

Dataset Scale and Societal Consistency Mediate Facial Impression Bias in Vision-Language AI

Robert Wolfe, Aayushi Dangol, Alexis Hiniker, Bill Howe

University of Washington
rwolfe3@uw.edu, adango@uw.edu, alexisr@uw.edu, billhowe@uw.edu

Abstract

Multimodal AI models capable of associating images and text hold promise for numerous domains, ranging from automated image captioning to accessibility applications for blind and low-vision users. However, uncertainty about harmful bias has in some cases limited their adoption and availability. In the present work, we study 43 CLIP vision-language models to determine whether they learn human-like facial impression biases, and we find evidence that such biases are reflected across three distinct CLIP model families. We show for the first time that the *degree* to which a bias is shared across a society predicts the degree to which it is reflected in a CLIP model. Human-like impressions of visually *unobservable* attributes, like trustworthiness and sexuality, emerge only in models trained on the largest dataset, indicating that a better fit to uncurated cultural data results in the reproduction of increasingly subtle social biases. Moreover, we use a hierarchical clustering approach to show that dataset size predicts the extent to which the underlying structure of facial impression bias resembles that of facial impression bias in humans. Finally, we show that Stable Diffusion models employing CLIP as a text encoder learn facial impression biases, and that these biases intersect with racial biases in Stable Diffusion XL-Turbo. While pretrained CLIP models may prove useful for scientific studies of bias, they will also require significant dataset curation when intended for use as general-purpose models in a zero-shot setting.

Introduction

OpenAI’s multimodal GPT-4 powers the beta version of Be My AI, an extension of the Be My Eyes app (Be My Eyes 2023a,b) that provides “instantaneous identification, interpretation, and conversational visual assistance” to blind and low-vision users. Until recently, the app allowed users to ask questions about images of people and receive live explanations. The temporary discontinuation of this feature was motivated by concern that GPT-4 “would say things it shouldn’t about people’s faces, such as assessing their gender or emotional state” (Hill 2023).

The decision belies a broader concern: that by learning to associate language and images, multimodal AI may make insufficiently informed judgments about human attributes based solely on a person’s face. When studied in human

subjects, this kind of inference is known as a “first impression” or “facial impression” bias (Todorov 2017), and it is known to affect consequential spheres of human social life such as criminal sentencing (Wilson and Rule 2015), employment decisions (Stoker, Garretsen, and Spreeuwers 2016), and political elections (Antonakis and Dalgas 2009). Such impressions can include traits like trustworthiness, which are unobservable from a person’s face and societally mediated to extent that they are consistent in a population (Todorov 2017). While psychologists have used computational geometry and supervised machine learning approaches to modeling facial impression biases (Blanz and Vetter 2023), it is not known whether semi-supervised vision-language AI models could inadvertently learn such biases in pretraining and propagate them to the many domains in which such models are used.

While features permitting facial image analysis are disabled in GPT-4, the opportunity to study facial impression bias is afforded by CLIP (“Contrastive Language Image Pre-training”), a state-of-the-art vision-language model that allows users to define text classes at inference using natural language (Radford et al. 2021). Rather than fine-tuning CLIP to model facial impressions similar to prior work using supervised learning, we study this bias in three families of pretrained CLIP models used in a wide range of multimodal computer vision tasks: the nine models trained by OpenAI (Radford et al. 2021); five “FaceCLIP” models post-trained for facial analysis (Zheng et al. 2022); and 29 “Scaling” models trained by Cherti et al. (2022) on systematically differing amounts of data, allowing for statistical analysis of the effects of model and dataset parameters on facial impression bias.

Analyzing whether CLIP models learn human-like facial impression biases requires a reliable source of human data. This research uses the authoritative One Million Impressions (OMI) dataset of Peterson et al. (2022), which includes 1,004 images of faces rated by human participants across 34 attributes, with which Peterson et al. (2022) learned a supervised model of facial impression biases. In the present work, we used CLIP to compute the similarity of each OMI image to text prompts for the 34 attributes, mimicking the task given to human subjects, and we compared the CLIP similarities to human subject ratings. We offer four primary findings:

1. **CLIP models learn societal facial impression biases, including for unobservable traits such as trustworthiness and sexuality.** Moreover, the extent to which an attribute




Figure 1: CLIP models learn human-like facial impression biases. The highest model-human correlations are obtained for intuitively visual categories that are broadly shared by a society (such as gender, age, and happiness). Models trained on the largest dataset (LAION-2B) exhibit more human-like biases than FaceCLIP or OpenAI models for most attributes.

bias is learned by a CLIP model is strongly correlated with the inter-rater reliability (IRR) of human judgments of the attribute (Spearman’s $\rho = .73$ for OpenAI models; $\rho = .76$ for FaceCLIP models; and $\rho = .72$ for Scaling models). A multiple linear regression predicting the similarity of CLIP bias to Human bias finds that the IRR of the attribute plays a larger role than any model-related variable, with $t(912) = 25.47, p < .001$. The extent to which a facial impression bias is learned by a model depends on how consistently it is shared in the population that produced the data on which the models trains.

2. **Dataset Scale is a significant predictor of facial impression bias in CLIP.** Comparison of model-human similarity in two groups of nine CLIP models trained on LAION-80M (80 million examples) and LAION-400M (407 million examples) yields large effect sizes ($d > 0.8$) and statistically significant ($p < .05$) paired samples t -tests for 17 of 34 attributes, indicating increases in the human similarity of bias in models trained on LAION-400M. Differences between models trained on LAION-2B (2.32 billion examples) and LAION-400M are mostly not significant, with the notable exception of *unobservable* attributes like trustworthiness ($d = 1.33, p < .001$) and sexuality ($d = 1.14, p < .001$). While models trained on larger datasets exhibit stronger task performance (Cherti et al. 2022), they also more faithfully reflect the biases of the population that produced the data.
3. **CLIP models learn human-like associations between facial impression biases.** Hierarchical clustering of CLIP and OMI attribute correlation matrices reveals similar

groupings of traits, including clusters related to ethnicity, and clusters grouping gender, sexuality, and age. Computing the normalized Frobenius inner product of CLIP correlation matrices with the OMI matrix reveals increasing similarity as pretraining data size increases, with a one-way ANOVA yielding $F(2) = 15.71, p < .001$, and large ($d > 0.8$) pairwise effect sizes between groups of Scaling-80M, 400M, and 2B models.

4. **Stable Diffusion text-to-image models employing CLIP as a text encoder learn facial impression biases that intersect with demographic biases.** Images generated by Stable Diffusion XL-Turbo (SDXL) are classified by a classifier fit using the OMI images. Classifications reflect human facial impression biases for subjective, observable attributes like attractiveness ($F1=.98$), and to a lesser extent for unobservable attributes like liberal ($F1=.68$) and smart ($F1=.65$). Applying the classifier to SDXL images generated for White and Black prompts reveals biases in SDXL differentially associating White individuals with traits like memorable, attractive, electable, and happy.

Training vision-language models on vast web-scraped datasets produces subtle yet consequential emergent biases. Such models may serve as useful tools for social science, potentially illuminating factors that contribute to human bias. However, the presence of such biases in vision-language AI also renders fraught these models’ real-world use, as they may subtly reinforce existing inequities in an online environment increasingly mediated by AI. Code and data are available at <https://github.com/wolferobert3/vl-facial-impressions-bias>.

Related Work

We review the related work on facial impression bias, vision-language AI, and the impact of scale in deep learning.

Facial Impression Bias

A wealth of psychological research indicates that humans make immediate judgments about the attributes of people they do not know based solely on facial appearance (Willis and Todorov 2006; Oh, Buck, and Todorov 2019; Charlesworth et al. 2019). Information inferred from faces includes character traits (like trustworthiness and outgoingness) and socially constructed group memberships (like gender and ethnicity) as well as relatively objective traits (like hair color and weight) (Todorov 2017; Peterson et al. 2022). Research on first-impression biases in humans has found that the inference of attributes from facial appearance plays a role in numerous consequential domains, including employment decisions (Stoker, Garretsen, and Spreuwers 2016; Graham, Harvey, and Puri 2017; Swider, Harris, and Gong 2021), criminal sentencing (Wilson and Rule 2015; Johnson and King 2017), and the election of political candidates (Antonakis and Dalgas 2009; Olivola and Todorov 2010; Lenz and Lawson 2011; Jäckle et al. 2020). While facial impression biases may be consistent among a population, inferences of unobservable attributes such as character traits are inaccurate and often reflect societal stereotypes (Sutherland and Young 2022; Todorov 2017). AI systems are increasingly employed to automate or mediate access to information in domains such as hiring (Li et al. 2021), political analysis and advertising (Papakyriakopoulos et al. 2022), and law (Choi et al. 2023), and to the extent that such systems reflect facial impression biases, they may have socially undesirable impacts.

Relationship to Social Group Biases

Some studies suggest a connection between first-impression biases and demographic traits such as gender and ethnicity. Oh, Buck, and Todorov (2019) find that gender biases associating men with competence are reflected in participant impressions of the competence of faces. Xie et al. (2021) find that the structure of impressions of novel faces is predicted by learned social stereotypes about gender and race. Peterson et al. (2022) find that facial impression biases are correlated with demographics, such that judgments of traits like “cuteness” are related to age. The relationship between first-impression bias and social stereotypes can have real-world consequences. For example, White phenotypic prototypicality (looking like the average White person) can moderate use of force by police (Kahn et al. 2016).

Computational Models of Facial Impression Bias

Peterson et al. (2022) use the OMI dataset to model facial impressions using the StyleGAN-2 network (Karras, Laine, and Aila 2019), and demonstrate its capacity to manipulate faces such that the average U.S. perceiver would consider them similar to an attribute (such as trustworthiness). They build on research on the scientific modeling of facial impression biases, which commonly utilizes techniques including landmark annotations of faces (Turk and Pentland 1991), parametric three-dimensional mesh modeling (Banz and Vetter

2023), geometric morphological analysis (Sano and Kawabata 2023), and supervised deep learning models (Yu and Suchow 2022). As noted by Peterson et al. (2022), creating a computational model of a *bias* differs from modeling the attribute itself (*i.e.*, trying to predict if an individual is trustworthy from their face, rather than whether the average person would *perceive* an individual as trustworthy), which would amount to physiognomy (y Arcas, Mitchell, and Todorov 2023) for an unobservable attribute like trustworthiness.

CLIP and Vision-Language AI

The present work studies CLIP, a multimodal vision-language model pretrained using a symmetric cross-entropy loss (Oord, Li, and Vinyals 2018; Zhang et al. 2020) to pair images with associated text captions (Radford et al. 2021). After pretraining, CLIP can rank, retrieve, or classify images based on association with text classes specified at inference rather than pre-selected at the time of training, making it a “zero-shot” vision-language model (Radford et al. 2021), as well as a good source for semantically rich embeddings (Wolfe and Caliskan 2022b). CLIP is composed of a language model (usually GPT2 (Radford et al. 2019)), and an image encoder, such as a Vision Transformer (“ViT”) (Dosovitskiy et al. 2020) or a ResNet (He et al. 2016). The language and image models are jointly pretrained, and representations are projected into a multimodal embedding space, in which cosine similarity quantifies the similarity between image and text (Radford et al. 2021). In addition to standard CLIP models, we study “FaceCLIP” models trained by Zheng et al. (2022), who introduce Facial Representation Learning (FaRL), which combines CLIP training with a masked image modeling objective (Xie et al. 2022) and trains on a faces-only subset of LAION-400M. Models trained using FaRL set state-of-the-art on downstream facial analysis tasks such as face parsing (Zheng et al. 2022).

Text-to-Image Generators

CLIP is an essential component for many generative text-to-image models. One of the first uses of a CLIP model was to provide training supervision to OpenAI’s first DALL-E image generator model (Ramesh et al. 2021). Other text-to-image generators like VQGAN-CLIP similarly use CLIP embedding space measurements in their objective function (Crowson et al. 2022). More recent image generators such as Stable Diffusion 2 employ CLIP models as text encoders (Rombach et al. 2022a), passing CLIP text embeddings to a U-Net or similar latent diffusion architecture capable of generating an image conditioned on those text embeddings. More recently, DALL-E 3 (“unCLIP”) decodes images directly from a CLIP embedding space, translating CLIP text embeddings into image embeddings, and inverting them (Ramesh et al. 2022). As discussed in the Data section, we study open-weight text-to-image generators utilizing a CLIP text encoder.

Impact of Scale in Deep Learning and in CLIP

Research shows that the impact of data scale on deep learning models is empirically predictable (Hestness et al. 2017) and that task performance scales with training dataset size (Sun

et al. 2017; Brown et al. 2020). Zhai et al. (2022) empirically demonstrate that both model and data scale impact visual task performance, and set new state of the art on Imagenet (Deng et al. 2009) by efficiently scaling a ViT. In CLIP models, Cherti et al. (2022) demonstrate a relationship between pretraining data scale and task performance. Prior work also demonstrates increases in hate speech in CLIP models trained on larger uncurated datasets (Birhane et al. 2024).

Bias in Vision-Language AI

Prior research identifies societal biases in CLIP models, which have been found to over-represent images of men in retrieval contexts (Wang, Liu, and Wang 2021), to exhibit social stereotypes when used in robotics applications (Hundt et al. 2022), and to reflect biased cultural defaults related to race and nationality (Wolfe and Caliskan 2022a,c). Generative text to image models typically reflect biases of the CLIP model from which they learn or decode (Luccioni et al. 2023; Bianchi et al. 2023), as do vision-language chatbots that use CLIP as an image encoder (Fraser and Kiritchenko 2024). As with other machine learning models (Hall et al. 2022), bias in CLIP can be traced back to its pretraining data, and audits of LAION-400M have found examples of racial and gender biases (Birhane, Prabhu, and Kahembwe 2021).

Synthetic Media for Vision-Language Research

Our work builds on prior work employing synthetic images to study societal biases. Similar to the OMI dataset studied in this research, Wolfe, Banaji, and Caliskan (2022) assessed classification biases in CLIP related to images of multiracial individuals by generating images of thousands of faces using StyleGAN-2. More recently, Fraser and Kiritchenko (2024) use Midjourney (Midjourney 2024) to generate pairs of images of Black and White or Male and Female individuals in identical attire and situations, which they use to assess bias in conversational vision-language models such as Instruct-BLIP (Dai et al. 2023). The use of synthetic data is well-motivated for bias research, where it may be undesirable to assign values related to traits like trustworthiness or smugness to the faces of human subjects in a dataset intended for public release.

Data

This research uses the One Million Impressions (OMI) dataset and 43 English-language CLIP models trained on web-scraped text-and-image datasets, as well as three text-to-image generators employing CLIP models as text encoders.

The One Million Impressions Dataset

The OMI dataset is a collection of 1,004 images of human faces produced by Peterson et al. (2022) using StyleGAN-2 (Karras et al. 2020). Each face is rated by 30 or more human participants on Amazon Mechanical Turk for 34 attributes. For each attribute, participants rate the face on a sliding scale, where one end represents one pole of an attribute binary (such as “trustworthy”) and the other end represents the opposing pole of the binary (such as “untrustworthy”). The OMI dataset records the mean participant rating for each of the 34 attributes. Consistent with Peterson et al. (2022), we



Figure 2: Examples from the OMI dataset repository at <https://github.com/jcpeterson/omi>, used as stimuli in our research.

use these ratings as measurements of human bias at a societal scale, against which CLIP associations can be compared.

CLIP Training Data

We study CLIP models pretrained on one of five datasets, ordered from smallest to largest:

- **LAION-Face:** A 20-million sample subset of human faces and captions filtered from LAION-400M (see below) using RetinaFace (Deng et al. 2019) and intended for training facial analysis models (Zheng et al. 2022).
- **LAION-80M:** An 80-million sample subset of LAION-2B (see below) created by Cherti et al. (2022) to study scaling behavior in CLIP.
- **LAION-Aesthetics:** A 120-million sample subset of aesthetically pleasing images from LAION-5B as determined using a CLIP model (Schuhmann et al. 2022).
- **WebImageText (WIT):** A web-scraped corpus of 400 million images and captions, constructed by Radford et al. (2021) from a query list using Wikipedia and WordNet.
- **LAION-400M** An open source collection of 407 million image-text pairs intended to replicate the WIT dataset (Schuhmann et al. 2021).
- **LAION-2B:** An open source English-language dataset of 2.32 billion image-text pairs (Schuhmann et al. 2022).

Pretrained CLIP Models

This research studies the following CLIP models:

- **OpenAI CLIP:** 9 models pretrained by Radford et al. (2021) on the WIT dataset.
- **Scaling CLIP:** 29 models pretrained by Cherti et al. (2022) on LAION-80M, LAION-400M, and LAION-2B to study CLIP scaling behavior.
- **FaceCLIP:** 5 models trained on the LAION-Face dataset of Zheng et al. (2022). FaceCLIP models post-train from pretrained OpenAI CLIP-ViT models.

Pretrained Stable Diffusion Models

This research studies three Stable Diffusion (SD) models:

- **Stable Diffusion XL-Turbo:** A high-resolution text-to-image generator employing adversarial distillation diffusion to speed up the rate of image generation (Sauer et al. 2023). Uses a CLIP-ViT-L and CLIP-ViT-bigG for its text encoder, and pretrains on an internal dataset.
- **Runway Stable Diffusion 1.5:** A high-resolution text-to-image generator finetuned on LAION-Aesthetics (Romach et al. 2022a). Uses a CLIP-ViT-L-14 as the text encoder, and pretrains on LAION-5B.

- **Stable Diffusion 2:** A text-to-image generator using a CLIP-ViT-H as the text encoder, and pretraining on a filtered subset of LAION-5B (Rombach et al. 2022a).

Approach

We used embeddings from 43 CLIP models to compare facial impression biases measured in CLIP to biases measured in humans by Peterson et al. (Peterson et al. 2022), and extended subspace projection methods from prior work to study bias in generative text-to-image models.

Obtaining Image and Text Embeddings

We obtained image embeddings for the 1,004 images in the OMI dataset after projection to each CLIP model’s text-image latent space. Text embeddings use the “a photo of *image class*” prompt recommended by Radford et al. (2021). Because OMI consists of images of faces, we modify this prompt to “a photo of someone who is *attribute*.” In keeping with the binary sliding scale of Peterson et al. (2022), we computed an image’s association with each attribute by subtracting its cosine similarity with one pole of the attribute binary (“a photo of someone who has dark hair”) from its similarity with opposing pole (“a photo of someone who has light hair”). Formally, given a model m from which embeddings are obtained, the association m_j^a of an image vector \vec{i}_j at index j of the OMI dataset with an attribute a is the difference of the vector’s cosine similarity with a positive pole text vector \vec{t}_{a^+} and its cosine similarity with a negative pole text vector \vec{t}_{a^-} :

$$m_j^a = \cos(\vec{i}_j, \vec{t}_{a^+}) - \cos(\vec{i}_j, \vec{t}_{a^-}) \quad (1)$$

Adjusting Prompts for Negation

CLIP may fail to adjust for negation in text prompts (Parcalabescu et al. 2021) and can behave like a visual bag-of-words model (Yuksekgonul et al. 2022). For example, CLIP might match the text “a photo with no apples” to a photo of apples, due to how unlikely it is for a text caption (*i.e.*, CLIP’s training supervision) to describe something not present in the photo. To adjust for this, negative pole prompts were chosen such that they did not simply negate the positive class. For example, the “outgoing” attribute uses “a photo of someone who is shy” as the negative text class, rather than “a photo of someone who is not outgoing.” This strategy is not viable for some attributes, like those related to ethnicity, which instead use “a photo of someone” as the negative prompt.

Computing CLIP Model-Human Similarity

We denote the ordered set of $n=1,004$ OMI images as I . The vector of associations \mathbf{m}^a for a model m with an attribute a for all images $i \in I$ is given by:

$$\mathbf{m}^a = (m_0^a, m_1^a, \dots, m_{n-1}^a, m_n^a) \quad (2)$$

Similarly, the vector of human-rated associations \mathbf{h}^a for attribute a for all images $i \in I$ is given by:

$$\mathbf{h}^a = (h_0^a, h_1^a, \dots, h_{n-1}^a, h_n^a) \quad (3)$$

where h_j^a denotes the OMI mean for image \vec{i}_j at index j . The similarity s_m^a of bias in a model m for attribute a to human bias is given by Spearman’s ρ :

$$s_m^a = \rho(\mathbf{m}^a, \mathbf{h}^a) \quad (4)$$

CAT: Correlated Attribute Test

We compute the correlation between two attributes in a CLIP model using a simple test we call the CAT. As above, the vector of associations \mathbf{m}^a for a model m is given by:

$$\mathbf{m}^a = (m_0^a, m_1^a, \dots, m_{n-1}^a, m_n^a) \quad (5)$$

The measurement $\text{CAT}_m(a, b)$ between attributes a and b in a model m is given by Spearman’s ρ :

$$\text{CAT}_m(a, b) = \rho(\mathbf{m}^a, \mathbf{m}^b) \quad (6)$$

Subspace Projection for Text-to-Image Models

We draw on subspace projection methods used by Bolukbasi et al. (2016) and Omrani Sabbaghi, Wolfe, and Caliskan (2023) to measure facial impression biases in generative text-to-image models. First, we first obtain image embeddings for the 1,004 OMI images from the top layer of a ViT-Large-Patch32-384 model pretrained on ImageNet. For each attribute a , we learn a weights vector w^a predicting the OMI attribute ratings h^a , corresponding to a semantic subspace in the embeddings for the attribute. We then use a generative model g to generate n images via a prompt corresponding to either the positive (a^+) or negative (a^-) pole of an attribute. We embed each generated image g_j at position j with the ViT-Large-Patch32-384 to obtain the vector \vec{g}_j , and compute an attribute association g_j^a as its projection product with w^a :

$$g_j^a = \frac{\vec{g}_j \cdot w^a}{\|w^a\|} \quad (7)$$

The vector of associations g^a for a generative model g with an attribute a is given by:

$$\mathbf{g}^a = (g_0^a, g_1^a, \dots, g_{n-1}^a, g_n^a) \quad (8)$$

Experiments

Four experiments test the existence of human-like facial impression bias in vision-language AI, with consideration given to Human IRR, model and dataset scale, and downstream impact in image generation.

Model vs. Human Biases

We tested whether OpenAI, Scaling, and FaceCLIP CLIP models reflect human-like facial impression biases. We obtained the human-model similarity s_m^a for each model m with each attribute a for the 34 OMI attributes. We then compared the mean human-model similarity for each group of models to the Human IRR for the attribute reported by Peterson et al. (2022), calculating Pearson’s ρ between Human IRR and model-human similarity for each of the 34 attributes. A large coefficient indicates that the more societally consistent




Figure 3: The similarity of CLIP bias to human bias is strongly correlated with human IRR, indicating that the societal consistency of a bias plays a significant role in whether a model learns it during semi-supervised pretraining.

a facial impression bias (*i.e.*, as Human IRR increases), the more likely the bias is to be learned during semi-supervised CLIP training. We also computed Pearson’s ρ pairwise between OpenAI, Scaling, and FaRL models to assess whether models trained on different datasets learn similar biases.

Effects of Dataset Scale

We calculated human-model similarity s_m^a for each model m with each attribute a for the 34 attributes studied, and we constructed a multiple linear regression to predict the model-human similarity s_m^a for a given CLIP model m and an attribute a . We examined the 27 CLIP models trained by Cherti et al. (2022) and produced via the combination of three CLIP architectures (ViT-B32, ViT-B16, and ViT-L14), three dataset sizes (80M, 400M, 2B), and three total training example counts (3B, 13B, 34B). Independent variables include Human IRR (from Peterson et al. (Peterson et al. 2022)), as well as Dataset Size, Model Parameter Count, and Total Training Examples (from Cherti et al. (2022)). We normalized variables with range outside of $(0, 1)$ by dividing by their max. We conducted post hoc comparisons between each level of scale (80m, 400m, 2b), using paired-samples t -tests.

Structure of Facial Impressions

We computed the correlation matrix C_m for the 27 CLIP models studied in the Dataset Scale analysis by obtaining $CAT_m(a, b)$ for every attribute pair $a \in A$ and $b \in A$. We computed a corresponding matrix C_h by obtaining correlations for the same attributes using the human ratings of the OMI dataset. We then measured the similarity of C_m and C_h based on the normalized Frobenius inner product $F_{m,h}$. We used a one-way ANOVA to test for differences in $F_{m,h}$ between the 80M, 400M, and 2B models. We used a paired-samples t -test to conduct post hoc comparisons.

Peterson et al. (2022) study the structure of facial impressions by computing the correlation matrix of OMI ratings and qualitatively examining its hierarchical structure. The present research also qualitatively compares the structure of the OMI correlation matrix to the hierarchically clustered correlation matrix of a CLIP-ViT-L-14 trained on LAION-2B for 13-billion total samples. If the structure of OMI biases is similar

to CLIP, we expect to observe similar attribute clusters at higher levels, with differences emerging in leaf nodes.

Generative Text-to-Image Models

Finally, we extended the analysis to SDXL-Turbo, SD2, and Runway SD1.5. We first studied the similarity of these models’ representations to human attribute ratings. To do so, we generated $N=25$ images for each attribute’s positive pole and 25 for its negative pole. We adjusted prompts for the models to “a realistic portrait photo of someone who is *attribute*,” because image generators may create cartoonish images that do not center the face. We extracted embeddings for these images and computed their projection products g^{a+} and g^{a-} . If generative text-to-image models reflect human-like facial impression biases, we expect images generated from a positive prompt to have positive projections, and images from a negative prompt to have negative projections. We thus frame model-human similarity as a classification problem, wherein images generated from the positive pole prompt receives a label of 1, and from the negative pole prompt receive a label of 0. The OMI subspace is positioned as a classifier, which predicts 1 where an image vector’s projection product is positive, and 0 where it is negative. We report Recall, Precision, and F1 Score. We validated this approach using the Outdoors attribute, a control group for measuring validity in the OMI dataset, obtaining F1=.94 for SDXL-Turbo

We then study social bias in Stable Diffusion XL-Turbo by projecting the positive prompt images generated for the White and Black attributes onto all 34 of the OMI attribute subspaces. For each subspace, we compute the White-Black differential bias by obtaining an effect size (Cohen’s d) between the projection products g^{White^+} and g^{Black^+} , and we then measure statistical significance using a paired samples t -test.

Results

Our results indicate that 1) CLIP models exhibit facial impression biases; 2) Human IRR is a significant predictor of which biases are learned; 3) models trained on larger datasets exhibit emergence of subjective facial impression biases, and

Model-Human Similarity of Facial Impression Bias in CLIP Models by Pretraining Dataset Size									
Measurement	Mean (Std)			Max			Cohen's <i>d</i>		
Attribute	2b	400m	80m	2b	400m	80m	2b-80m	400m-80m	2b-400m
Happy	.88 (0.02)	.86 (0.03)	.76 (0.06)	.91	.90	.84	1.54*	1.36*	.73
Gender	.85 (0.02)	.87 (0.02)	.83 (0.04)	.87	.91	.89	.61	1.08*	-.84*
Age	.83 (0.08)	.84 (0.05)	.73 (0.11)	.93	.90	.90	.94	1.03*	-.04
Asian	.82 (0.03)	.81 (0.05)	.73 (0.06)	.85	.86	.79	1.40*	1.13*	.43
Hispanic	.82 (0.02)	.80 (0.02)	.67 (0.07)	.85	.84	.78	1.62*	1.50*	1.07*
Outdoors	.81 (0.01)	.79 (0.04)	.65 (0.10)	.82	.84	.75	1.51*	1.39*	.57
Pacific Islander	.76 (0.03)	.74 (0.04)	.62 (0.11)	.81	.81	.77	1.28*	1.15*	.50
Middle Eastern	.72 (0.02)	.68 (0.05)	.57 (0.11)	.75	.78	.71	1.43*	1.17*	.88
Native American	.68 (0.06)	.67 (0.06)	.53 (0.13)	.78	.78	.77	1.20*	1.15*	.14
Weight	.68 (0.03)	.69 (0.03)	.63 (0.05)	.72	.72	.67	1.05*	1.24*	-.40
Hair-Color	.66 (0.07)	.61 (0.09)	.50 (0.12)	.76	.70	.67	1.31*	.92	.67
Cute	.65 (0.10)	.67 (0.08)	.41 (0.15)	.78	.76	.56	1.35*	1.46*	-.31
Long-Haired	.63 (0.10)	.63 (0.11)	.43 (0.16)	.76	.75	.67	1.20*	1.16*	.06
Gay	.49 (0.07)	.41 (0.06)	.30 (0.11)	.57	.51	.42	1.46*	1.09*	1.14*
Attractive	.48 (0.08)	.50 (0.11)	.28 (0.21)	.60	.65	.57	1.11*	1.11*	-.16
Electable	.47 (0.22)	.50 (0.11)	.30 (0.25)	.68	.65	.60	.72*	.94	-.16
Smart	.42 (0.15)	.45 (0.14)	.25 (0.11)	.62	.59	.50	1.09*	1.23*	-.20
Black	.41 (0.09)	.35 (0.11)	.37 (0.12)	.51	.49	.51	.46	-.10	.59
Smug	.38 (0.11)	.19 (0.20)	.07 (0.15)	.50	.45	.25	1.50*	.63	1.02
Trustworthy	.36 (0.18)	.04 (0.18)	-.06 (0.15)	.59	.35	.15	1.56*	.60	1.33*
Skin-Color	.36 (0.14)	.37 (0.13)	.28 (0.11)	.59	.58	.43	.58	.70	-.11
Outgoing	.33 (0.13)	.21 (0.17)	.16 (0.16)	.52	.35	.42	1.02*	.31	.74
Privileged	.26 (0.15)	.21 (0.12)	.05 (0.07)	.49	.44	.16	1.35*	1.31*	.33
Godly	.20 (0.09)	.26 (0.17)	.25 (0.10)	.34	.51	.42	-.57	.01	-.42
Liberal	.14 (0.15)	.22 (0.15)	.23 (0.19)	.32	.45	.50	-.49	-.05	-.51
Typical	.13 (0.10)	.11 (0.16)	.05 (0.13)	.28	.28	.20	.72	.46	.16
Dorky	.13 (0.23)	.17 (0.22)	.05 (0.21)	.38	.52	.33	.37	.52	-.15
Familiar	.12 (0.12)	.04 (0.08)	-.00 (0.14)	.32	.14	.16	.90	.42	.75
Well-Groomed	.12 (0.26)	-.04 (0.21)	.11 (0.18)	.43	.44	.31	.03	-.71	.62
Dominant	.11 (0.23)	.06 (0.18)	.00 (0.29)	.47	.34	.29	.41	.25	.23
Memorable	.08 (0.11)	-.05 (0.15)	.04 (0.12)	.24	.13	.18	.35	-.67	.92
White	.00 (0.28)	.12 (0.21)	-.03 (0.20)	.45	.45	.26	.12	.67	-.46
Looks-Like-You	-.13 (0.14)	-.10 (0.14)	-.11 (0.16)	.15	.06	.11	-.17	.08	-.26
Alert	-.29 (0.10)	-.17 (0.10)	-.14 (0.09)	-.16	-.03	.05	-1.25*	-.31	-1.05

Table 1: 17 of 34 OMI attributes exhibit statistically significant differences and large effect sizes between model groups.




Figure 4: CLIP models exhibit significant Spearman's ρ between Mean Model-Human Similarity and OMI IRR.

more human-like associations among impressions; and 4) text-to-image generators exhibit facial impression biases and undesirable social biases associating preferred attributes with images of White individuals.

CLIP Models Reflect Human Biases

As shown in Figure 1, OpenAI, FaceCLIP, and Scaling CLIP models exhibit human-like facial impression biases. Relatively objective attributes like age, hair-color, and happiness

exhibit high model-human similarity, as do many socially constructed attributes such as gender and cuteness. Notable exceptions include Black, White, and skin color attributes, which fall short of expectations based on Human IRR, as visualized in Figure 3. Model-human similarities of traits like Trustworthiness, Electability, and Intelligence are statistically significant but lower, consistent with the lower IRR of these attributes. That CLIP learns these biases at all is noteworthy: to our knowledge, this is the first research to document unobservable facial impression biases learned by a semi-supervised vision-language model (rather than a supervised model of facial impression biases) that are consistent with human societal biases.

As described in Figure 4, all three families of models exhibit strong correlations ranging from .72 to .76 between the mean model-human similarity of a trait and its Human IRR. Coefficients are larger between OpenAI and FaceCLIP models than with Scaling CLIP models, likely a result of FaceCLIP post-training from OpenAI base models.

Dataset Scale

A multiple linear regression finds that only Human IRR and Dataset Size are statistically significant predictors of model-human similarity. As described in Table 2, Human IRR plays the larger role of the two independent variables, as evidenced by a much larger coefficient and *t*-value. Total Training Samples, Image Parameters, and Text Parameters are not statistically significant predictors of facial impression bias in




Figure 5: The structure of facial impression biases in CLIP-ViT-L-14 mirrors that of human facial impression biases quantified in the OMI dataset. Clusters related to ethnicity emerge in each, as do clusters grouping gender, sexuality, and smugness.

Statistic	
Adj. R^2	.420
F-Statistic	134.0
N	918
DoF Residuals	912
DoF Model	5

Ind. Variable	Coef	t	$p < t $
Human IRR	1.1204	25.470	.001
Dataset Size	.0877	4.484	.001
Total Samples	-.0132	-.599	.549
Image Params	-.5615	-.150	.881
Text Params	.7898	.144	.885
Constant	-.7542	-.434	.665

Table 2: Fitting a linear regression to model-human similarity coefficients reveals that Human IRR and Dataset Scale are significant predictors a bias will be learned by a CLIP model.

vision-language models.

Table 1 describes the mean (with standard deviation) and maximum model-human similarity for each dataset size, and it reports Cohen’s d between the groups of models trained on the three dataset sizes. Effect sizes obtained between the 400M and the 80M level are large and statistically significant for 17 out of 34 attributes, with large absolute differences between attribute means, such as .41 for Cute at the 80M level vs. .67 at the 400M level. For most attributes, comparisons are not statistically significant between the 2B and 400M levels, though they may return small or medium effect sizes. A notable exception emerges for several unobservable attributes, including Trustworthiness and Sexuality, which exhibit large effect sizes and statistically significant differences between the 2B level and the 400M level. While observ-

able attributes such as Happiness reflect human ratings well at the 80M level, and more subjective but still visually observable attributes like Cute are reflected consistently at the 400M level, it is not until the 2B level that CLIP models reflect subjective and visually *unobservable* attributes such as Trustworthiness. The results indicate that increases in the scale of the pretraining data have more significant effects for learning subtle societal biases reflecting attributes with *lower* IRR. Models trained on additional data approximate a distribution that more closely reflects the perceptions of society as a whole, learning to use the biased visual heuristics present in the human-authored captions in the pretraining data, even for unobservable attributes.

Differences in $F_{m,h}$ by Dataset Scale			
Groups Compared	Cohen’s d	t	$p <$
2b-400m	.89	2.76	.05
400m-80m	1.30	4.10	.01
2b-80m	1.54	4.38	.01

Table 3: The structure of facial impression bias is more similar to human bias in CLIP models trained on larger datasets.

Structure of Facial Impressions

Results indicate that dataset size impacts the extent to which the *structure* of facial impression bias in CLIP reflects the structure of the facial impression bias in humans. Figure 5 visualizes the hierarchical similarities between the attribute cross-correlation matrix for CLIP-ViT-L-14 (the most commonly used CLIP model as of this writing) and the OMI attribute cross-correlation matrix. The most salient similarities between the two include a cluster of correlated racial and ethnic identities, such as Hispanic, Middle-Eastern, Native American, and Pacific Islander, as well as a cluster grouping

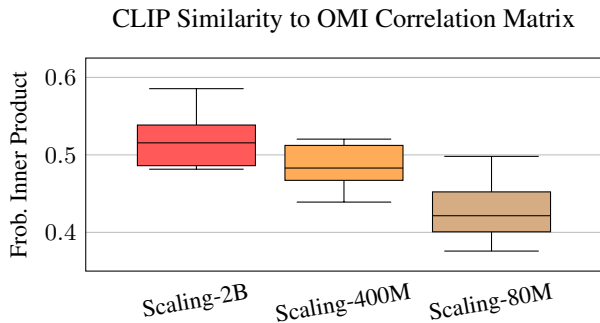


Figure 6: Scaling-2B CLIP models exhibit the greatest structural similarity to human facial impression biases.

together the Smugness, Gender, Sexuality, and Age attributes. There are also differences between the model and human ratings: while Trustworthy is correlated with Cute in OMI, it is correlated with Smart and Happy in CLIP. Similarities between CLIP attribute correlations and OMI attribute correlations are more evident at the higher levels of the hierarchy, with differences appearing toward the leaf nodes.

A one-way ANOVA provides evidence that similarities in the structure of facial impression biases increase for models trained on larger datasets, with $F(2) = 15.71, p < .001$ after Bonferonni correction, demonstrating statistically significant differences in $F_{m,h}$ between the Scaling-2B, Scaling-400M, and Scaling-80M models. As described in Table 3, we observe statistically significant differences and large effect sizes both between the 80M and 400M levels and between the 400M and 2B levels. Figure 6 further illustrates the differences among the three model groups, showing that the magnitude of difference is greater between the 80M level and the 400M level than between the 400M level and the 2B level.

Generative Text-to-Image Models

As shown in Table 4, we observe more variance in text-to-image generator F1 scores than in CLIP model-human similarities. SDXL has the most human-like associations, with Spearman’s $\rho = .60, p < .001$ between Human IRR and SDXL F1 scores. Runway-SD1.5 is also correlated with IRR, with $\rho = .50, p < .01$, while SD2 is not significantly correlated with IRR, with $\rho = .32, p = .06$. Notably, the Attractive attribute has the highest F1 score for SDXL, whereas it is the 16th most human-similar attribute in CLIP models, suggesting the importance of representing beauty for user-facing image generators, which often undergo additional training to better reflect user aesthetic preferences (Rombach et al. 2022b). With the exception of Liberal, unobservable traits rank in the bottom half of F1 scores for SDXL, and Trustworthy is lowest of any trait. Though SD2 and Runway-SD1.5 are more human-like than SDXL for trustworthiness, the results suggest that exploiting biased heuristics may be more straightforward for a classifier like CLIP.

Images generated by Stable Diffusion XL-Turbo also bear signs of racial bias: as shown in Figure 7, we observe statistically significant differences indicating that generated images of White individuals are more likely to be perceived as domi-

Attribute F1 Scores by Stable Diffusion Model			
Attribute	SDXL	SD2	Runway-SD1.5
attractive	0.98	0.7	0.65
outdoors	0.94	0.67	0.60
well-groomed	0.91	0.66	0.59
hair-color	0.83	0.79	0.77
weight	0.76	0.57	0.37
long-haired	0.76	0.67	0.65
black	0.75	0.76	0.76
white	0.74	0.43	0.53
asian	0.70	0.63	0.69
middle-eastern	0.69	0.70	0.64
cute	0.69	0.58	0.55
happy	0.68	0.79	0.77
islander	0.68	0.66	0.69
age	0.68	0.68	0.72
liberal	0.68	0.60	0.62
skin-color	0.67	0.68	0.61
alert	0.67	0.57	0.59
gender	0.66	0.64	0.62
smart	0.65	0.57	0.71
dominant	0.65	0.68	0.67
hispanic	0.65	0.66	0.69
native	0.64	0.69	0.68
electable	0.61	0.65	0.60
dorky	0.53	0.60	0.60
looks-like-you	0.53	0.68	0.49
smug	0.51	0.58	0.57
memorable	0.48	0.66	0.59
privileged	0.47	0.60	0.58
gay	0.37	0.56	0.40
godly	0.31	0.65	0.59
typical	0.29	0.63	0.63
familiar	0.27	0.68	0.51
outgoing	0.24	0.66	0.71
trustworthy	0.12	0.51	0.64

Table 4: SDXL-Turbo reflects human facial impression biases, with Spearman’s $\rho = .60$ between IRR and SDXL F1.

nant, privileged, memorable, attractive, electable, and happy than images of Black individuals, which are more likely to be perceived as more liberal and heavy (vs. thin). Note that many of these relationships are *not* observed in the OMI correlation clusters seen in Figure 5, indicating that they originate not with the OMI dataset but with the text-to-image model.

Discussion

Our results make clear the inter-connection of visual perception in AI with the human social world: where a facial impression bias is more consistently shared among humans, CLIP is also more likely to learn it. Facial impression biases have notable consequences in professional and civic life (Stoker, Garretsen, and Spreeuwers 2016), and prove difficult to dislodge even after intervention (Jaeger et al. 2020). While CLIP may serve as a tool for studying such biases, it may also reinforce or amplify these biases in society, especially given their presence in user-facing image generators.

Scale and Bias

Training on larger datasets results in emergent and amplified facial impression biases. CLIP models exhibit more human-like biases related to trustworthiness and sexuality when trained on LAION-2B, and nearly every OMI attribute increases in model-human similarity between the 80M level and the 400M level. That model parameterization plays no detectable role in facial impression bias underlines that what

Differential White-Black Bias in Stable-Diffusion-XL-Turbo




Figure 7: Stable Diffusion XL-Turbo exhibits differential White-Black biases, such that generated images of White individuals project as more dominant, electable, and attractive. We plot effect size and significance ($* < .05$, $** < .01$, $*** < .001$) below each significant comparison.

CLIP models have learned is a biased visual heuristic reflected in the training data, not a more precise representation of an objectively detectable attribute. This also suggests that greater attention to the characteristics of the training data is advisable when pretraining CLIP systems intended for use in downstream applications without fine-tuning, or for supervising other models, including text to image generators using CLIP as a text encoder.

Ramifications of Human-like Models

Amid the excitement over vision-language models like GPT-4 that convincingly imitate aspects of human intelligence (Bubeck et al. 2023), the emergence of subtle biases in multi-modal models trained on the largest datasets elicits a corrective question: is more “human-like” always better? Approximating the distribution of societal associations present in the pretraining dataset as completely as possible may be useful for providing a more general-purpose model. However, as zero-shot vision-language models continue to become more accessible to the general public, including via conversational interfaces that mimic text-based interaction with a human being (Li et al. 2023; Liu et al. 2024), monitoring for subtle emergent biases in need of mitigation is likely to become a more pressing concern. Humans interacting with fluent, mostly debiased models may in fact be less skeptical of *subtle* reflections of societal bias than of the more blatant and clearly offensive misrepresentations of demographic groups seen in previous generations of models.

Implications for Computational Social Science

The present work is also notable in its consequences for computational social science. By our reckoning, Peterson et al. (2022) spent tens of thousands of dollars to collect the human subject data needed to learn a supervised model of facial impressions; CLIP produces a model of facial impressions as a side effect of pretraining. While CLIP provides a less precise model of these biases than the supervised counterpart of Peterson et al. (2022), our research nonetheless

suggests that CLIP models might play a role similar to static word embeddings, which social scientists now routinely employ in computational studies of human attitudes, including in research that generalizes the findings of human subjects experiments (Morehouse et al. 2023; Caliskan et al. 2022), or which quantifies shifts in human attitudes over decades (Borenstein et al. 2023; Garg et al. 2018). That CLIP models reflect facial impression bias suggests that they could model other complex sociocultural phenomena not observable via text embeddings alone, a possibility that might be further explored in future work.

Limitations

While participants in the study of Peterson et al. (2022) were reflective of U.S. demographics as a whole, this also means that a majority of perceivers identified as White, as is clear from the correlation of White and Looks-Like-You OMI attributes in Figure 5. Such a demographic skew may render attribute ratings sensitive to correlated biases, given that prior work observes a relationship between social bias and face impressions (Xie et al. 2021). In addition, while we adopt now-standard prompts specified by OpenAI when introducing CLIP Radford et al. (2021), significantly changing these prompts may induce variance in the results.

Conclusion

The present work demonstrates that three families of CLIP models reflect human facial impression biases, and that human agreement and dataset scale predict how faithfully a CLIP model reproduces a bias. Our findings illustrate the importance of pretraining data for fair and usable zero-shot vision-language models, and they further underscore the interconnection of human bias and AI representations. While concerning for zero-shot applications of CLIP, especially where the model may be used in a user-facing system, our results also demonstrate the remarkable utility of the model for the computational study of societal bias.

Ethical Considerations

Caution is warranted in a discussion of the findings of this research. While CLIP models trained on larger amounts of data more faithfully encode a human perceptual bias, this does not thereby mean that the models have accurately made an inference about a person based solely on an image of their face. The present work produces no artifact intended for predicting a person’s attributes, and *should not be interpreted* as scientific support for the use of vision-language models to predict unobservable attributes. Rather, it should serve as a scientific treatment of emergent social bias arising from large-scale pretraining.

Researcher Positionality

Our research team includes two individuals with extensive experience in the quantitative study of AI bias and ethics, as well as two individuals with extensive experience in human-computer interaction and value-sensitive design. The team consists of two members who identify as men, and two who identify as women, one of whom is a woman of color. All team members are researchers at the same university in the United States. We sought to include a diversity of perspectives and kinds of expertise given the sensitive subject matter of our work.

Adverse Impacts

While we have not trained a new model or built any system to facilitate facial analysis, it is possible that an individual may nonetheless attempt to use the methods described in this paper to make unwarranted inferences about individuals based solely on their face. However, these inferences would certainly be less precise than those obtained using a supervised model trained explicitly for facial analysis. Our approach to these potential adverse impacts is to make clear that facial impressions constitute *biases*, and that they intersect with other harmful demographic biases. We hope that demonstrating the presence of these biases in CLIP and Stable Diffusion will serve as a warning of the potential societal effects of using zero-shot vision-language models to classify or represent human beings.

References

- Antonakis, J.; and Dalgas, O. 2009. Predicting elections: Child’s play! *Science*, 323(5918): 1183–1183.
- Be My Eyes. 2023a. Announcing ‘Be My AI,’ Soon Available for Hundreds of Thousands of Be My Eyes Users. <https://www.bemyeyes.com/blog/announcing-be-my-ai>. [Accessed 02-29-2024].
- Be My Eyes. 2023b. See the world together. <https://www.bemyeyes.com/>. [Accessed 02-29-2024].
- Bianchi, F.; Kalluri, P.; Durmus, E.; Ladhak, F.; Cheng, M.; Nozza, D.; Hashimoto, T.; Jurafsky, D.; Zou, J.; and Caliskan, A. 2023. Easily accessible text-to-image generation amplifies demographic stereotypes at large scale. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 1493–1504.
- Birhane, A.; Dehdashtian, S.; Prabhu, V.; and Boddeti, V. 2024. The Dark Side of Dataset Scaling: Evaluating Racial Classification in Multimodal Models. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*, 1229–1244.
- Birhane, A.; Prabhu, V. U.; and Kahembwe, E. 2021. Multimodal datasets: misogyny, pornography, and malignant stereotypes. *arXiv preprint arXiv:2110.01963*.
- Blanz, V.; and Vetter, T. 2023. A morphable model for the synthesis of 3D faces. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, 157–164. Association for Computing Machinery.
- Bolukbasi, T.; Chang, K.-W.; Zou, J. Y.; Saligrama, V.; and Kalai, A. T. 2016. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. *Advances in neural information processing systems*, 29: 4349–4357.
- Borenstein, N.; Stańczak, K.; Rolskov, T.; Perez, N. d. S.; Käfer, N. K.; and Augenstein, I. 2023. Measuring Intersectional Biases in Historical Documents. *arXiv preprint arXiv:2305.12376*.
- Brown, T. B.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.
- Bubeck, S.; Chandrasekaran, V.; Eldan, R.; Gehrke, J.; Horvitz, E.; Kamar, E.; Lee, P.; Lee, Y. T.; Li, Y.; Lundberg, S.; et al. 2023. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*.
- Caliskan, A.; Ajay, P. P.; Charlesworth, T.; Wolfe, R.; and Banaji, M. R. 2022. Gender bias in word embeddings: A comprehensive analysis of frequency, syntax, and semantics. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*, 156–170.
- Charlesworth, T. E.; Hudson, S.-k. T.; Cogsdill, E. J.; Spelke, E. S.; and Banaji, M. R. 2019. Children use targets’ facial appearance to guide and predict social behavior. *Developmental Psychology*, 55(7): 1400.
- Cherti, M.; Beaumont, R.; Wightman, R.; Wortsman, M.; Ilharco, G.; Gordon, C.; Schuhmann, C.; Schmidt, L.; and Jitsev, J. 2022. Reproducible scaling laws for contrastive language-image learning. *arXiv preprint arXiv:2212.07143*.
- Choi, J. H.; Hickman, K. E.; Monahan, A.; and Schwarcz, D. 2023. Chatgpt goes to law school. *Available at SSRN*.
- Crowson, K.; Biderman, S.; Kornis, D.; Stander, D.; Hallahan, E.; Castricato, L.; and Raff, E. 2022. VQGAN-CLIP: Open Domain Image Generation and Editing with Natural Language Guidance. *arXiv preprint arXiv:2204.08583*.
- Dai, W.; Li, J.; Li, D.; Tiong, A. M. H.; Zhao, J.; Wang, W.; Li, B.; Fung, P.; and Hoi, S. 2023. InstructBLIP: Towards General-purpose Vision-Language Models with Instruction Tuning. *arXiv:2305.06500*.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.

- Deng, J.; Guo, J.; Zhou, Y.; Yu, J.; Kotsia, I.; and Zafeiriou, S. 2019. Retinaface: Single-stage dense face localisation in the wild. *arXiv preprint arXiv:1905.00641*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*.
- Fraser, K. C.; and Kiritchenko, S. 2024. Examining Gender and Racial Bias in Large Vision-Language Models Using a Novel Dataset of Parallel Images. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (EACL)*.
- Garg, N.; Schiebinger, L.; Jurafsky, D.; and Zou, J. 2018. Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences*, 115(16): E3635–E3644.
- Graham, J. R.; Harvey, C. R.; and Puri, M. 2017. A corporate beauty contest. *Management Science*, 63(9): 3044–3056.
- Hall, M.; van der Maaten, L.; Gustafson, L.; Jones, M.; and Adcock, A. 2022. A systematic study of bias amplification. *arXiv preprint arXiv:2201.11706*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hestness, J.; Narang, S.; Ardalani, N.; Diamos, G.; Jun, H.; Kianinejad, H.; Patwary, M. M. A.; Yang, Y.; and Zhou, Y. 2017. Deep learning scaling is predictable, empirically. *arXiv preprint arXiv:1712.00409*.
- Hill, K. 2023. OpenAI Worries About What Its Chatbot Will Say About People’s Faces. *The New York Times*.
- Hundt, A.; Agnew, W.; Zeng, V.; Kacianka, S.; and Gombolay, M. 2022. Robots enact malignant stereotypes. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, 743–756.
- Jäckle, S.; Metz, T.; Wenzelburger, G.; and König, P. D. 2020. A catwalk to congress? Appearance-based effects in the elections to the US House of Representatives 2016. *American Politics Research*, 48(4): 427–441.
- Jaeger, B.; Todorov, A. T.; Evans, A. M.; and van Beest, I. 2020. Can we reduce facial biases? Persistent effects of facial trustworthiness on sentencing decisions. *Journal of Experimental Social Psychology*, 90: 104004.
- Johnson, B. D.; and King, R. D. 2017. Facial profiling: Race, physical appearance, and punishment. *Criminology*, 55(3): 520–547.
- Kahn, K. B.; Goff, P. A.; Lee, J. K.; and Motamed, D. 2016. Protecting whiteness: White phenotypic racial stereotypicality reduces police use of force. *Social Psychological and Personality Science*, 7(5): 403–411.
- Karras, T.; Laine, S.; and Aila, T. 2019. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4401–4410.
- Karras, T.; Laine, S.; Aittala, M.; Hellsten, J.; Lehtinen, J.; and Aila, T. 2020. Analyzing and Improving the Image Quality of StyleGAN. In *Proc. CVPR*.
- Lenz, G. S.; and Lawson, C. 2011. Looking the part: Television leads less informed citizens to vote based on candidates’ appearance. *American Journal of Political Science*, 55(3): 574–589.
- Li, J.; Li, D.; Savarese, S.; and Hoi, S. 2023. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*.
- Li, L.; Lassiter, T.; Oh, J.; and Lee, M. K. 2021. Algorithmic hiring in practice: Recruiter and HR Professional’s perspectives on AI use in hiring. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, 166–176.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2024. Visual instruction tuning. *Advances in neural information processing systems*, 36.
- Luccioni, A. S.; Akiki, C.; Mitchell, M.; and Jernite, Y. 2023. Stable bias: Analyzing societal representations in diffusion models. *arXiv preprint arXiv:2303.11408*.
- Midjourney. 2024. Midjourney. <https://www.midjourney.com/home>. [Accessed 03-01-2024].
- Morehouse, K.; Rouduri, V.; Cunningham, W.; and Charlesworth, T. 2023. Traces of Human Attitudes in Contemporary and Historical Word Embeddings (1800-2000). *Research Square Preprint*.
- Oh, D.; Buck, E. A.; and Todorov, A. 2019. Revealing hidden gender biases in competence impressions of faces. *Psychological Science*, 30(1): 65–79.
- Olivola, C. Y.; and Todorov, A. 2010. Elected in 100 milliseconds: Appearance-based trait inferences and voting. *Journal of nonverbal behavior*, 34: 83–110.
- Omran, S.; Wolfe, R.; and Caliskan, A. 2023. Evaluating biased attitude associations of language models in an intersectional context. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, 542–553.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Papakyriakopoulos, O.; Tessono, C.; Narayanan, A.; and Kshirsagar, M. 2022. How algorithms shape the distribution of political advertising: Case studies of Facebook, Google, and TikTok. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*, 532–546.
- Parcalabescu, L.; Cafagna, M.; Muradjan, L.; Frank, A.; Calixto, I.; and Gatt, A. 2021. VALSE: A task-independent benchmark for vision and language models centered on linguistic phenomena. *arXiv preprint arXiv:2112.07566*.
- Peterson, J. C.; Uddenberg, S.; Griffiths, T. L.; Todorov, A.; and Suchow, J. W. 2022. Deep models of superficial face judgments. *Proceedings of the National Academy of Sciences*, 119(17): e2115228119.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.;

- et al. 2021. Learning transferable visual models from natural language supervision. *arXiv preprint arXiv:2103.00020*.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I.; et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8): 9.
- Ramesh, A.; Dhariwal, P.; Nichol, A.; Chu, C.; and Chen, M. 2022. Hierarchical Text-Conditional Image Generation with CLIP Latents. *arXiv preprint arXiv:2204.06125*.
- Ramesh, A.; Pavlov, M.; Goh, G.; Gray, S.; Voss, C.; Radford, A.; Chen, M.; and Sutskever, I. 2021. Zero-shot text-to-image generation. *arXiv preprint arXiv:2102.12092*.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022a. High-Resolution Image Synthesis With Latent Diffusion Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10684–10695.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022b. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Sano, T.; and Kawabata, H. 2023. A computational approach to investigating facial attractiveness factors using geometric morphometric analysis and deep learning. *Scientific reports*, 13(1): 19797.
- Sauer, A.; Lorenz, D.; Blattmann, A.; and Rombach, R. 2023. Adversarial diffusion distillation. *arXiv preprint arXiv:2311.17042*.
- Schuhmann, C.; Beaumont, R.; Vencu, R.; Gordon, C.; Wightman, R.; Cherti, M.; Coombes, T.; Katta, A.; Mullis, C.; Wortsman, M.; et al. 2022. Laion-5b: An open large-scale dataset for training next generation image-text models. *arXiv preprint arXiv:2210.08402*.
- Schuhmann, C.; Vencu, R.; Beaumont, R.; Kaczmarczyk, R.; Mullis, C.; Katta, A.; Coombes, T.; Jitsev, J.; and Komatsuzaki, A. 2021. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*.
- Stoker, J. I.; Garretsen, H.; and Spreeuwers, L. J. 2016. The facial appearance of CEOs: Faces signal selection but not performance. *PloS one*, 11(7): e0159950.
- Sun, C.; Shrivastava, A.; Singh, S.; and Gupta, A. 2017. Revisiting unreasonable effectiveness of data in deep learning era. In *Proceedings of the IEEE international conference on computer vision*, 843–852.
- Sutherland, C. A.; and Young, A. W. 2022. Understanding trait impressions from faces. *British Journal of Psychology*, 113(4): 1056–1078.
- Swider, B. W.; Harris, T. B.; and Gong, Q. 2021. First impression effects in organizational psychology. *Journal of Applied Psychology*.
- Todorov, A. 2017. Face value. In *Face Value*. Princeton University Press.
- Turk, M.; and Pentland, A. 1991. Eigenfaces for recognition. *Journal of cognitive neuroscience*, 3(1): 71–86.
- Wang, J.; Liu, Y.; and Wang, X. E. 2021. Are Gender-Neutral Queries Really Gender-Neutral? Mitigating Gender Bias in Image Search. *arXiv preprint arXiv:2109.05433*.
- Willis, J.; and Todorov, A. 2006. First impressions: Making up your mind after a 100-ms exposure to a face. *Psychological science*, 17(7): 592–598.
- Wilson, J. P.; and Rule, N. O. 2015. Facial trustworthiness predicts extreme criminal-sentencing outcomes. *Psychological science*, 26(8): 1325–1331.
- Wolfe, R.; Banaji, M. R.; and Caliskan, A. 2022. Evidence for Hypodescent in Visual Semantic AI. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*.
- Wolfe, R.; and Caliskan, A. 2022a. American==White in Multimodal Language-and-Image AI. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*.
- Wolfe, R.; and Caliskan, A. 2022b. Contrastive Visual Semantic Pretraining Magnifies the Semantics of Natural Language Representations. *Association for Computational Linguistics*.
- Wolfe, R.; and Caliskan, A. 2022c. Markedness in Visual Semantic AI. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*.
- Xie, S. Y.; Flake, J. K.; Stolier, R. M.; Freeman, J. B.; and Hehman, E. 2021. Facial impressions are predicted by the structure of group stereotypes. *Psychological Science*, 32(12): 1979–1993.
- Xie, Z.; Zhang, Z.; Cao, Y.; Lin, Y.; Bao, J.; Yao, Z.; Dai, Q.; and Hu, H. 2022. Simmim: A simple framework for masked image modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9653–9663.
- y Arcas, B. A.; Mitchell, M.; and Todorov, A. 2023. Physiognomy in the Age of AI. *Feminist AI: Critical Perspectives on Algorithms, Data, and Intelligent Machines*, 208.
- Yu, Y.; and Suchow, J. 2022. Deep tensor factorization models of first impressions. In *SVRHM 2022 Workshop@ NeurIPS*.
- Yuksekgonul, M.; Bianchi, F.; Kalluri, P.; Jurafsky, D.; and Zou, J. 2022. When and Why Vision-Language Models Behave like Bags-Of-Words, and What to Do About It? In *The Eleventh International Conference on Learning Representations*.
- Zhai, X.; Kolesnikov, A.; Houlsby, N.; and Beyer, L. 2022. Scaling vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12104–12113.
- Zhang, Y.; Jiang, H.; Miura, Y.; Manning, C. D.; and Langlotz, C. P. 2020. Contrastive learning of medical visual representations from paired images and text. *arXiv preprint arXiv:2010.00747*.
- Zheng, Y.; Yang, H.; Zhang, T.; Bao, J.; Chen, D.; Huang, Y.; Yuan, L.; Chen, D.; Zeng, M.; and Wen, F. 2022. General facial representation learning in a visual-linguistic manner. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18697–18709.