Image Representations Learned With Unsupervised Pre-Training Contain Human-like Biases

Ryan Steed
ryansteed@cmu.edu
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA

Aylin Caliskan
aylin@gwu.edu
George Washington University
Washington, District of Columbia, USA

ABSTRACT

Recent advances in machine learning leverage massive datasets of unlabeled images from the web to learn general-purpose image representations for tasks from image classification to face recognition. But do unsupervised computer vision models automatically learn implicit patterns and embed social biases that could have harmful downstream effects? We develop a novel method for quantifying biased associations between representations of social concepts and attributes in images. We find that state-of-the-art unsupervised models trained on ImageNet, a popular benchmark image dataset curated from internet images, automatically learn racial, gender, and intersectional biases. We replicate 8 previously documented human biases from social psychology, from the innocuous, as with insects and flowers, to the potentially harmful, as with race and gender. Our results closely match three hypotheses about intersectional bias from social psychology. For the first time in unsupervised computer vision, we also quantify implicit human biases about weight, disabilities, and several ethnicities. When compared with statistical patterns in online image datasets, our findings suggest that machine learning models can automatically learn bias from the way people are stereotypically portrayed on the web.

CCS CONCEPTS

• **Computing methodologies** → *Unsupervised learning; Transfer learning; Machine learning.*

KEYWORDS

implicit bias, unsupervised learning, computer vision

ACM Reference Format:

Ryan Steed and Aylin Caliskan. 2021. Image Representations Learned With Unsupervised Pre-Training Contain Human-like Biases. In *Conference on Fairness, Accountability, and Transparency (FAccT '21)*, March 3–10, 2021, Virtual Event, Canada. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3442188.3445932>

1 INTRODUCTION

Can machines learn social biases from the way people are portrayed in image datasets? Companies and researchers regularly use machine learning models trained on massive datasets of images scraped

Figure 1: Unilever using AI-powered job candidate assessment tool HireVue [35].

from the web for tasks from face recognition [40] to image classification [66]. To reduce costs, many practitioners use state-of-the-art models “pre-trained” on large datasets to help solve other machine learning tasks, a powerful approach called *transfer learning* [68]. For example, HireVue used similar state-of-the-art computer vision and natural language models to evaluate job candidates’ video interviews, potentially discriminating against candidates based on race, gender, or other social factors [35]. In this paper, we show how models trained on unlabeled images scraped from the web embed human-like biases, including racism and sexism.

Where most bias studies focus on supervised machine learning models, we seek to quantify learned patterns of implicit social bias in unsupervised image representations. Studies in supervised computer vision have highlighted social biases related to race, gender, ethnicity, sexuality, and other identities in tasks including face recognition, object detection, image search, and visual question answering [10, 43, 47, 52, 61, 76]. These algorithms are used in important real-world settings, from applicant video screening [35, 60] to autonomous vehicles [28, 52], but their harmful downstream effects have been documented in applications such as online ad delivery [67] and image captioning [38].

Our work examines the growing set of computer vision methods in which no labels are used during model training. Recently, pre-training approaches adapted from language models have dramatically increased the quality of unsupervised image representations [3, 12–15, 20, 36, 50]. With *fine-tuning*, practitioners can pair these general-purpose representations with labels from their domain to accomplish a variety of supervised tasks like face recognition or image captioning. We hypothesize that 1) like their counterparts in language, these unsupervised image representations also contain human-like social biases, and 2) these biases correspond to stereotypical portrayals of social group members in training images.

Results from natural language support this hypothesis. Several studies show that word embeddings, or representations, learned

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

FAccT '21, March 3–10, 2021, Virtual Event, Canada

ACM ISBN 978-1-4503-8309-7/21/03.

<https://doi.org/10.1145/3442188.3445932>

automatically from the way words co-occur in large text corpora exhibit human-like biases [7, 11, 26]. Word embeddings acquire these biases via statistical regularities in language that are based on the co-occurrence of stereotypical words with social group signals. Recently, new deep learning methods for learning context-specific representations sharply advanced the state-of-the-art in natural language processing (NLP) [19, 58, 59]. Embeddings from these pre-trained models can be fine-tuned to boost performance in downstream tasks such as translation [23, 24]. As with static embeddings, researchers have shown that embeddings extracted from contextualized language models also exhibit downstream racial and gender biases [4, 34, 69, 79].

Recent advances in NLP architectures have inspired similar unsupervised computer vision models. We focus on two state-of-the-art, pre-trained models for image representation, iGPT [14] and SimCLRv2 [13]. We chose these models because they hold the highest fine-tuned classification scores, were pre-trained on the same large dataset of Internet images, and are publicly available. iGPT, or Image GPT, borrows its architecture from GPT-2 [59], a state-of-the-art unsupervised language model. iGPT learns representations for pixels (rather than for words) by pre-training on many unlabeled images [14]. SimCLRv2 uses deep learning to construct image representations from ImageNet by comparing augmented versions of the training images [13, 15].

Do these unsupervised computer vision models embed human biases like their counterparts in natural language? If so, what are the origins of this bias? In NLP, embedding biases have been traced to word co-occurrences and other statistical patterns in text corpora used for training [6, 9, 11]. Both our models are pre-trained on ImageNet 2012, the most widely-used dataset of curated images scraped from the web [64]. In image datasets and image search results, researchers have documented clear correlations between the presence of individuals of a certain gender and the presence of stereotypical objects; for instance, the category “male” co-occurs with career and office related content such as ties and suits whereas “female” more often co-occurs with flowers in casual settings [43, 73]. As in NLP, we expect that these patterns of bias in the pre-training dataset will result in implicitly embedded bias in unsupervised models, even without access to labels during training.

This paper presents the Image Embedding Association Test (iEAT), the first systematic method for detecting and quantifying social bias learned automatically from unlabeled images.

- We find statistically significant racial, gender, and intersectional biases embedded in two state-of-the-art unsupervised image models pre-trained on ImageNet [64], iGPT [14] and SimCLRv2 [13].
- We test for 15 previously documented human and machine biases that have been studied for decades and validated in social psychology and conduct the first machine replication of Implicit Association Tests (IATs) with picture stimuli [31].
- In 8 tests, our machine results match documented human biases, including 4 of 5 biases also found in large language models. The 7 tests which did not show significant human-like biases are from IATs with only small samples of picture stimuli.
- With 16 novel tests, we show how embeddings from our model confirm several hypotheses about intersectional bias from social psychology [29].

- We compare our results to statistical analyses of race and gender in image datasets. Unsupervised models seem to learn bias from the ways people are commonly portrayed in images on the web.
- We present a qualitative case study of how image generation, a downstream task utilizing unsupervised representations, exhibits a bias towards the sexualization of women.

2 RELATED WORK

Various tests have been constructed to quantify bias in unsupervised natural language models [4, 11, 48, 79], but to our knowledge, there are no principled tests for measuring bias embedded in *unsupervised* computer vision models. Wang et al. [73] develop a method to automatically recognize bias in visual datasets but still rely on human annotations. Our method uses no annotations whatsoever. In NLP, there are several systematic approaches to measuring unsupervised bias in word embeddings [8, 11, 34, 45, 48, 69]. Most of these tests take inspiration from the well-known IAT [31, 32]. Participants in the IAT are asked to rapidly associate stimuli, or exemplars, representing two target concepts (e.g. “flowers” and “insects”) with stimuli representing evaluative attributes (e.g. “pleasant” and “unpleasant”) attribute [31]. Assuming that the cognitive association task is easier when the strength of implicit association between the target concept and attributes is high, the IAT quantifies bias as the latency of response [31] or the rate of classification error [53]. Stimuli may take the form of words, pictures, or even sounds [55], and there are several IATs with picture-only stimuli [55].

Notably, Caliskan et al. [11] adapt the heavily-validated IAT [31] from social psychology to machines by testing for the mathematical association of word embeddings rather than response latency. They present a systematic method for measuring language biases associated with social groups, the Word Embedding Association Test (WEAT). Like the IAT, the WEAT measures the effect size of bias in static word embeddings by quantifying the relative associations of two sets of target stimuli (e.g., {“woman,” “female”} and {“man,” “male”}) that represent social groups with two sets of evaluative attributes (e.g., {“science,” “mathematics”} and {“arts,” “literature”}). For validation, two WEATs quantify associations towards flowers vs. insects and towards musical instruments vs. weapons, both accepted baselines Greenwald et al. [31]. Greenwald et al. [31] refer to these baseline biases as “universally” accepted stereotypes since they are widely shared across human subjects and are not potentially harmful to society. Other WEATs measure social group biases such as sexist and racist associations or negative attitudes towards the elderly or people with disabilities. In any modality, implicit biases can potentially be prejudiced and harmful to society. If downstream applications use these representations to make consequential decisions about human beings, such as automated video job interview evaluations, machine learning may perpetuate existing biases and exacerbate historical injustices [18, 60].

The original WEAT [11] uses *static* word embedding models such as word2vec [49] and GloVe [57], each trained on Internet-scale corpora composed of billions of tokens. Recent work extends the WEAT to *contextualized* embeddings: dynamic representations based on the context in which a token appears. May et al. [48] insert targets and attributes into sentences like “This is a[n] <word>” and applying WEAT to the vector representation for the whole sentence,

with the assumption that the sentence template used is “semantically bleached” (such that the only meaningful content in the sentence is the inserted word). Tan and Celis [69] extract the contextual word representation for the token of interest before pooling to avoid confounding effects at the sentence level; in contrast, Bommasani et al. [8] find that pooling tends to improve representational quality for bias evaluation. Guo and Caliskan [34] dispense with sentence templates entirely, pooling across n word-level contextual embeddings for the same token extracted from random sentences. Our approach is closest to these latter two methods, though we pool over images rather than words.

3 APPROACH

In this paper, we adapt bias tests designed for contextualized word embeddings to the image domain. While language transformers produce contextualized *word* representations to solve the next *token* prediction task, an image transformer model like iGPT generates *image* representations to solve the next *pixel* prediction task [14]. Unlike words and tokens, pixels do not explicitly correspond to semantic concepts (objects or categories) as words do. In language, a single token (e.g. “love”) corresponds to the target concept or attribute (e.g. “pleasant”). But in images, no single pixel corresponds to a semantically meaningful concept. To address the abstraction of semantic representation in the image domain, we propose the Image Embedding Association Test (iEAT), which modifies contextualized word embedding tests to compare pooled image-level embeddings. The goal of the iEAT is to measure the biases embedded during unsupervised pre-training by comparing the relative association of image embeddings in a systematic process. Chen et al. [14] and Chen et al. [15] show through image classification that unsupervised image features are good representations of object appearance and categories; we expect they will also embed information gleaned from the common co-occurrence of certain objects and people and therefore contain related social biases.

Our approach is summarized in Figure 2. The iEAT uses the same formulas for the test statistic, effect size d , and p -value as the WEAT [11], described in Section 3.3. Section 3.1 summarizes our approach to replicating several different IATs; Section 3.2 describes several novel intersectional iEATs. Section 3.3 describes our test statistic, drawn from embedding association tests like the WEAT.

3.1 Replication of Bias Tests

In this paper, we validate the iEAT by replicating as closely as possible several common IATs. These tests fall into two broad categories: valence tests, in which two target concepts are tested for association with “pleasant” and “unpleasant” images; and stereotype tests, in which two target concepts are tested for association with a pair of stereotypical attributes (e.g. “male” vs. “female” “career” vs. “family”). To closely match the ground-truth human IAT data and validate our method, our replications use the same concepts as the original IATs (listed in Table 1). Because some IATs rely on verbal stimuli, we adapt them to images, using image stimuli from the IATs when available. When no previous studies use image stimuli, we map the non-verbal stimuli to images using the data collection method described in Section 5.

Figure 2: Example iEAT replication of the Insect-Flower IAT [31], which measures the differential association between flowers vs. insects and pleasantness vs. unpleasantness.

Many of these bias tests have been replicated for machines in the language domain; for the first time, we also replicate tests with image-only stimuli, including the Asian and Native American IATs. Most of these tests were originally administered in controlled laboratory settings [31, 32], and all except for the Insect-Flower IAT have also been tested on the Project Implicit website at <http://projectimplicit.org> [32, 33, 54]. Project Implicit has been available worldwide for over 20 years; in 2007, the site had collected more than 2.5 million IATs. The average effect sizes (which are based on samples so large the power is nearly 100%) for these tests are reproduced in Table 1. To establish a principled methodology, all the IAT verbal and original image stimuli for our bias tests were replicated exactly from this online IAT platform [56]. We will treat these results, along with the laboratory results from the original experiments [31], as ground-truth for human biases that serve as validation benchmarks for our methods (Section 6).

3.2 Intersectional iEATs

We also introduce several new tests for intersectional valence bias and bias at the intersection of gender stereotypes and race. Intersectional stereotypes are often even more severe than their constituent stereotypes [17]. Following Tan and Celis [69], we anchored comparison on White males, the group with the most representation, and compared against White females, Black males, and Black females, respectively (Table 2). Drawing on social psychology [29], we pose three hypotheses about intersectional bias:

- *Intersectionality hypothesis*: tests at the intersection of gender and race will reveal emergent biases not explained by the sum of biases towards race and gender alone.
- *Race hypothesis*: biases between racial groups will be more similar to differential biases between the men than between the women.

- *Gender hypothesis*: biases between men and women will be most similar to biases between White men and White women.

3.3 Embedding Association Tests

Though our stimuli are images rather than words, we can use the same statistical method for measuring biased associations between image representations [11] to quantify a standardized effect size of bias. We follow Caliskan et al. [11] in describing the WEAT here.

Let X and Y be two sets of target concepts embeddings of size N_t , and let A and B be two sets of attribute embeddings of size N_a . For example, the Gender-Career IAT tests for the differential association between the concepts “male” (A) and “female” (B) and the attributes “career” (X) and “family” (Y). Generally, experts in social psychology and cognitive science select stimuli that are typically representative of various concepts. In this case, A contains embeddings for verbal stimuli such as “boy,” “father,” and “man,” while X contains embeddings for verbal stimuli like “office” and “business.” These linguistic, visual, and sometimes auditory stimuli are proxies for the aggregate representation of a concept in cognition. Embedding association tests use these unambiguous stimuli as semantic representations to study biased associations between the concepts being represented. Since the stimuli are chosen by experts to most accurately represent concepts, they are not polysemous or ambiguous tokens. We use these expert-selected stimuli as the basis for our tests in the image domain.

The test statistic measures the differential association of the target concepts X and Y with the attributes A and B

$$s(X, Y, A, B) = \sum_{x \in X} s(x, A, B) - \sum_{y \in Y} s(y, A, B) \quad (1)$$

where $s(w, A, B)$ is the differential association of w with the attributes, quantified by the cosine similarity of vectors

$$s(w, A, B) = \text{mean}_{a \in A} \cos(w, a) - \text{mean}_{b \in B} \cos(w, b)$$

We test the significance of this association with a permutation test¹ over all possible equal-size partitions $\{(X_i, Y_i)\}_i$ of $X \cup Y$ to generate a null hypothesis as if no biased associations existed. The one-sided p -value measures the unlikelihood of the null hypothesis

$$p = \Pr[s(X_i, Y_i, A, B) > s(X, Y, A, B)]$$

and the effect size, a standardized measure of the separation between the relative association of X and Y with A and B , is

$$d = \frac{\text{mean}_{x \in X} s(x, A, B) - \text{mean}_{y \in Y} s(y, A, B)}{\text{sd}_{w \in X \cup Y} s(w, A, B)}$$

A larger effect size indicates a larger differential association; for instance, the large effect size d in Table 1 for the gender-career bias example above indicates that in human respondents, “male” is strongly associated with “career” attributes compared to “female,” which is strongly associated with “family” attributes. Note that these effect sizes cannot be directly compared to effect sizes in human IATs, but the significance levels are uniformly high. Human IATs measure individual people’s associations; embedding association tests measure the aggregate association in the representation space learned from the training set. In general, significance increases with

the number of stimuli; an insignificant result does not necessarily indicate a lack of bias.

One important assumption of the iEAT is that categories can be meaningfully represented by groups of images, such that the association bias measured refers to the categories of interest and not some other, similar-looking categories. Thus, a positive test result indicates only that there is an association bias between the corresponding samples’ sets of target images and attribute images. To generalize to associations between abstract social concepts requires that the samples adequately represent the categories of interest. Section 5 details our procedure for selecting multiple, representative stimuli, following validated approaches from prior work [31].

We use an adapted version of May et al. [48]’s Python WEAT implementation. All code, pre-trained models, and data used to produce the figures and results in this paper can be accessed at github.com/ryansteed/ieat.

4 COMPUTER VISION MODELS

To explore what kinds of biases may be embedded in image representations generated in unsupervised settings, where class labels are not available for images, we focus on two computer vision models published in summer 2020, iGPT and SimCLRv2. We extract representations of image stimuli with these two pre-trained, unsupervised image representation models. We choose these particular models because they achieve state-of-the-art performance in *linear evaluation* (a measure of the accuracy of a linear image classifier trained on embeddings from each model). iGPT is the first model to learn from pixel co-occurrences to generate image samples and perform image completion tasks.

4.0.1 Pre-training Data. Both models are pre-trained on ImageNet 2012, a large benchmark dataset for computer vision tasks [64].² ImageNet 2012 contains 1.2 million annotated images of 200 object classes, including a person class; even if the annotated object is not a person, a person may appear in the image. For this reason, we expect the models to be capable of generalizing to stimuli containing people [63, 64]. While there are no publicly available pre-trained models with larger training sets, and the “people” category of ImageNet is no longer available, this dataset is a widely used benchmark containing a comprehensive sample of images scraped from the web, primarily Flickr [64]. We assume that the portrayals of people in ImageNet are reflective of the portrayal of people across the web at large, but a more contemporary study is left to future work. CIFAR-100, a smaller classification database, was also used for linear evaluation and stimuli collection [44].

4.0.2 Image Representations. Both models are *unsupervised*: neither use any labels during training. Unsupervised models learn to produce embeddings based on the implicit patterns in the entire training set of image features. Both models incorporate neural networks with multiple hidden layers (each learning a different level of abstraction) and a projection layer for some downstream task. For linear classification tasks, features can be drawn directly from layers in the base neural network. As a result, there are various ways to

¹We use an exact, non-parametric permutation test over all possible partitions. There are no normality assumptions about the distribution of the null hypothesis.

²Both models were tested on the Tensorflow version of ILSVRC 2012, available at <https://www.tensorflow.org/datasets/catalog/imagenet2012>.

extract image representations, each encoding a different set of information. We follow Chen et al. [14] and Chen et al. [15] in choosing the features for which linear evaluation scores are highest such that the features extracted contain high-quality, general-purpose information about the objects in the image. Below, we describe the architecture and feature extraction method for each model.

4.1 iGPT

The Image Generative Pre-trained Transformer (iGPT) model is a novel, NLP-inspired approach to unsupervised image representation. We chose iGPT for its high linear evaluation scores, minimalist architecture, and strong similarity to GPT-2 [59], a transformer-based architecture that has found great success in the language domain. Transformers learn patterns in the way individual tokens in an input sequence appear with other tokens in the sequence [71]. Chen et al. [14] apply a structurally simple, highly parameterized version of the GPT-2 generative language pre-training architecture [59] to the image domain for the first time. GPT-2 uses the “contextualized embeddings” learned by a transformer to predict the next token in a sequence and generate realistic text [59]. Rather than autoregressively predict the next entry in a sequence of tokens as GPT-2 does, iGPT predicts the next entry in a flattened sequence of pixels. iGPT is trained to autoregressively complete cropped images, and feature embeddings extracted from the model can be used to train a state-of-the-art linear classifier [14].

We use the largest open-source version of this model, iGPT-L 32x32, with $L = 48$ layers and embedding size 1536. All inputs are restricted to 32x32 pixels; the largest model, which takes 64x64 input, is not available to the public. Original code and checkpoints for this model were obtained from its authors at github.com/openai/image-gpt. iGPT is composed of L blocks

$$\begin{aligned} n^l &= \text{layer_norm}(h^l) \\ a^l &= h^l + \text{multihead_attention}(n^l) \\ h^{l+1} &= a^l + \text{mlp}(\text{layer_norm}(a^l)) \end{aligned}$$

where h^l is the input tensor to the l^{th} block. In the final layer, called the *projection head*, Chen et al. [14] learn a projection from $n^L = \text{layer_norm}(h^L)$ to a set of logits parameterizing the conditional distributions across the sequence dimension. Because this final layer is designed for autoregressive pixel prediction, the final layer may not contain the optimal representations for object recognition tasks. Chen et al. [14] obtain the best linear classification results using embeddings extracted from a middle layer - specifically, somewhere near the 20th layer [14]. A linear classifier trained on these features is much more accurate than one trained on the next-pixel embeddings [14]. Such “high-quality” features from the middle of the network f^l are obtained by average-pooling the layer norm across the sequence dimension:

$$f^l = \langle n_i^l \rangle_i \quad (2)$$

Chen et al. [14] then learn a set of *class* logits from f^l for their fine-tuned, supervised linear classifier, but we will just use the embeddings f^{20} . In general, we prefer these embeddings over embeddings from other layers for two reasons: 1) they can be more closely compared to the SimCLRv2 embeddings, which are also

optimal for fine-tuning a linear classifier; 2) we hypothesize that embeddings with higher linear evaluation scores will also be more likely to embed biases, since stereotypical portrayals typically incorporate certain objects and scenes (e.g. placing men with sports equipment). In Appendix C, we try another embedding extraction strategy and show that this hypothesis is correct.

4.2 SimCLR

The Simple Framework for Contrastive Learning of Visual Representations (SimCLR) [15, 16] is another state-of-the-art unsupervised image classifier. We chose SimCLRv2 because it has a state-of-the-art open source release and for variety in architecture: unlike iGPT, SimCLRv2 utilizes a traditional neural network for image encoding, ResNet [37]. SimCLRv2 extracts representations in three stages: 1) data augmentation (random cropping, random color distortions, and Gaussian blur); 2) an encoder network, ResNet [37]; 3) mapping to a latent space for contrastive learning, which maximizes agreement between the different augmented views [15]. These representations can be used to train state-of-the-art linear image classifiers [15, 16]. We use the largest pre-trained open-source version (the model with the highest linear evaluation scores) of SimCLRv2 [16], obtained from its authors at github.com/google-research/simclr. This pre-trained model uses a 50-layer ResNet with width $3\times$ and selective kernels (which have been shown to increase linear evaluation accuracy), and it was also pre-trained on ImageNet [64].

As with iGPT, we extract the embeddings identified by Chen et al. [15] as “high-quality” features for linear evaluation. Following [15], let $\tilde{x}_i$ and $\tilde{x}_j$ be two data augmentations (random cropping, random color distortion, and random Gaussian blur) of the same image. The base encoder network $f(\cdot)$ is a network of L layers

$$h_i = f(\tilde{x}_i) = \text{ResNet}(\tilde{x}_i) \quad (3)$$

where $h_i \in \mathbb{R}^d$ is the output after the average pooling layer. During pre-training, SimCLRv2 utilizes an additional layer: a projection head $g(\cdot)$ that maps h_i to a latent space for contrastive loss. The contrastive loss function can be found in [15].

After pre-training, Chen et al. [15] discard the projection head $g(\cdot)$, using the average pool output $f(\cdot)$ for linear evaluation. Note that the projection head $g(h)$ is still necessary for pre-training high-quality representations (it improves linear evaluation accuracy by over 10%); but Chen et al. [15] find that training on h rather than $z = g(h)$ also improves linear evaluation accuracy by more than 10%. We follow suit, using h_i (the average pool output of ResNet) to represent our image stimuli, which has dimensionality 2,048. High dimensionality is not a great obstacle; association tests have been used with embeddings as large as 4,096 dimensions [48].

5 STIMULI

To replicate the IATs, we systematically compiled a representative set of image stimuli for each of the concepts, or categories, listed in Table 1. Rather than attempting to specify and justify new constructs, we adhere as closely as possible to stimuli defined and employed by well-validated psychological studies. For each category (e.g. “male” or “science”) in each IAT (e.g. Gender-Science), we drew representative images from either 1) the original IAT stimuli,

if the IAT used picture stimuli [56], 2) the CIFAR-100 dataset [44], or 3) a Google Image Search.

This section describes how we obtained a set of images that meaningfully represent some target concept (e.g. “male”) or attribute (e.g. “science”) as it is normally, or predominantly, portrayed in society and on the web. We follow the stimuli selection criteria outlined in foundational prior work to collect the most typical and accurate exemplars [31, 32]. For picture-IATs with readily available image stimuli, we accept those stimuli as representative and exactly replicate the IAT conditions, with two exceptions: 1) the weapon-tool IAT picture stimuli include outdated objects (e.g. cutlass, Walkman), so we chose to collect an additional, modernized set of images; 2) the disability IAT utilizes abstract symbols, so we collected a replacement set of images of real people for consistency with the training set. For IATs with verbal stimuli, we use Google Image Search as a proxy for the predominant portrayal of words (expressed as search terms) on the web (described in Section 5.1). Human IATs employ the same philosophy: for example, the Gender-Science IAT uses common European American names to represent male and female, because the majority of names in the U.S. are European American [54]. We follow the same approach in replicating the human IATs for machines in the vision domain.

One consequence of the stimuli collection approach outlined in Section 5.1 is that our test set will be biased towards certain demographic groups, just as the Human IATs are biased towards European American names. For example, Kay et al. [43] showed that in 2015, search results for powerful occupations like CEO systematically under-represented women. In a case like this, we would expect to underestimate bias towards minority groups. For example, since we expect Gender-Science biases to be higher for non-White women, a test set containing more White women than non-White would exhibit lower overall bias than a test set containing an equal number of stimuli from white and non-White women. Consequently, tests on Google Image Search stimuli would be expected to result in under-estimated stereotype-congruent bias scores. While under-representation in the test set does not pose a major issue for measuring normative concepts, we cannot use the same datasets to test for intersectional bias. For those iEATs, we collected separate, equal-sized sets of images with search terms based on the categories White male, White female, Black male, and Black female, since none of the IATs specifically target these intersectional groups.

5.1 Verbal to Image Stimuli

One key challenge of our approach is representing social constructs and abstract concepts such as “male” or “pleasantness” in images. A Google Image Search for “pleasantness” returns mostly cartoons and pictures of the word itself. We address this difficulty by adhering as closely as possible to the verbal IAT stimuli, to ensure the validity of our replication. In verbal IATs, this is accomplished with “buckets” of verbal exemplars that include a variety of common-place and easy-to-process realizations of the concept in question. For example, in the Gender-Science IAT, the concept “male” is defined by the verbal stimuli “man,” “son,” “father,” “boy,” “uncle,” “grandpa,” “husband,” and “male” [77]. To closely match the representations

tested by these IATs, we use these sets of words to search for substitute image stimuli that portray one of these words or phrases. For the vast majority of exemplars, we were able to find direct visualizations of the stimuli as an isolated person, object, or scene. For example, Figure 2 depicts sample image stimuli corresponding to the verbal stimuli “orchid” (for category “flower”), “centipede” (“insect”), “sunset” (“pleasant”), and “morgue” (“unpleasant”).³

We collected images for each verbal stimulus from either CIFAR-100⁴ or Google Image Search according to a systematic procedure detailed in Appendix B. This procedure controls for image characteristics that might confound the category we are attempting to define (e.g. lighting, background, dominant colors, placement) in several ways: 1) we collected more than one for each verbal stimulus, in case of idiosyncrasies in the images collected; 2) for stimuli referring to an object or person, we chose images that isolated the object or person of interest against a plain background, unless the object filled the whole image; 3) when an attribute stimulus refers to a group of people, we chose only images where the target concepts were evenly represented in the attribute images;⁵ 4) for the picture-IATs, we accepted the original image stimuli to exactly reconstruct the original test conditions. We also did not alter the original verbal stimuli, relying instead on the construct validity of the original IAT experiments.⁶ For each verbal stimulus, Appendix B lists corresponding search terms and the precise number of images collected. All the images used to represent the concepts being tested are available at github.com/ryansteed/ieat.

5.2 Choosing Valence Stimuli

Valence, the intrinsic pleasantness or goodness of things, is one of the principal dimensions of affect and cognitive heuristics that shape attitudes and biases [31]. Many IATs quantify implicit bias by comparing two social groups to the valence attributes “pleasant” vs. “unpleasant.” Here, positive valence will denote “pleasantness” and negative valence will denote “unpleasantness.” The verbal exemplars for valence vary slightly from test to test. Rather than create a new set of image stimuli for each valence IAT, we collected one, large consolidated set from an experimentally validated database [5] of low and high valence words (e.g. “rainbow,” “morgue”) commonly used in the valence IATs. To quantify norms, [5] asked human participants to rate these non-social words for “pleasantness” and “imagery” in a controlled laboratory setting. Because some of the words for valence do not correspond to physical objects, we collected images for verbal stimuli with high valence and imagery scores. We used the same procedure as for all the other verbal stimuli (described above in Section 5.1). The full list of verbal valence stimuli can be found in Appendix A.

³In the original IATs, the category set sizes N_I and N_A range from 5-15 exemplars. We collected $n \approx 5$ images for each exemplar such that N_I and N_A are 30-50. Significance could be increased by including more stimuli, at the risk of diluting the test set with less-representative images from farther down in the search results.

⁴We first check for test images in CIFAR-100 because iGPT performs well in out-of-sample linear evaluation on this dataset [15].

⁵For example, for the “family” attribute in the Gender-Career test, we chose only images of families with equal numbers of men and women.

⁶One exception: the Gender-Career IAT used specific male- and female-sounding names, rather than general exemplars like “man” or “father” as in the Gender-Science IAT. We use the general exemplars for both tests.

6 EVALUATION

We evaluate the validity of iEAT by comparing the results to human and natural language biases measured in prior work. We obtain stereotype-congruent results for baseline, or “universal,” biases. We also introduce a simple experiment to test how often the iEAT incorrectly finds bias in a random set of stimuli.

Predictive Validity. We posit that iEAT results have predictive validity if they correspond to ground-truth IAT results for humans or WEAT results in word embeddings. In this paper, we validate the iEAT by replicating several human IATs as closely as possible (as described in Section 5) and comparing the results. We find that embeddings extracted from at least one of the two models we test display significant bias for 8 of the 15 ground-truth human IATs we replicate (Section 7). The insignificant biases are likely due to small sample sizes. We also find evidence supporting each of the intersectional hypotheses listed in Section 3.2, which have also been empirically validated in a study with human participants [29].

Baselines. As a baseline, we replicate a “universal” bias test presented in the first paper introducing the IAT [31]: the association between flower vs. insects and pleasant vs. unpleasant. If human-like biases are encoded in unsupervised image models, we would expect a strong and statistically significant flower-insect valence bias, for two reasons: 1) as Greenwald et al. [31] conjecture, this test measures a close-to-universal baseline human bias; 2) our models (described in Section 4) achieve state-of-the-art performance when classifying simple objects including flowers and bees.⁷ The presence of universal bias and absence of random bias suggests our conclusions are valid for other social biases.

Specificity. Prior work on embedding association tests does not evaluate the false positive rate. To validate the specificity of our significance estimation, we created 1,000 random partitions of $X \cup Y \cup A \cup B$ from the flower-insect test to evaluate true positive detection. Our false positive rate is roughly bounded by the p -value: 10.3% of these random tests resulted in a false positive at $p < 10^{-1}$; 1.2% were statistically significant false positives at $p < 10^{-2}$.

7 EXPERIMENTS AND RESULTS

In correspondence with the human IAT, we find several significant racial biases and gender stereotypes, including intersectional biases, shared by both iGPT and SimCLRv2 when pre-trained on ImageNet.

7.1 iEATs

Effect sizes and p -values from the permutation test for each bias type measurement are reported in Table 1 and interpreted below.

7.1.1 Widely Accepted Biases. First, we apply the iEAT to the widely accepted baseline Insect-Flower IAT, which measures the association of insects and flowers with pleasantness and unpleasantness, respectively. As hypothesized, we find that embeddings from both models contain significant positive biases in the same direction as the human participants, associating flowers with pleasantness and insects with unpleasantness, with $p < 10^{-1}$ (Table 1). Notably, the magnitude of bias is greater for SimCLRv2 (effect size 1.69, $p < 10^{-3}$) than for iGPT (effect size 0.34, $p < 10^{-1}$). In general,

SimCLRv2 embeddings contain stronger biases than iGPT embeddings but do not contain as many kinds of bias. We conjecture that because SimCLRv2 transforms images before training (including color distortion and blurring) and is more architecturally complex than iGPT [15], its embeddings become more suitable for concrete object classification as opposed to implicit social patterns.

7.1.2 Racial Biases. Both models display statistically significant racial biases, including both valence and stereotype biases. The racial attitude test, which measures the differential association of images of European Americans vs. African Americans with pleasantness and unpleasantness, shows no significant biases. But embeddings extracted from both models exhibit significant bias for the Arab-Muslim valence test, which measures the association of images of Arab-Americans vs. others with pleasant vs. unpleasant images. Also, embeddings extracted with iGPT exhibit strong bias large effect size (effect size 1.26, $p < 10^{-2}$) for the Skin Tone test, which compares valence associations with faces of lighter and darker skin tones. These findings relate to anecdotal examples of software that claim to make faces more attractive by lightening their skin color. Both iGPT and SimCLRv2 embeddings also associate White people with tools and Black people with weapons in both classical and modernized versions of the Weapon IAT.

7.1.3 Gender Biases. There are statistically significant gender biases in both models, though not for both stereotypes we tested. In the Gender-Career test, which measures the relative association of the category “male” with career attributes like “business” and “office” and the category “female” with family-related attributes like “children” and “home,” embeddings extracted from both models exhibit significant bias (iGPT effect size 0.62, $p < 10^{-2}$, SimCLRv2 effect size 0.74, $p < 10^{-3}$). This finding parallels Kay et al. [43]’s observation that image search results for powerful occupations like CEO systematically under-represented women. In the Gender-Science test, which measures the association of “male” with “science” attributes like math and engineering and “female” with “liberal arts” attributes like art and writing, only iGPT displays significant bias (effect size 0.44, $p < 10^{-1}$).

7.1.4 Other Biases. For the first time, we attempt to replicate several other tests measuring weight stereotypes and attitudes towards the elderly or people with disabilities. iGPT displays an additional bias (effect size 1.67, $p = 10^{-4}$) towards the association of thin people with pleasantness and overweight people with unpleasantness. We found no significant bias for the Native American or Asian American stereotype tests, the Disability valence test, or the Age valence test. For reference, significant age biases have been detected in static word embeddings; the others have not been tested because they use solely image stimuli [11]. Likely, the target sample sizes for these tests are too low; all three of these tests use picture stimuli from the original IAT, which are all limited to fewer than 10 images. Replication with an augmented test set is left to future work. Note that lack of significance in a test, even if the sample size is sufficiently large, does not indicate the embeddings from either model are definitively bias-free. While these tests did not *confirm* known human biases regarding foreigners, people with disabilities, and the elderly, they also did not *contradict* any known human-like biases.

⁷A linear image classifier trained on iGPT embeddings reaches 88.5% accuracy on CIFAR-100; SimCLRv2 embeddings reach 89% accuracy [14].

Table 1: iEAT tests for the association between target concepts X vs. Y (represented by n_t images each) and attributes A vs. B (represented by n_a images each) in embeddings generated by an unsupervised model. Effect sizes d represent the magnitude of bias, colored by conventional small (0.2), medium (0.5), and large (0.8). Permutation p -values indicate significance. Reproduced from Nosek et al. [56], the original human IAT effect sizes are all statistically significant with $p < 10^{-8}$; they can be compared to our effect sizes in sign but not in magnitude.

	X	Y	A	B	n_t	n_a	Model	iEAT d	iEAT p	IAT d
Age [†]	Young	Old	Pleasant	Unpleasant	6	55	iGPT	0.42	0.24	1.23
							SimCLR	0.59	0.16	1.23
Arab-Muslim	Other	Arab-Muslim	Pleasant	Unpleasant	10	55	iGPT	0.86	0.02	0.33
							SimCLR	1.06	$< 10^{-2}$	0.33
Asian [§]	European American	Asian American	American	Foreign	6	6	iGPT	0.25	0.34	0.62
							SimCLR	0.47	0.21	0.62
Disability [†]	Disabled	Abled	Pleasant	Unpleasant	4	55	iGPT	-0.02	0.53	1.05
							SimCLR	0.38	0.34	1.05
Gender-Career	Male	Female	Career	Family	40	21	iGPT	0.62	$< 10^{-2}$	1.1
							SimCLR	0.74	$< 10^{-3}$	1.1
Gender-Science	Male	Female	Science	Liberal Arts	40	21	iGPT	0.44	0.02	0.93
							SimCLR	-0.10	0.67	0.93
Insect-Flower	Flower	Insect	Pleasant	Unpleasant	35	55	iGPT	0.34	0.07	1.35
							SimCLR	1.69	$< 10^{-3}$	1.35
Native [§]	European American	Native American	U.S.	World	8	5	iGPT	-0.33	0.73	0.46
							SimCLR	-0.19	0.65	0.46
Race [†]	European American	African American	Pleasant	Unpleasant	6	55	iGPT	-0.62	0.85	0.86
							SimCLR	-0.57	0.83	0.86
Religion	Christianity	Judaism	Pleasant	Unpleasant	7	55	iGPT	0.37	0.25	-0.34
							SimCLR	0.36	0.26	-0.34
Sexuality	Gay	Straight	Pleasant	Unpleasant	9	55	iGPT	-0.03	0.52	0.74
							SimCLR	0.04	0.47	0.74
Skin-Tone [†]	Light	Dark	Pleasant	Unpleasant	7	55	iGPT	1.26	$< 10^{-2}$	0.73
							SimCLR	-0.19	0.71	0.73
Weapon [§]	White	Black	Tool	Weapon	6	7	iGPT	0.86	0.07	1.0
							SimCLR	1.38	$< 10^{-2}$	1.0
Weapon (Modern)	White	Black	Tool	Weapon	6	9	iGPT	0.88	0.06	N/A
							SimCLR	1.28	0.01	N/A
Weight [†]	Thin	Fat	Pleasant	Unpleasant	10	55	iGPT	1.67	$< 10^{-3}$	1.83
							SimCLR	-0.30	0.74	1.83

[§] Originally a picture-IAT (image-only stimuli). [†] Originally a mixed-mode IAT (image and verbal stimuli).

7.2 Intersectional Biases

7.2.1 Intersectional Valence. Intersectional valence tests with the iGPT embeddings are the most consistent with social psychology, exhibiting results predicted by the intersectionality, race, and gender hypotheses listed in Section 3 [29]. Overall, iGPT embeddings contain a positive valence bias towards White people and a negative valence bias towards Black people (effect size 1.16, $p < 10^{-3}$), as in the human Race IAT [56]. As predicted by the race hypothesis, the same bias is significant but less severe for both White males vs. Black males (iGPT effect size 0.88, $p < 10^{-2}$) and White males vs. Black females (iGPT effect size 0.83, $p < 10^{-2}$), and the White female vs. Black female bias is insignificant; in general, race biases are more similar to the race biases between men. We hypothesize that as in text corpora, computer vision datasets are dominated by the majority social groups (men and White).

As predicted by the gender hypothesis, our results also conform with the theory that females are associated with positive valence when compared to males [22], but only when those groups are White (iGPT effect size 0.79, $p < 10^{-2}$); there is no significant valence bias for Black females vs. Black males. This insignificant

result might be due to the under-representation of Black people in the visual embedding space. The largest differential valence bias of all our tests emerges between White females and Black males; White females are associated with pleasant valence and Black males with negative valence (iGPT effect size 1.46, $p < 10^{-3}$).

7.2.2 Intersectional Stereotypes. We find significant but contradictory intersectional differences in gender stereotypes (Table 2). For Gender-Career stereotypes, the iGPT-encoded bias for White males vs. Black females is insignificant though there is a bias (effect size 0.81, $p < 10^{-3}$) for male vs. female in general. There is significant Gender-Career stereotype bias between embeddings of White males vs. White females (iGPT effect size 0.97, $p < 10^{-3}$), even higher than the general case; this result conforms to the race hypothesis, which predicts gender stereotypes are more similar to the stereotypes between Whites than between Blacks. The career-family bias between White males and Black males is reversed; embeddings for images of Black males are more associated with career and images of White men with family (iGPT effect size 0.89, $p < 10^{-2}$). One explanation for this result is under-representation; there are likely fewer photos depicting Black men with non-stereotypical male attributes.

Table 2: iEAT tests for the association between intersectional group X vs. Y (represented by n_t images each) and attributes A vs. B (represented by n_a images each) in embeddings produced by an unsupervised model. Effect sizes d represent the magnitude of bias, colored by conventional small (0.2), medium (0.5), and large (0.8). Permutation p -values indicate significance.

	X	Y	A	B	n_t	n_a	d	p
Gender-Career (MF)	Male	Female	Career	Family	40	21	0.81	$< 10^{-3}$
Gender-Career (WMBF)	White Male	Black Female			20	21	0.20	0.27
Gender-Career (WMBM)	Black Male	White Male	Career	Family	20	21	0.89	$< 10^{-2}$
Gender-Career (WMWF)	White Male	White Female			20	21	0.97	$< 10^{-3}$
Gender-Science (MF)	Male	Female	Science	Liberal Arts	40	21	0.00	0.50
Gender-Science (WMBF)	White Male	Black Female			20	21	0.80	$< 10^{-2}$
Gender-Science (WMBM)	White Male	Black Male	Science	Liberal Arts	20	21	0.49	0.06
Gender-Science (WMWF)	White Male	White Female			20	21	-0.37	0.88
Valence (BFBM)	Black Female	Black Male	Pleasant	Unpleasant	20	55	0.17	0.29
Valence (BW)	White	Black			40	55	1.16	$< 10^{-3}$
Valence (FM)	Female	Male	Pleasant	Unpleasant	40	55	0.39	0.04
Valence (WBBF)	White Female	Black Female			20	55	1.51	$< 10^{-3}$
Valence (WBBM)	White Female	Black Male	Pleasant	Unpleasant	20	55	1.46	$< 10^{-3}$
Valence (WMBF)	White Male	Black Female			20	55	0.83	$< 10^{-2}$
Valence (WMBM)	White Male	Black Male	Pleasant	Unpleasant	20	55	0.88	$< 10^{-2}$
Valence (WMWF)	White Female	White Male			20	55	0.79	$< 10^{-2}$

Unexpectedly, the intersectional test of male vs. female (with equal representation for White and Black people) reports no significant Gender-Science bias, though the normative test (with unequal representation) does (Table 1). Nevertheless, race-science stereotypes do emerge when White males are compared to Black males (iGPT effect size 0.49, $p < 10^{-1}$) and, to an even greater extent, when White males are compared to Black females (iGPT effect size 0.80, $p < 10^{-2}$), confirming the intersectional hypothesis [29]. But visual Gender-Science biases do not conform to the race hypothesis; the gender stereotype between White males and White females is insignificant, though the overall male vs. female bias is not.

7.3 Origins of Bias

7.3.1 Bias in Web Images. Do these results correspond with our hypothesis that biases are learned from the co-occurrence of social group members with certain stereotypical or high-valence contexts? Both our models were pre-trained on ImageNet, which is composed of images collected from Flickr and other Internet sites [64]. Yang et al. [78] show that the ImageNet categories unequally represent race and gender; for instance, the “groom” category may contain mostly White people. Under-representation in the training set could explain why, for instance, White people are more associated with pleasantness and Black people with unpleasantness. There is a similar theory in social psychology: most bias takes the form of in-group favoritism, rather than out-group derogation [39]. In image datasets, favoritism could take the form of unequal representation and have similar effects. For example, one of the exemplars for “pleasantness” is “wedding,” a positive-valence, high imagery word [5]; if White people appear with wedding paraphernalia more often than Black people, they could be automatically associated with a concept like “pleasantness,” even though no explicit labels for “groom” and “White” are available during training.

Likewise, the portrayal of different social groups in context may be automatically learned by unsupervised image models. Wang et al. [73] find that in OpenImages (also scraped from Flickr) [46],

a similar benchmark classification dataset, a higher proportion of “female” images are set in the scene “home or hotel” than “male” images. “male” is more often depicted in “industrial and construction” scenes. This difference in portrayal could account for the Gender-Career biases embedded in unsupervised image embeddings. In general, if the portrayal of people in Internet images reflects human social biases that are documented in cognition and language, we conclude that unsupervised image models could automatically learn human-like biases from large collections of online images.

7.3.2 Bias in Autoregression. Though the next-pixel prediction features contained very little significant bias, they may still propagate stereotypes in practice. For example, the incautious and unethical application of a generative model like iGPT could produce biased depictions of people. As a qualitative case study, we selected 5 male- and 5 female-appearing artificial faces from a database [1] generated with StyleGAN [42]. We decided to use images of non-existent people to avoid perpetuating any harm to real individuals. We cropped the portraits below the neck and used iGPT to generate 8 different completions (with the temperature hyperparameter set to 1.0, following Chen et al. [14]). We found that completions of woman *and* men are often sexualized: for female faces, 52.5% of completions featured a bikini or low-cut top; for male faces, 7.5% of completions were shirtless or wore low-cut tops, while 42.5% wore suits or other career-specific attire. One held a gun. This behavior might result from the sexualized portrayal of people, especially women, in internet images [30] and serves as a reminder of computer vision’s controversial history with Playboy centerfolds and objectifying images [41]. To avoid promoting negative biases, Figure 3 shows only an example of male-career associations in completions of a GAN-generated face.

8 DISCUSSION

By testing for bias in unsupervised models pre-trained on a widely used large computer vision dataset, we show how biases may

Figure 3: Example of career associations in image completion of a male face with iGPT, pre-trained on ImageNet.

be learned automatically from images and embedded in general-purpose representations. Not only do we observe human-like biases in the majority of our tests, but we also detect 4 of the 5 human biases replicated in natural language [11]. Caliskan et al. [11] show that artifacts of the societal status quo, such as occupational gender statistics, are imprinted in online text and mimicked by machines. We suggest that a similar phenomenon is occurring for online images. One possible culprit is confirmation bias [65], the tendency of individuals to consume and produce content conforming to group norms. Self-supervised models exhibit the same tendency [2].

In addition to confirming human and natural language machine biases in the image domain, the iEAT measures visual biases that may implicitly affect humans and machines but cannot be captured in text corpora. Foroni and Bel-Bahar [25] conjecture that in humans, picture-IATs and word-IATs measure different mental processes. More research is needed to explore biases embedded in images and investigate their origins, as Brunet et al. [9] suggest for language models. Tenney et al. [70] show that contextual representations learn syntactic and semantic features from the context. Voita et al. [72] explain the change of vector representations among layers based on the compression/prediction trade-off perspective. Advances in this direction would contribute to our understanding of the causal factors behind visual perception and biases related to cognition and language acquisition.

Our methods come with some limitations. The biases we measure are in large part due to patterns learned from the pre-training data, but ImageNet 2012 does not necessarily represent the entire population of images currently produced and circulated on the Internet. Additionally, ImageNet 2012 is intended for object detection, not distinguishing people’s social attributes, and both our models were validated for non-person object classification.⁸ The largest version

⁸Recently, Yang et al. [78] proposed updates to improve fairness and representation in the ImageNet “person” category that could change our results.

of iGPT (not publicly available) was pre-trained on 100 million additional web images [14]. Given the financial and carbon costs of the computation required to train highly parameterized models like iGPT, we did not train our own models on larger-scale corpora. Complementary iEAT bias testing with unsupervised models pre-trained on an updated version of ImageNet could help quantify the effectiveness of dataset de-biasing strategies.

A model like iGPT, pre-trained on a more comprehensive private dataset from a platform like Instagram or Facebook, could encode much more information about contemporary social biases. Clearview AI reportedly scraped over 3 billion images from Facebook, YouTube, and millions of other sites for their face recognition model [40]. Dosovitskiy et al. [21] recently trained a very similar transformer model on Google’s JFT-300M, a 300 million image dataset scraped from the web [66]. Further research is needed to determine how architecture choices affect embedded biases and how dataset filtering and balancing techniques might help [74, 75]. Previous metric-based and adversarial approaches generally require labeled datasets [73–75]. Our method avoids the limitations of laborious manual labeling.

Though models like these may be useful for quantifying contemporary social biases as they are portrayed in vast quantities of images on the Internet, our results suggest the use of unsupervised pre-training on images at scale is likely to propagate harmful biases. Given the high computational and carbon cost of model training at scale, transfer learning with pre-trained models is an attractive option for practitioners. But our results indicate that patterns of stereotypical portrayal of social groups do affect unsupervised models, so careful research and analysis are needed before these models make consequential decisions about individuals and society. Our method can be used to assess task-agnostic biases contained in a dataset to enhance transparency [27, 51], but bias mitigation for unsupervised transfer learning is a challenging open problem.

9 CONCLUSIONS

We develop a principled method for measuring bias in unsupervised image models, adapting embedding association tests used in the language domain. With image embeddings extracted by state-of-the-art unsupervised image models pre-trained on ImageNet, we successfully replicate validated bias tests in the image domain and document several social biases, including severe intersectional bias. Our results suggest that unsupervised image models learn human biases from the way people are portrayed in images on the web. These findings serve as a caution for computer vision practitioners using transfer learning: pre-trained models may embed all types of harmful human biases from the way people are portrayed in training data, and model design choices determine whether and how those biases are propagated into harms downstream.

ACKNOWLEDGMENTS

This material is based on research partially supported by the U.S. National Institute of Standards and Technology (NIST) Grant 60NANB20D212. 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] 2021. Generated Photos. <https://generated.photos>
- [2] Eric Arazo, Diego Ortego, Paul Albert, Noel E. O'Connor, and Kevin McGuinness. 2020. Pseudo-Labeling and Confirmation Bias in Deep Semi-Supervised Learning. In *Proceedings of the International Joint Conference on Neural Networks*. Institute of Electrical and Electronics Engineers Inc., 1–8. <https://doi.org/10.1109/IJCNN48605.2020.9207304>
- [3] Philip Bachman, R Devon Hjelm, and William Buchwalter. 2019. Learning Representations by Maximizing Mutual Information Across Views. In *Advances in Neural Information Processing Systems*, H Wallach, H Larochelle, A Beygelzimer, F d Alché-Buc, E Fox, and R Garnett (Eds.), Vol. 32. Curran Associates, Inc., 15535–15545. <https://proceedings.neurips.cc/paper/2019/file/ddf354219aac374fd40b7e760ee5bb7-Paper.pdf>
- [4] Christine Basta, Marta R Costa-jussà, and Noe Casas. 2019. Evaluating the Underlying Gender Bias in Contextualized Word Embeddings. In *Proceedings of the First Workshop on Gender Bias in Natural Language Processing*. Association for Computational Linguistics, Florence, Italy, 33–39. <https://doi.org/10.18653/v1/W19-3805>
- [5] Francis S. Bellezza, Anthony G. Greenwald, and Mahzarin R. Banaji. 1986. Words high and low in pleasantness as rated by male and female college students. *Behavior Research Methods, Instruments, & Computers* 18, 3 (5 1986), 299–303. <https://doi.org/10.3758/BF03204403>
- [6] Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. Language (Technology) is Power: A Critical Survey of "Bias" in NLP. *arXiv preprint arXiv:2005.14050* (2020).
- [7] Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. 2016. Man is to Computer Programmer as Woman is to Homemaker? Debiasing Word Embeddings. In *Advances in Neural Information Processing Systems* 29, D D Lee, M Sugiyama, U V Luxburg, I Guyon, and R Garnett (Eds.). Curran Associates, Inc., 4349–4357. <http://papers.nips.cc/paper/6228-man-is-to-computer-programmer-as-woman-is-to-homemaker-debiasing-word-embeddings.pdf>
- [8] Rishi Bommasani, Kelly Davis, and Claire Cardie. 2020. Interpreting Pretrained Contextualized Representations via Reductions to Static Embeddings. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics (ACL), 4758–4781. <https://doi.org/10.18653/v1/2020.acl-main.431>
- [9] Marc-Etienne Brunet, Colleen Alkalay-Houlihan, Ashton Anderson, and Richard Zemel. 2019. Understanding the Origins of Bias in Word Embeddings. In *Proceedings of the 36th International Conference on Machine Learning*. PMLR, 803–811. <http://proceedings.mlr.press/v97/brunet19a.html>
- [10] Joy Buolamwini and Timnit Gebru. 2018. Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency (Proceedings of Machine Learning Research, Vol. 81)*, Sorelle A Friedler and Christo Wilson (Eds.). PMLR, New York, NY, USA, 77–91. <http://proceedings.mlr.press/v81/buolamwini18a.html>
- [11] Aylin Caliskan, Joanna J Bryson, and Arvind Narayanan. 2017. *Semantics Derived Automatically from Language Corpora Contain Human-like Biases*. Technical Report 6334. Science. 183–186 pages. <https://doi.org/10.1126/science.aal4230>
- [12] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. 2020. End-to-End Object Detection with Transformers. *arXiv preprint arXiv:2005.12872* (2020).
- [13] George H Chen. 2020. Deep Kernel Survival Analysis and Subject-Specific Survival Time Prediction Intervals. In *Proceedings of the 5th Machine Learning for Healthcare Conference (Proceedings of Machine Learning Research, Vol. 126)*, Finale Doshi-Velez, Jim Fackler, Ken Jung, David Kale, Rajesh Ranganath, Byron Wallace, and Jenna Wiens (Eds.). PMLR, Virtual, 537–565. <http://proceedings.mlr.press/v126/chen20a.html>
- [14] Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. 2020. Generative Pretraining From Pixels. In *Proceedings of the 37th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 119)*, Hal Daumé III and Aarti Singh (Eds.). PMLR, 1691–1703. <http://proceedings.mlr.press/v119/chen20s.html>
- [15] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A Simple Framework for Contrastive Learning of Visual Representations. In *Proceedings of the 37th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 119)*, Hal Daumé III and Aarti Singh (Eds.). PMLR, 1597–1607. <http://proceedings.mlr.press/v119/chen20j.html>
- [16] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey Hinton. 2020. Big self-supervised models are strong semi-supervised learners. *arXiv preprint arXiv:2006.10029* (2020).
- [17] Kimberle Crenshaw. 1990. Mapping the margins: Intersectionality, identity politics, and violence against women of color. *Stan. L. Rev.* 43 (1990), 1241.
- [18] Maria De-Arteaga, Alexey Romanov, Hanna Wallach, Jennifer Chayes, Christian Borgs, Alexandra Chouldechova, Sahin Geyik, Krishnamurthy Kenthapadi, and Adam Tauman Kalai. 2019. Bias in bios: A case study of semantic representation bias in a high-stakes setting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. ACM, 120–128.
- [19] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [20] Jeff Donahue and Karen Simonyan. 2019. Large Scale Adversarial Representation Learning. In *Advances in Neural Information Processing Systems*, H Wallach, H Larochelle, A Beygelzimer, F d Alché-Buc, E Fox, and R Garnett (Eds.), Vol. 32. Curran Associates, Inc., 10542–10552. <https://proceedings.neurips.cc/paper/2019/file/18cdf49ea54e0c29238fcc95f76ce41-Paper.pdf>
- [21] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xi-aohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=YicbFdNTTy>
- [22] Alice H Eagly, Antonio Mladinic, and Stacey Otto. 1991. Are women evaluated more favorably than men?: An analysis of attitudes, beliefs, and emotions. *Psychology of Women Quarterly* 15, 2 (1991), 203–216.
- [23] Dumitru Erhan, Yoshua Bengio, Aaron Courville, Pierre-Antoine Manzagol, Pascal Vincent, and Samy Bengio. 2010. Why Does Unsupervised Pre-training Help Deep Learning? *Journal of Machine Learning Research* 11, 19 (2010), 625–660. <http://jmlr.org/papers/v11/erhan10a.html>
- [24] Dumitru Erhan, Pierre-Antoine Manzagol, Yoshua Bengio, Samy Bengio, and Pascal Vincent. 2009. The Difficulty of Training Deep Architectures and the Effect of Unsupervised Pre-Training. In *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics (Proceedings of Machine Learning Research, Vol. 5)*, David van Dyk and Max Welling (Eds.). PMLR, Hilton Clearwater Beach Resort, Clearwater Beach, Florida USA, 153–160. <http://proceedings.mlr.press/v5/erhan09a.html>
- [25] Francesco Foroni and Tarik Bel-Bahar. 2010. Picture-IAT versus word-IAT: Level of stimulus representation influences on the IAT. *European Journal of Social Psychology* 40, 2 (3 2010), 321–337. <https://doi.org/10.1002/ejsp.626>
- [26] Nikhil Garg, Londa Schiebinger, Dan Jurafsky, and James Zou. 2018. Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences of the United States of America* 115, 16 (4 2018), E3635–E3644. <https://doi.org/10.1073/pnas.1720347115>
- [27] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. 2018. Datasheets for datasets. *arXiv preprint arXiv:1803.09010* (2018).
- [28] Andreas Geiger, Philip Lenz, and Raquel Urtasun. 2012. Are we ready for autonomous driving? the KITTI vision benchmark suite. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. IEEE, 3354–3361. <https://doi.org/10.1109/CVPR.2012.6248074>
- [29] 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.
- [30] Kaitlin A Graff, Sarah K Murnen, and Anna K Krause. 2013. Low-cut shirts and high-heeled shoes: Increased sexualization across time in magazine depictions of girls. *Sex roles* 69, 11–12 (2013), 571–582.
- [31] A G Greenwald, D E McGhee, and J L Schwartz. 1998. Measuring Individual Differences in Implicit Cognition: The Implicit Association Test. *Journal of Personality and Social Psychology* 74, 6 (6 1998), 1464–80. <http://www.ncbi.nlm.nih.gov/pubmed/9654756>
- [32] Anthony G. Greenwald, Brian A. Nosek, and Mahzarin R. Banaji. 2003. Understanding and Using the Implicit Association Test: I. An Improved Scoring Algorithm. *Journal of Personality and Social Psychology* 85, 2 (8 2003), 197–216. <https://doi.org/10.1037/0022-3514.85.2.197>
- [33] Anthony G Greenwald, T Andrew Poehlman, Eric Luis Uhlmann, and Mahzarin R Banaji. 2009. Understanding and using the Implicit Association Test: III. Meta-analysis of predictive validity. *Journal of personality and social psychology* 97, 1 (2009), 17.
- [34] Wei Guo and Aylin Caliskan. 2020. Detecting Emergent Intersectional Biases: Contextualized Word Embeddings Contain a Distribution of Human-like Biases. *arXiv preprint arXiv:2006.03955* (2020).
- [35] Drew Harwell. 2019. A face-scanning algorithm increasingly decides whether you deserve the job. <https://www.washingtonpost.com/technology/2019/10/22/ai-hiring-face-scanning-algorithm-increasingly-decides-whether-you-deserve-job/>
- [36] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, 9729–9738.
- [37] 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 (CVPR)*. IEEE, 770–778. <http://image-net.org/challenges/LSVRC/2015/>
- [38] Lisa Anne Hendricks, Kaylee Burns, Kate Saenko, Trevor Darrell, and Anna Rohrbach. 2018. Women also snowboard: Overcoming bias in captioning models. In *European Conference on Computer Vision*. 793–811.

- [39] Miles Hewstone, Mark Rubin, and Hazel Willis. 2002. Intergroup bias. *Annual review of psychology* 53, 1 (2002), 575–604.
- [40] Kashmir Hill. 2020. The Secretive Company That Might End Privacy as We Know It. <https://www.nytimes.com/2020/01/18/technology/clearview-privacy-facial-recognition.html>
- [41] Corinne Iozzio. 2016. The Playboy Centerfold That Revolutionized Image-Processing Research. *The Atlantic* (2 2016). <https://www.theatlantic.com/technology/archive/2016/02/lena-image-processing-playboy/461970/>
- [42] Tero Karras, Samuli Laine, and Timo Aila. 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.
- [43] Matthew Kay, Cynthia Matuszek, and Sean A. Munson. 2015. Unequal Representation and Gender Stereotypes in Image Search Results for Occupations. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems - CHI '15*. ACM Press, New York, New York, USA, 3819–3828. <https://doi.org/10.1145/2702123.2702520>
- [44] Alex Krizhevsky. 2009. *Learning multiple layers of features from tiny images*. Technical Report. University of Toronto. <https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf>
- [45] Keita Kurita, Nidhi Vyas, Ayush Pareek, Alan W Black, and Yulia Tsvetkov. 2019. Measuring Bias in Contextualized Word Representations. In *Proceedings of the First Workshop on Gender Bias in Natural Language Processing*. Association for Computational Linguistics (ACL), Florence, Italy, 166–172. <https://doi.org/10.18653/v1/w19-3823>
- [46] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Tom Duerig, and others. 2018. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *arXiv preprint arXiv:1811.00982* (2018).
- [47] Varun Manjunatha, Nirat Saini, and Larry S Davis. 2019. Explicit Bias Discovery in Visual Question Answering Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE.
- [48] Chandler May, Alex Wang, Shikha Bordia, Samuel R. Bowman, and Rachel Rudinger. 2019. On Measuring Social Biases in Sentence Encoders. In *Proceedings of the 2019 Conference of the North American Association for Computational Linguistics*, Stroudsburg, PA, USA, 622–628. <https://doi.org/10.18653/v1/N19-1063>
- [49] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781* (2013).
- [50] Ishan Misra and Laurens Van Der Maaten. 2020. Self-Supervised Learning of Pretext-Invariant Representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 6707–6717.
- [51] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. Model Cards for Model Reporting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT* '19)*. Association for Computing Machinery, New York, NY, USA, 220–229. <https://doi.org/10.1145/3287560.3287596>
- [52] Francesco Nex and Fabio Remondino. 2014. UAV for 3D mapping applications: A review. , 15 pages. <https://doi.org/10.1007/s12518-013-0120-x>
- [53] Brian A. Nosek and Mahzarin R. Banaji. 2001. The GO/NO-GO Association Task. *Social Cognition* 19, 6 (12 2001), 625–664. <https://doi.org/10.1521/soco.19.6.625.20886>
- [54] Brian A. Nosek, Mahzarin R. Banaji, and Anthony G. Greenwald. 2002. Harvesting implicit group attitudes and beliefs from a demonstration web site. *Group Dynamics* 6, 1 (2002), 101–115. <https://doi.org/10.1037/1089-2699.6.1.101>
- [55] Brian A Nosek, Anthony G Greenwald, and Mahzarin R Banaji. 2007. The Implicit Association Test at Age 7: A Methodological and Conceptual Review. In *Automatic processes in social thinking and behavior*, J. A. Bargh (Ed.). Psychology Press, Chapter 6, 265–292.
- [56] Brian A. Nosek, Frederick L. Smyth, Jeffrey J. Hansen, Thierry Devos, Nicole M. Lindner, Kate A. Ranganath, Colin Tucker Smith, Kristina R. Olson, Dolly Chugh, Anthony G. Greenwald, and Mahzarin R. Banaji. 2007. Pervasiveness and correlates of implicit attitudes and stereotypes. *European Review of Social Psychology* 18, 1 (11 2007), 36–88. <https://doi.org/10.1080/10463280701489053>
- [57] Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. GloVe: Global Vectors for Word Representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, Doha, Qatar, 1532–1543. <https://doi.org/10.3115/v1/D14-1162>
- [58] Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep Contextualized Word Representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. Association for Computational Linguistics, New Orleans, Louisiana, 2227–2237. <https://doi.org/10.18653/v1/N18-1202>
- [59] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog* 1, 8 (2019), 9.
- [60] Manish Raghavan, Solon Barocas, Jon Kleinberg, and Karen Levy. 2020. Mitigating bias in algorithmic hiring: Evaluating claims and practices. In *FAT* 2020 - Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. Association for Computing Machinery, Inc, 469–481. <https://doi.org/10.1145/3351095.3372828>
- [61] Inioluwa Deborah Raji, Timnit Gebru, Margaret Mitchell, Joy Buolamwini, Joonseok Lee, and Emily Denton. 2020. Saving face: Investigating the ethical concerns of facial recognition auditing. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*. ACM, 145–151.
- [62] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. 2019. Do ImageNet Classifiers Generalize to ImageNet?. In *Proceedings of the 36th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 97)*, Kamalika Chaudhuri and Ruslan Salakhutdinov (Eds.). PMLR, 5389–5400. <http://proceedings.mlr.press/v97/recht19a.html>
- [63] Olga Russakovsky, Jia Deng, Zhiheng Huang, Alexander C Berg, and Li Fei-Fei. 2013. Detecting avocados to zucchinis: what have we done, and where are we going?. In *Proceedings of the IEEE International Conference on Computer Vision*. IEEE, 2064–2071.
- [64] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision* 115, 3 (12 2015), 211–252. <https://doi.org/10.1007/s11263-015-0816-y>
- [65] Stefan Schweiger, Aileen Oeberst, and Ulrike Cress. 2014. Confirmation bias in web-based search: A randomized online study on the effects of expert information and social tags on information search and evaluation. *Journal of Medical Internet Research* 16, 3 (2014), –undefined. <https://doi.org/10.2196/jmir.3044>
- [66] Chen Sun, Abhinav Shrivastava, Saurabh Singh, and Abhinav Gupta. 2017. Revisiting unreasonable effectiveness of data in deep learning era. In *Proceedings of the IEEE international conference on computer vision*. IEEE, 843–852.
- [67] L Sweeney. 1997. Weaving technology and policy together to maintain confidentiality (vol 25, pg 2, 1997). *Journal Of Law Medicine & Ethics* 25, 4 (1997), 327.
- [68] Chuanqi Tan, Fuchun Sun, Tao Kong, Wenchang Zhang, Chao Yang, and Chunfang Liu. 2018. A survey on deep transfer learning. In *International conference on artificial neural networks*. IEEE, 270–279.
- [69] Yi Chern Tan and L Elisa Celis. 2019. Assessing Social and Intersectional Biases in Contextualized Word Representations. In *Advances in Neural Information Processing Systems, H Wallach, H Larochelle, A Beygelzimer, F d Alché-Buc, E Fox, and R Garnett (Eds.)*, Vol. 32. Curran Associates, Inc., 13230–13241. <https://proceedings.neurips.cc/paper/2019/file/201d546992726352471cfa6b0df0a48-Paper.pdf>
- [70] Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R Thomas McCoy, Najuoung Kim, Benjamin Van Durme, Samuel R Bowman, Dipanjan Das, and others. 2019. What do you learn from context? probing for sentence structure in contextualized word representations. *arXiv preprint arXiv:1905.06316* (2019).
- [71] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin. 2017. Attention Is All You Need. In *Advances in Neural Information Processing Systems*, I Guyon, U V Luxburg, S Bengio, H Wallach, R Fergus, S Vishwanathan, and R Garnett (Eds.), Vol. 30. Curran Associates, Inc., 5998–6008. <https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf>
- [72] Elena Voita, Rico Sennrich, and Ivan Titov. 2019. The bottom-up evolution of representations in the transformer: A study with machine translation and language modeling objectives. *arXiv preprint arXiv:1909.01380* (2019).
- [73] A Wang, A Narayanan, and O Russakovsky. 2020. REVISE: A Tool for Measuring and Mitigating Bias in Visual Datasets. In *European Conference on Computer Vision*.
- [74] 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 International Conference on Computer Vision*. IEEE, 5310–5319.
- [75] Zeyu Wang, Klint Qinami, Ioannis Christos Karakozis, Kyle Genova, Prem Nair, Kenji Hata, and Olga Russakovsky. 2020. Towards fairness in visual recognition: Effective strategies for bias mitigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, 8919–8928.
- [76] Benjamin Wilson, Judy Hoffman, and Jamie Morgenstern. 2019. Predictive Inequity in Object Detection. *arXiv preprint arXiv:1902.11097* (2 2019). <http://arxiv.org/abs/1902.11097>
- [77] Kaiyuan Xu, Brian Nosek, and Anthony Greenwald. 2014. Data from the Race Implicit Association Test on the Project Implicit Demo Website. *Journal of Open Psychology Data* 2, 1 (3 2014), e3. <https://doi.org/10.5334/jopd.ac>
- [78] Kaiyu Yang, Klint Qinami, Li Fei-Fei, Jia Deng, and Olga Russakovsky. 2020. Towards Fairer Datasets: Filtering and Balancing the Distribution of the People Subtree in the ImageNet Hierarchy. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAT* '20)*. Association for Computing Machinery, New York, NY, USA, 547–558. <https://doi.org/10.1145/3351095.3375709>

- [79] Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai Wei Chang. 2017. Men also like shopping: Reducing gender bias amplification using corpus-level constraints. In *EMNLP 2017 - Conference on Empirical Methods in Natural Language Processing, Proceedings*. Association for Computational Linguistics (ACL), 2979–2989. <https://doi.org/10.18653/v1/d17-1323>