30s. Equation 1 measurement has been used previously to measure bias in language models such as BERT [14] and T5 [42], which were shown to embed names statistically more associated with underrepresented racial and gender groups such that they are more self-similar across contexts, reflecting that the range of representations is more narrowly defined [61]. To our knowledge, this is the first application of the measurement to study bias in multimodal AI.

### Markedness According to Race or Ethnicity, Gender, and Age.

For each image in the FairFace dataset, we obtain a projected image embedding. Then, we obtain projected text embeddings corresponding to every race or ethnicity label, every gender label, and every age label in the FairFace dataset, as described in Section 3. Radford et al. [40] find that, because single-word labels are uncommon in CLIP's dataset, the model performs better in the zero-shot setting when used with the prompt "a photo of a [image class]." To that end, we use the prompt "a photo of a [race or ethnicity] person" for race or ethnicity labels; "a photo of a [gender] person" for gender labels; and "a photo of a person between [age range minimum] and [age range maximum] years old" for age labels. For the "More than 70" age range, we use the prompt "a photo of a person more than 70 years old." Finally, we obtain a projected text embedding for the text label "a photo of a person," with no information related to race or ethnicity, gender, or age.

The cosine similarity is then computed between each projected text embedding and the projected image embedding. Cosine similarity is a direct proxy to the probability that an image is accurately associated with a text label in CLIP, and is used by the model to rank, retrieve, and classify images [40, 55]. For each image, we observe whether the cosine similarity of the projected image embedding is higher with the "a photo of a person" label than with any of the race or ethnicity labels; than with any of the gender labels; and than

<sup>2</sup>In the FairFace dataset, "Indian" refers not to people of Indigenous American ancestry, but to people whose ancestry originates in India.

with any of the age range labels. Our research is not concerned with whether CLIP has classified in accordance with the FairFace label for a set of socially defined categories. Rather, it measures whether the model associates an image of a person with *any* of the labels corresponding to a socially defined category, or if the image is left unmarked. Results report the percentage of the time a social group is unmarked (*i.e.*, the "a photo of a person" label is ranked highest by CLIP according to cosine similarity) for a given social category. For example, results will show that individuals who are White and Male and between the ages of 30 and 39 remain unmarked based on race (the "a photo of a person" label is preferred) 46.6% of the time. We report results for each race or ethnicity, gender, and age range in the FairFace dataset, and examine results from an intersectional lens to observe whether, for example, the age of a person has an effect on whether they are marked with a gender label.


## 5 RESULTS

The evidence indicates that CLIP embeds images of women, of relatively younger and older people, and of people belonging to racial and ethnic groups other than White such that they are more self-similar. CLIP consistently prefers to mark the race or ethnicity of racial groups other than White, the gender of women, and age of relatively young or old people, while leaving images of people who are White, Male, and middle-aged unmarked.

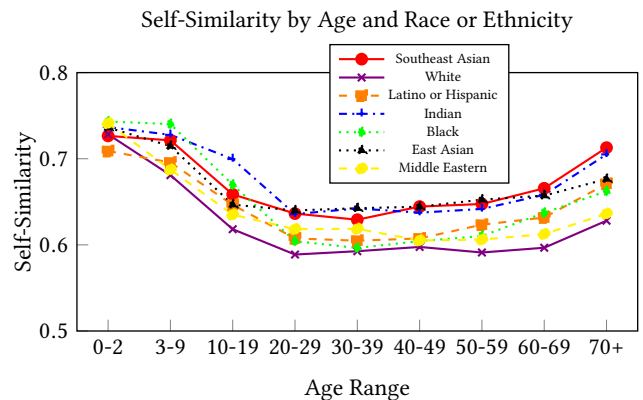
**Representational Uniformity.** The evidence reveals biases of representational uniformity which differ based on race or ethnicity, gender, and age. Table 1 indicates that the most variation ( $s = 0.573$ ) exists in the multimodal image embeddings of White individuals, while embeddings of Indian, East Asian, and Southeast Asian individuals are the most self-similar to each other, with  $s > 0.620$ . Table 3 indicates that embedded images of Male individuals are less self-similar ( $s = 0.566$ ) than embedded images of Female individuals ( $s = 0.592$ ). Table 2 describes differences in self-similarity based on age. Most notably, images of individuals between the ages of 3 and 9 are much more self-similar ( $s = 0.669$ ) than those of individuals between 20 and 29, 30 and 39, 40 and 49, and 50 and 59 ( $s \in [0.573, 0.575]$ ). A similar phenomenon manifests for individuals between the ages of 10 and 19, for which self-similarity is higher ( $s = 0.611$ ) than for images of individuals between 20 and 59, despite the significant physical variation of people aged 10-19.

Examining the data from an intersectional lens reveals additional disparities. Figure 1 indicates that images of Female individuals are more self-similar when compared within the same age range against Male individuals, a difference which is first evident for the 20–29 age range. With each increase in age range, the difference in self-similarity between Female individuals and Male individuals also increases: from 0.036 in the 40-49 age range ( $s = 0.619$  vs.  $s = 0.583$ ), to 0.051 in the 50-59 age range ( $s = 0.635$  vs.  $s = 0.584$ ), to 0.065 in the 60-69 age range ( $s = 0.651$  vs.  $s = 0.586$ ), to 0.071 in the 70-79 age range ( $s = 0.679$  vs.  $s = 0.608$ ). The data indicate a gender bias which affects representations of Female individuals as age increases.

Figure 2 shows that embedded images of White individuals are the least self-similar at every age except for 0-2. As with the gender bias, increases in age exacerbate already existing differences based on race. For individuals between 40 and 49, the difference between



**Figure 1: At each increase in age range beyond 20-29, the disparity in self-similarity between Female individuals and Male individuals increases, indicating that age exacerbates differences based on gender in CLIP representations.**



**Figure 2: At each increase in age range beyond 40-49, the disparity in self-similarity between White individuals and people belonging to the most self-similar race or ethnicity (Southeast Asian) also increases, indicating that age exacerbates differences based on race or ethnicity in CLIP representations.**

the most self-similar race or ethnicity (Southeast Asian,  $s = 0.644$ ) and the least self-similar race or ethnicity (White,  $s = 0.598$ ) is 0.046; this difference increases with each increase in age, such that self-similarity for individuals who are White and more than 70 is 0.628, while self-similarity for individuals who are Southeast Asian and more than 70 is 0.713, a difference of 0.085.

We face a challenge in comprehensibly describing the intersectional effects of self-similarity given the significant diversity in race and ethnicity, gender, and age in the dataset we use. A partial description of these results is given in Table 4, which reports the self-similarity of projected image embeddings for the most self-similar and the least self-similar social groups. In addition to these tables, we visualize the difference in self-similarity by age for individuals who are White and Male and for individuals who are Black and Female. The selection of these groups is not arbitrary. Crenshaw

Self-Similarity by Race or Ethnicity							
Metric	White	Latino/a	Middle Eastern	Black	Southeast Asian	East Asian	Indian
Self-Similarity	0.573	0.589	0.594	0.604	0.623	0.623	0.626

**Table 1: CLIP representations of White individuals are less self-similar than representations of people of any other race or ethnicity, indicating that CLIP directs less attention to the race of White individuals when encoding multimodal representations, as it does not cluster them as closely together based on the shared characteristic of race.**

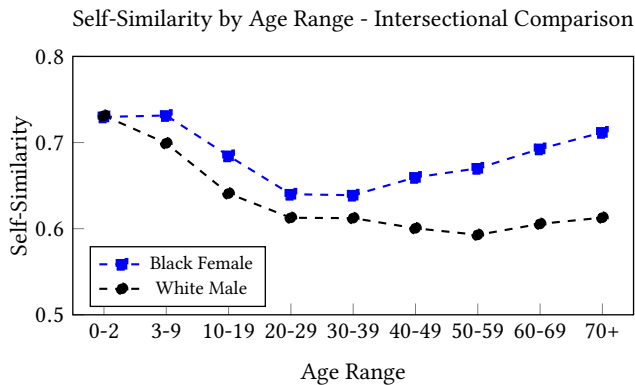
Self-Similarity by Age Range									
Metric	0-2	3-9	10-19	20-29	30-39	40-49	50-59	60-69	70+
Self-Similarity	0.692	0.669	0.611	0.575	0.574	0.573	0.574	0.583	0.623

**Table 2: CLIP representations of individuals under 20 or over 70 are more self-similar than representations of people of any other age, indicating that CLIP directs more attention to the age of relatively younger and older individuals.**

Self-Similarity by Gender		
Metric	Male	Female
Self-Similarity	0.566	0.592

**Table 3: CLIP representations of Female individuals are more self-similar than representations of Male individuals, indicating that the model directs more attention to gender in forming representations of women.**

[11] introduce intersectionality by describing the overlapping axes of discrimination faced by Black women, while White men have enjoyed economic, political, and cultural advantages based on race and gender. If an intersectional effect exists which compounds the effects noted for race, gender, and age, we would expect to be able to observe this via a comparison of these two social groups.



**Figure 3: At every age range after 0-2, CLIP representations of individuals who are Black and Female are more self-similar than representations of individuals who are White and Male, and the disparity between the groups is greater than the disparities between Male individuals and Female depicted in Figure 1.**

Figure 3 demonstrates that disparities in self-similarity are exacerbated in the intersection of race, gender, and age. The same divergence in self-similarity exists as was observed for gender in Figure 1. However, the magnitude of the difference is greater at every age range, including much greater differences for individuals

between the ages of 3 and 39, which are relatively small when observed solely for gender. Concretely, the self-similarity of images of individuals who are White and Male and 70 or older is 0.613, while the self-similarity of images of individuals who are Black and Female and 70 or older is 0.712, a difference of 0.099, higher than the difference based on gender for individuals 70 and older (.071) and the largest difference based on race for individuals 70 and older (.085).

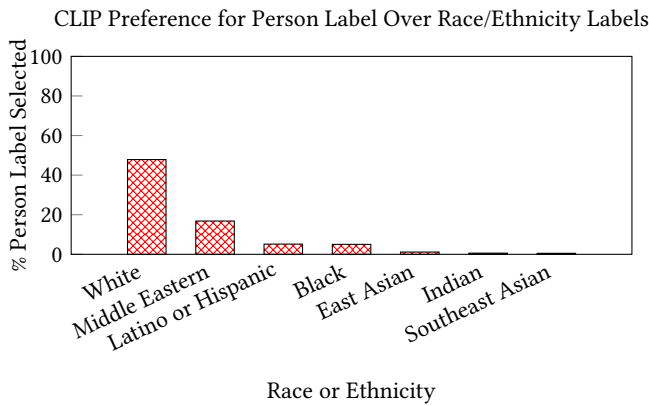
Table 4 indicates that the ten least self-similar social groups all reflect Male individuals; that they reflect only three of the seven races or ethnicities in the dataset (White, Black, and Latino or Hispanic); and that eight of the ten groups come from age ranges between 20 and 59 years old, with no age ranges under the age of 20 represented. Six of the ten social groups reflect individuals who are White and Male, including the only two groups of individuals over the age of 60. The most self-similar social groups all reflect people who are younger than age of 10 or older than the age of 70. Unlike the ten least self-similar groups, six of the ten most self-similar groups reflect Female individuals, and only one of the groups reflects White individuals.

**Markedness Based on Race or Ethnicity, Gender, and Age.** As shown in Figure 4, CLIP ranks the “a photo of a person” label highest 47.9% of the time for images of White individuals. For images of individuals of every other race or ethnicity in FairFace, CLIP ranks “a photo of a person” highest less than 16.8% of the time, and less than 1.5% of the time for East Asian, Indian, or Southeast Asian individuals. CLIP ranks “a photo of a person” highest 26.7% of the time for images of Male individuals, but 15.2% of the time for images of Female individuals, as shown in Figure 5, reflecting a preference to mark the gender of Female individuals.


As shown in Figure 6, CLIP ranks the “a photo of a person” label highest 48.4% of the time for individuals between the ages of 20-29, 57.9% of the time for images of individuals between the ages of

Self-Similarity by Social Group			
Most Self-Similar Social Groups		Least Self-Similar Social Groups	
Social Group	Self-Similarity	Social Group	Self-Similarity
Black Male 0-2	0.767	White Male 70+	0.613
Black Male 3-9	0.759	White Male 20-29	0.613
Southeast Asian Female 70+	0.757	Latino or Hispanic Male 30-39	0.612
East Asian Female 0-2	0.751	White Male 30-39	0.612
Indian Female 0-2	0.751	Black Male 50-59	0.608
Middle Eastern Male 0-2	0.748	White Male 60-69	0.606
Southeast Asian Female 0-2	0.747	Black Male 30-39	0.603
Indian Female 70+	0.746	White Male 40-49	0.600
White Female 0-2	0.744	Black Male 20-29	0.600
Indian Male 3-9	0.742	White Male 50-59	0.593

**Table 4: The most self-similar social groups reflect individuals younger than 10 or older than 70, and predominantly Female individuals. The least self-similar social groups all reflect Male individuals, and predominantly White individuals and individuals between the ages of 20 and 59.**




**Figure 4: CLIP prefers a label with race or ethnicity omitted less than 6% of the time for five of seven racial or ethnic groups, but prefers a label with no information about race or ethnicity 47.9% of the time for White individuals.**



**Figure 5: CLIP prefers the "person" label over labels with gender information 26.7% of the time for Male individuals, compared to 15.2% of the time for Female individuals, indicating that CLIP marks the gender of women at higher rates than men.**


30-39, 56.1% for images of individuals 40-49, and 39.2% of the time for individuals 50-59. CLIP ranks the "a photo of a person" label highest less than 1% of the time for images of individuals under 2 years old or between 3 years old and 9 years old, and 6.7% of the time for images of individuals over 70 years old. The data suggests a bias which preferentially draws attention to the age of individuals younger than 20 or older than 59.



**Figure 6: CLIP prefers the "person" label over labels with age information at least 39% of the time for individuals between 20 and 29, 30 and 39, 40 and 49, and 50 and 59, but less than 7% of the time for individuals over 70, and less than 1% of the time for individuals under 10.**

Examination of the results based on gender and age yields additional insight. Figure 7 shows that, at every age range, CLIP prefers the "a photo of a person" label at a higher rate for Male individuals than for Female individuals. For both gender groups, images of individuals under the age of 10 or over the age of 70 are the least likely to be marked based on gender. Images of Female individuals between the ages of 30 and 39 are the most likely to be marked based on gender, with a gender label ranked higher than the "person" label 91.5% of the time. Images of Female individuals between the ages of 10 and 19, 20 and 29, 40 and 49, and 50 and 59 are marked based on gender more than 85.0% of the time. The


largest disparities between the Male and Female groups occur in the 3-9 and 10-19 age ranges. In the 3-9 age range, 59.0% of Male individuals are most associated with the "a photo of a person" label, compared to 38.1% of Female individuals. In the 10-19 age range, 31.1% of Male individuals are most associated with the "a photo of a person" label, compared to 13.3% of Female individuals, reflecting that CLIP marks the gender of Female individuals at an earlier age than it does Male individuals. Differences between Male individuals and Female individuals diminish in the 50-59 and 60-69 age ranges, before diverging again in the 70-79 age range, with 39.2% of Female images unmarked compared to 51.1% of Male images.



**Figure 7: Female individuals are more likely to be marked with a gender label at every age range, and are marked based on gender at earlier ages than are Male individuals.**

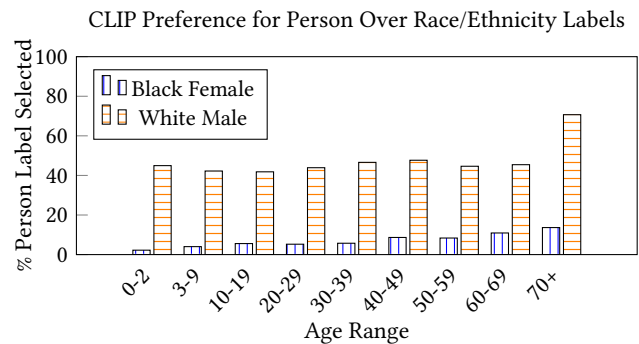
As shown in Figure 8, the least likely age ranges to be marked based on gender are the most likely to be marked based on age. At least 99% of individuals under the age of 10 are marked based on age, regardless of gender. In the 10-19 age range, CLIP ranks the "a photo of a person" label highest 20.7% of the time for Female individuals, compared to 12.3% of the time for Male individuals. In this age range, Female individuals are more likely to be marked based on gender, but less likely to be marked based on age, than Male individuals. Among Female individuals, CLIP ranks the "a photo of a person" label highest most frequently for people between 30 and 39, at 54.3%; among Male individuals, CLIP ranks the "a photo of a person" label highest most frequently for people between 40 and 49, at 62.2%. The most significant disparities between Male individuals and Female individuals occur in the 40-49 (62.2% vs. 45.1% unmarked) and 50-59 (47.4% vs. 25.3% unmarked) age ranges.

Finally, we report results based on race or ethnicity, gender, and age for individuals who are White and Male, and for individuals who are Black and Female. As shown in Figure 9, CLIP ranks the "a photo of a person" label highest at least 40% of the time in every age range for individuals who are White and Male, and less than 15% of the time in all age ranges for individuals who are Black and Female. A pattern emerges for individuals who are Black and Female: as age increases, the probability of CLIP preferring the "a photo of a person" label over a label denoting race or ethnicity also increases, from 2.2% at 0-2, to 5.7% at 30-39, to 13.7% at more than 70. The pattern is less consistent for individuals who are White and Male, but increases significantly, from 45.4% to 70.7%, from the 60-69 age range to the more than 70 age range.




**Figure 8: Female individuals are less likely to be marked with an age label prior to the age of 30, but are more likely to be marked with an age label after the age of 30, with the most significant disparity in the 50-59 age range.**

Figure 10 shows that, until the 50-59 age range, individuals who are White and Male are less frequently marked based on gender than are individuals who are Black and Female, with the largest disparities in the 0-2 and 3-9 age ranges. In the 50-59, 60-69, and more than 70 age ranges, individuals who are Black and Female are marked based on gender less frequently than individuals who are White and Male. This is the only circumstance in which individuals who are Black and Female are marked less frequently than individuals who are White and Male, for gender, age, or race or ethnicity labels. Figure 11 shows that, in every age range, individuals who are White and Male are less likely to be marked based on age than individuals who are Black and Female. For all age ranges from 30-39 through more than 70, the disparity is similar but more severe than the disparity observed based solely on gender, and is most pronounced at the 50-59 age range (60.8% vs. 20.2% unmarked). These results suggest that, as with measurements of self-similarity, already existing biases related to CLIP's preference to mark are magnified in the intersection of race, gender, and age.




**Figure 9: In every age range, CLIP prefers a label with race or ethnicity omitted for people who are White and Male at higher rates than for individuals who are Black and Female. As age increases, the probability of individuals who are Black and Female receiving a label with race or ethnicity omitted also increases.**





**Figure 10: Individuals who are Black and Female are less likely than individuals who are White and Male to receive a label with gender omitted before the age of 20, but are more likely to receive the "person" label (with gender omitted) after the age of 50.**



**Figure 11: In every age range after 3-9, individuals who are White and Male are more likely to receive the "person" label (with age omitted) than individuals who are Black and Female. The gender disparity in age labeling observed in Figure 8 is further exacerbated when the Male group is also White, and the Female group is also Black.**

## 6 DISCUSSION

Our results provide evidence of biases related to the representation of gender, age, and race or ethnicity in the CLIP embedding space. While prior work has demonstrated semantic biases in CLIP [1, 22, 40, 55], our results indicate a fundamental form of descriptive bias in the model, related to who is marked, and according to what socially defined characteristics they are marked. Representations of Male individuals; White individuals; and individuals between the ages of 20 and 59 are consistently among the least self-similar, and the least likely to be marked. Representations of individuals who are Female; who belong to underrepresented racial or ethnic groups; and who are under the age of 20 or over the age of 60 are consistently the most self-similar, and the most likely to be marked.

Differentially greater self-similarity is not a desirable property in a multimodal embedding space, as it reflects that the embeddings cluster more tightly around the gender, age, or race or ethnicity

of the people depicted in the images. The disparities between the most self-similar and least self-similar social groups suggests that independently impactful biases based on race, gender, and age compound to multiply impact groups possessing more than one of these constituent characteristics. Results based on self-similarity are most comparable to prior work finding that low-frequency names which are statistically more associated with women and with underrepresented races and ethnicities are more self-similar across contexts in language models, and are represented using a smaller region of the embedding space than names more associated with men and with people who are White [61]. Our results indicate that such effects are not limited to linguistic representations, but extend to language-and-vision embedding spaces wherein representations have conceptual features. Moreover, using a large dataset of faces overcomes the ambiguity of measuring the bias of a linguistic representation like a name.

Our results demonstrate the effects of natural language supervision in shaping a multimodal embedding space. Among the most consistent results is that age has significant impact on whether the model prefers to mark not only age but also gender and, to a lesser extent, race or ethnicity. Among Female individuals, 91.5% of those between 30 and 39 are marked based on gender, compared to 60.8% of those more than 70. This suggests that the model arranges its embedding space based not on fine-grained recognition of visual information, but on the categories used to describe similar images in the training data. Moreover, such results suggest that age should be considered as a factor in interpreting the results of research on bias in AI, even when that research does not intend to test primarily the effects of age.

The results of the markedness experiment appear to form patterns similar to a normal distribution across age ranges, with a mean centered at 30-39, and shifted based on gender or on race or ethnicity. Given that this experiment quantifies how frequently "a photo of a person" was selected over more descriptive labels, the normality of these distributions suggests that CLIP represents certain characteristics (White, Male, between 20 and 59) as more central to the concept of "person," and other characteristics as different enough that they need to be marked using more descriptive language.

**Limitations and Future Work.** Our research uses the FairFace dataset to study bias in CLIP. While this is the most balanced and diverse dataset of human images of which we are aware, it is limited in that it uses a small number of socially defined categories for gender and for race or ethnicity, and assigns categories based not on self-identification but on the perception of annotators. In adopting this dataset as a source of images, our results are necessarily contextualized within those categories. Diverse, balanced, and ethically obtained datasets of human images are needed to permit less constrained research designs at a similar scale (more than 80,000 images) to our work. Our research examines only one visual semantic model, and future work will be needed to assess whether our results generalize across architectures. As the field of language-and-vision AI matures, systematic study across architectures may reveal additional insight into the bias of markedness. Moreover, CLIP embeds highly contextual sentence representations, and the results of the markedness experiment may be affected by using different prompts or social categories. We have adopted a

principled approach using only the prompt specified by Radford et al. [40] and the social categories defined by Karkkainen and Joo [29]. Future work might explore how systematically varying such settings affects CLIP’s preference to mark.

Radford et al. [40] note that CLIP was first intended to have the capabilities of a zero-shot caption generator, *i.e.*, a model capable of not just matching an image with a label but of producing that caption using a language model [40]. Moreover, one of the first uses of CLIP was to train DALL-E, a zero-shot text-to-image transformer capable of generating novel visual representations from text input [43]. As such generative models are developed, research might be directed to understanding what socially defined categories are made conspicuous or invisible in the text and images generated by AI, in ways which could impact society. Finally, future work might explore the connection between self-similarity and frequency of representation in multimodal training corpora. Prior work suggests that low frequency leads to high self-similarity in language models [61], and our results identify a similar effect in multimodal AI.

## 7 CONCLUSION

We demonstrate that a multimodal visual semantic AI model learns to unequally mark gender, age, and race or ethnicity. Biases learned in visual semantic embedding spaces are likely to affect the next generation of state-of-the-art AI applications, which build on the ability of such models to associate images with text in a zero-shot setting.

**Acknowledgements:** This material is based on research partially supported by the U.S. National Institute of Standards and Technology (NIST) Grant 60NANB20D212T. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of NIST.

## REFERENCES

- [1] Sandhini Agarwal, Gretchen Krueger, Jack Clark, Alec Radford, Jong Wook Kim, and Miles Brundage. 2021. Evaluating CLIP: Towards Characterization of Broader Capabilities and Downstream Implications. *arXiv preprint arXiv:2108.02818* (2021).
- [2] Edwin L Battistella et al. 1990. *Markedness: The evaluative superstructure of language*. SUNY Press.
- [3] Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 610–623.
- [4] Abeba Birhane, Vinay Uday Prabhu, and Emmanuel Kahembwe. 2021. Multimodal datasets: misogyny, pornography, and malignant stereotypes. *arXiv preprint arXiv:2110.01963* (2021).
- [5] Geoffrey C Bowker and Susan Leigh Star. 2000. *Sorting things out: Classification and its consequences*. MIT press.
- [6] Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*. PMLR, 77–91.
- [7] Aylin Caliskan, Joanna J Bryson, and Arvind Narayanan. 2017. Semantics derived automatically from language corpora contain human-like biases. *Science* 356, 6334 (2017), 183–186.
- [8] Aylin Caliskan and Molly Lewis. [n. d.]. Social biases in word embeddings and their relation to human cognition. ([n. d.]).
- [9] Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. 2020. Generative pretraining from pixels. In *International Conference on Machine Learning*. PMLR, 1691–1703.
- [10] Patricia Hill Collins and Sirma Bilge. 2020. *Intersectionality*. John Wiley & Sons.
- [11] Kimberlé Crenshaw. 1989. Demarginalizing the intersection of race and sex: A black feminist critique of antidiscrimination doctrine, feminist theory and antiracist politics. *u. Chi. Legal f.* (1989), 139.
- [12] Kimberle Crenshaw. 1990. Mapping the margins: Intersectionality, identity politics, and violence against women of color. *Stan. L. Rev.* 43 (1990), 1241.
- [13] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 248–255.
- [14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, Minneapolis, Minnesota, 4171–4186. <https://doi.org/10.18653/v1/N19-1423>
- [15] Jesse Dodge, Maarten Sap, Ana Marasovic, William Agnew, Gabriel Ilharco, Dirk Groeneveld, and Matt Gardner. 2021. Documenting the english colossal clean crawled corpus. *arXiv preprint arXiv:2104.08758* (2021).
- [16] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xi-aohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*.
- [17] Wolfgang U Dressler. 1985. On the predictiveness of Natural Morphology 1. *Journal of linguistics* 21, 2 (1985), 321–337.
- [18] Thomas B Fitzpatrick. 1988. The validity and practicality of sun-reactive skin types I through VI. *Archives of dermatology* 124, 6 (1988), 869–871.
- [19] Andrea Frome, Greg Corrado, Jonathon Shlens, Samy Bengio, Jeffrey Dean, Marc’Aurelio Ranzato, and Tomas Mikolov. 2013. Devise: A deep visual-semantic embedding model. (2013).
- [20] Negin Ghavami and Letitia Anne Peplau. 2013. An intersectional analysis of gender and ethnic stereotypes: Testing three hypotheses. *Psychology of Women Quarterly* 37, 1 (2013), 113–127.
- [21] Talmy Givón. 1991. Markedness in grammar: Distributional, communicative and cognitive correlates of syntactic structure. *Studies in Language. International Journal sponsored by the Foundation “Foundations of Language”* 15, 2 (1991), 335–370.
- [22] Gabriel Goh, Nick Cammarata, Chelsea Voss, Shan Carter, Michael Petrov, Ludwig Schubert, Alec Radford, and Chris Olah. 2021. Multimodal neurons in artificial neural networks. *Distill* 6, 3 (2021), e30.
- [23] Joseph Harold Greenberg. 1963. Universals of language. (1963).
- [24] Nitzan Guetta, Asaf Shabtai, Inderjeet Singh, Satoru Momiyama, and Yuval Elovici. 2021. Dodging Attack Using Carefully Crafted Natural Makeup. *arXiv preprint arXiv:2109.06467* (2021).
- [25] Wei Guo and Aylin Caliskan. 2021. Detecting emergent intersectional biases: Contextualized word embeddings contain a distribution of human-like biases. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. 122–133.
- [26] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [27] Roman Jakobson. 1972. Verbal communication. *Scientific American* 227, 3 (1972), 72–81.
- [28] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc V Le, Yunhsuan Sung, Zhen Li, and Tom Duerig. 2021. Scaling Up Visual and Vision-Language Representation Learning With Noisy Text Supervision. *arXiv e-prints* (2021), arXiv–2102.
- [29] Kimmo Karkkainen and Jungseock Joo. 2021. FairFace: Face Attribute Dataset for Balanced Race, Gender, and Age for Bias Measurement and Mitigation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. 1548–1558.
- [30] Eugenia Kim, De’Aira Bryant, Deepak Srikanth, and Ayanna Howard. 2021. Age Bias in Emotion Detection: An Analysis of Facial Emotion Recognition Performance on Young, Middle-Aged, and Older Adults. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. 638–644.
- [31] James Leibold. 2020. Surveillance in China’s Xinjiang region: Ethnic sorting, coercion, and inducement. *Journal of Contemporary China* 29, 121 (2020), 46–60.
- [32] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2019. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557* (2019).
- [33] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. ViLBERT: pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*. 13–23.
- [34] Chandler May, Alex Wang, Shikha Bordia, Samuel Bowman, and Rachel Rudinger. 2019. On Measuring Social Biases in Sentence Encoders. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. 622–628.
- [35] Willi Mayerthaler. 1988. *Morphological naturalness*. Karoma Pub.
- [36] Norman Mu, Alexander Kirillov, David Wagner, and Saining Xie. 2021. SLIP: Self-supervision meets Language-Image Pre-training. *arXiv preprint arXiv:2112.12750* (2021).

- [37] Moin Nadeem, Anna Bethke, and Siva Reddy. 2020. Stereoset: Measuring stereotypical bias in pretrained language models. *arXiv preprint arXiv:2004.09456* (2020).
- [38] Pandu Nayak. 2021. MUM: A new AI milestone for understanding information. <https://blog.google/products/search/introducing-mum/>
- [39] Joon Sung Park, Michael S Bernstein, Robin N Brewer, Ece Kamar, and Meredith Ringel Morris. 2021. Understanding the Representation and Representativeness of Age in AI Data Sets. *arXiv preprint arXiv:2103.09058* (2021).
- [40] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. *arXiv preprint arXiv:2103.00020* (2021).
- [41] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog* 1, 8 (2019), 9.
- [42] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *Journal of Machine Learning Research* 21, 140 (2020), 1–67.
- [43] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. 2021. Zero-shot text-to-image generation. *arXiv preprint arXiv:2102.12092* (2021).
- [44] Morgan Klaus Scheuerman, Madeleine Pape, and Alex Hanna. 2021. Auto-essentialization: Gender in automated facial analysis as extended colonial project. *Big Data & Society* 8, 2 (2021), 20539517211053712.
- [45] Christoph Schuhmann. 2021. LAION-400-Million Open Dataset. <https://laion.ai/laion-400-open-dataset/>
- [46] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. 2018. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 2556–2565.
- [47] Emily Sheng, Kai-Wei Chang, Prem Natarajan, and Nanyun Peng. 2019. The Woman Worked as a Babysitter: On Biases in Language Generation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. 3407–3412.
- [48] Richard Socher, Milind Ganjoo, Christopher D Manning, and Andrew Ng. 2013. Zero-Shot Learning Through Cross-Modal Transfer. In *Advances in Neural Information Processing Systems*. 935–943.
- [49] Ryan Steed and Aylin Caliskan. 2021. Image representations learned with unsupervised pre-training contain human-like biases. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 701–713.
- [50] Yi Chern Tan and L Elisa Celis. 2019. Assessing social and intersectional biases in contextualized word representations. *NeurIPS* (2019).
- [51] Deborah Tannen. 1993. Marked women, unmarked men. *The New York Times Magazine* 20, 1993 (1993), 52–54.
- [52] Yonglong Tian, Dilip Krishnan, and Phillip Isola. 2019. Contrastive Representation Distillation. In *International Conference on Learning Representations*.
- [53] Saurabh Tiwary. 2021. Turing Bletchley: A Universal Image Language Representation model by Microsoft. <https://www.microsoft.com/en-us/research/blog/turing-bletchley-a-universal-image-language-representation-model-by-microsoft/>
- [54] Nikolai Sergeevich Trubetzkoy. 1969. Principles of phonology. (1969).
- [55] Jialu Wang, Yang Liu, and Xin Eric Wang. 2021. Are Gender-Neutral Queries Really Gender-Neutral? Mitigating Gender Bias in Image Search. *arXiv preprint arXiv:2109.05433* (2021).
- [56] Tianlu Wang, Jieyu Zhao, Mark Yatskar, Kai-Wei Chang, and Vicente Ordonez. 2019. Balanced datasets are not enough: Estimating and mitigating gender bias in deep image representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 5310–5319.
- [57] Zirui Wang, Jiahui Yu, Adams Wei Yu, Zihang Dai, Yulia Tsvetkov, and Yuan Cao. 2021. SimVLM: Simple Visual Language Model Pretraining with Weak Supervision. *arXiv preprint arXiv:2108.10904* (2021).
- [58] Linda R Waugh. 1982. Marked and unmarked: A choice between unequals in semiotic structure. (1982).
- [59] Benjamin Wilson, Judy Hoffman, and Jamie Morgenstern. 2019. Predictive inequity in object detection. *arXiv preprint arXiv:1902.11097* (2019).
- [60] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. Transformers: State-of-the-Art Natural Language Processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. Association for Computational Linguistics, Online, 38–45. <https://www.aclweb.org/anthology/2020.emnlp-demos.6>
- [61] Robert Wolfe and Aylin Caliskan. 2021. Low Frequency Names Exhibit Bias and Overfitting in Contextualizing Language Models. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. 518–532.
- [62] Robert Wolfe and Aylin Caliskan. 2022. Contrastive Visual Semantic Pretraining Magnifies the Semantics of Natural Language Representations. *Association for Computational Linguistics* (2022).
- [63] Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V Le. 2020. Self-training with noisy student improves imagenet classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10687–10698.
- [64] Yuhao Zhang, Hang Jiang, Yasuhide Miura, Christopher D Manning, and Curtis P Langlotz. 2020. Contrastive learning of medical visual representations from paired images and text. *arXiv preprint arXiv:2010.00747* (2020).