

# Gender Artifacts in Visual Datasets

Nicole Meister<sup>1†\*</sup>, Dora Zhao<sup>2†\*</sup>, Angelina Wang<sup>3</sup>, Vikram V. Ramaswamy<sup>3</sup>,  
Ruth Fong<sup>3</sup>, Olga Russakovsky<sup>3</sup>

<sup>1</sup>Stanford University <sup>2</sup>Sony AI <sup>3</sup>Princeton University

\*Equal contribution †Work done as a student at Princeton University

{nmeister, dorothyz}@stanford.edu

## Abstract

Gender biases are known to exist within large-scale visual datasets and can be reflected or even amplified in downstream models. Many prior works have proposed methods for mitigating gender biases, often by attempting to remove gender expression information from images. To understand the feasibility and practicality of these approaches, we investigate what “gender artifacts” exist in large-scale visual datasets. We define a “gender artifact” as a visual cue correlated with gender, focusing specifically on cues that are learnable by a modern image classifier and have an interpretable human corollary. Through our analyses, we find that gender artifacts are ubiquitous in the COCO and OpenImages datasets, occurring everywhere from low-level information (e.g., the mean value of the color channels) to higher-level image composition (e.g., pose and location of people). Further, bias mitigation methods that attempt to remove gender actually remove more information from the scene than the person. Given the prevalence of gender artifacts, we claim that attempts to remove these artifacts from such datasets are largely infeasible as certain removed artifacts may be necessary for the downstream task of object recognition. Instead, the responsibility lies with researchers and practitioners to be aware that the distribution of images within datasets is highly gendered and hence develop fairness-aware methods which are robust to these distributional shifts across groups.

## 1. Introduction

It has been well established that machine learning systems contain gender biases [12, 15, 31, 76, 82, 83]. For instance, image tagging systems have labelled similarly posed female politicians differently than their male counterparts [64], and image search engines can return stereotypical results mirroring harmful gender roles [38, 48, 50].

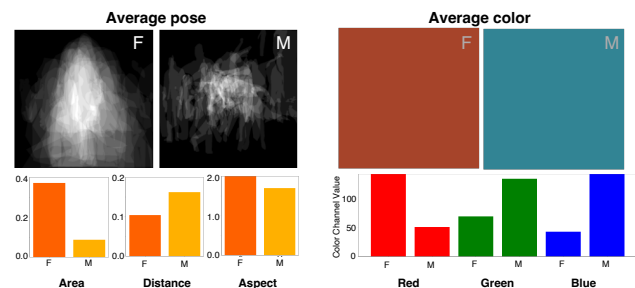


Figure 1. We aim to understand gender artifacts in image datasets by training a classifier to distinguish between images containing a person labelled “female” vs. “male”. We find that such a classifier frequently performs significantly above random chance even when trained and evaluated on modified versions of images (e.g. where a person’s *appearance* is obscured). On the left, we trained a CNN on COCO [42] images with only the person *segmentation mask* visible; on the right, we trained a logistic regression classifier using only 3 features (i.e. *average color*) per COCO image. Both classifiers perform significantly above random chance, illustrating how deeply gender artifacts are embedded within the dataset. We visualize the classifiers’ top 20 (most “female”) and bottom 20 (most “male”) images (top), and further show the differences between their segmentation masks (left bottom) in normalized area, distance from the center, and aspect ratio, and the differences between the R, G and B features (right bottom). These results suggest that gender artifacts are pervasive in computer vision datasets, and current “fairness through blindness” methods that aim to “remove” gendered information may not be as effective as hoped.

A prevailing assumption in the fairness community is that biases in models originate in biases in the input data. Thus, there has been an effort to mitigate these dataset biases (and thus also model biases). While many biases stem from the lack of representation of certain demographic groups, including women [13, 82, 83], naively balancing gender distributions has been shown to be insufficient for mitigating model biases [76]. As a result, a large vein of work has attempted to pursue “fairness through blindness” [22] which operates under the assumption that remov-

ing gendered information can remove the potential for gender bias in the dataset. These attempts include using segmentation masks to occlude the person in images [9, 31, 70] or removing image features correlated with gender [75, 76].

In our work, we focus not on mitigating these dataset biases, but rather on exploring to what extent gendered information can truly be removed from the dataset. To do so, we develop a framework to identify *gender artifacts*, or visual cues correlated with gender labels. We use the term “artifact” to emphasize how these correlations may not be related to gender presentation, rather simply a result of the dataset collection process. While there can be infinitely many potential correlations in an image, we focus on those that are learnable (i.e., result in a learnable difference in a model) and interpretable (i.e., have an interpretable human corollary, such as color and shape).<sup>1</sup> Overall, we do not claim (nor believe) the existence of a specific gender artifact necessarily leads to disparate performance on downstream tasks for different groups (a common measure of fairness [13, 63, 80]), however, the *source* of disparate performance lies in gender artifacts which are critical to study.

To discover gender artifacts, we use a gender classifier, which we refer to as a “gender artifact model,” since the model is not predicting gender so much as it is predicting correlations to gender expression in the training dataset. Our goal is to understand what the model uses as predictive features and how variations of the dataset with particular gender artifacts removed may affect its performance. Thus, we systematically manipulate the datasets (e.g., occlude the person, occlude the background, etc.) and train our model on this data to identify potential gender artifacts. This framework for discovering gender artifacts is more complex than prior work [74], as it goes beyond analyzing annotated attributes to include higher level perceptual components of the image (e.g., resolution, color) and person and background artifacts to gain a deeper understanding of where gender artifacts exist in visual datasets.

It is important to note we do not condone the use of automated gender recognition in practice. We perform gender classification only to understand differences in data distributions, *not* for the sake of classifying gender itself. We never use the model outputs as the end goal and no causal claims about the observed correlation between an attribute and gender labels can be made. Instead, the outputs are a means for better understanding the data. Furthermore, our use of the term “gender” when referencing people in image datasets refers to the binarized perceived gender expression. Since most image datasets do not include self-reported gender identity, we rely on external annotators’ perception as gender labels [63]. To be consistent with prior works [13, 31, 76, 82, 83], we use “male” and “female” to refer to gender expression. Finally, we do not make any nor-

mative claims about the gender artifacts (i.e., if it is good or bad); we focus simply on identifying which artifacts exist.

Using our proposed framework, we perform an in-depth analysis into understanding the wide variety of gender artifacts present in two popular image datasets: Common Objects in Context (COCO) [42] and OpenImages [40]. Our results show that gender artifacts are *everywhere*, from the shape of the person segmentation mask (Fig. 1) to randomly selected contextual objects. Even the *average color* of the image is sufficient to distinguish between the genders: when each image is represented by just three features, the mean pixel value of the red, green, and blue color channels, the classifier is able to achieve an AUC of  $58.0\% \pm 0.4$  in COCO and  $59.1\% \pm 0.4$  in OpenImages.

These findings have the following implications:

- Many prior works have proposed mitigation methods that attempt to remove gender expression information from the image [1, 9, 31, 70, 76]. In contrast, we show **gender artifacts are so intricately embedded in visual datasets that attempts to remove them via mitigation techniques may be a futile endeavor**. This is evidenced by experiments where even after a person is *entirely occluded* with a rectangular mask and the background is visible, the gender artifact model is nevertheless able to reliably distinguish which of the two genders is present in the image, achieving an AUC of 70.8% in COCO [42] and 63.0% in OpenImages [40]. Using a popular adversarial approach [76] as a case study, we find 76.8% of the removed information comes from the scene rather than the person.
- We more realistically advise practitioners to **adopt fairness-aware models** [22, 80] and disaggregated evaluation metrics [5], which explicitly account for potential discrepancies between protected groups.
- Since there are so many salient artifacts in image datasets that are (spuriously) correlated with gender, our findings point to an **incoherence of gender prediction**. Any time a model predicts “gender,” we should wonder what the prediction refers to; it could easily be relying on the background scene color, rather than societally meaningful forms of gender expression.

## 2. Related Work

**Understanding model biases.** Understanding where biases arise in models is important to inform mitigation strategies. One proposed method is experimentally manipulating features in an image to isolate sources of gender biases. This has been done using simple image processing techniques, such as changing the image brightness [47], manually finding counterfactual examples [69], or using generative adversarial networks (GANs) to synthetically manipulate at-

<sup>1</sup>See Appendix A for details on artifact constraints.

tributes [3, 21, 55, 67]. Another approach is to use interpretability methods, such as attention heatmaps, to understand where the model focuses when making a prediction [31]. While we use classifiers as a tool for understanding gender artifacts, we focus on how these gender artifacts manifest in *datasets* rather than in the models.

**Identifying dataset biases.** The presence of bias in datasets has been well-studied. While dataset bias can be analyzed along many axes both with respect to demographic [13, 51, 74, 82] and non-demographic attributes [19, 72, 74], we focus on the presence of social biases in visual datasets. Image datasets have been found to include imbalanced demographic representation [13, 79, 82], stereotypical portrayals [15, 64, 73], and even harmful or offensive content [10, 82]. Looking more closely into gender artifacts, prior works have considered differences with respect to demographic attributes [10, 13, 82] and contextual objects [10, 74, 83] (e.g., instance counts, distance from objects).

**Mitigating dataset and model biases.** Prior works have proposed interventions at both the dataset and algorithmic level. Proposed techniques for mitigating dataset bias include manual data cleanup [79] and applications of synthetic methods, such as GANs [55, 60, 67]. Algorithmic interventions have also been proposed to mitigate gender biases. One common approach is to remove gendered information from the image by blurring or occluding pixels corresponding to people [1, 9, 31, 70, 76]. For example, Hendricks et al. [31] developed techniques encouraging the model to be “confused” when predicting gendered words if the person is occluded, and Wang et al. [76] use adversarial methods that learn to obscure image features that correspond with gender. Other examples of algorithmic mitigation techniques include domain independent training [77] and corpus-level constraints [83]. In contrast, we focus on analyzing where gender artifacts may *originate* in visual datasets. We evaluate what these interventions that remove gender artifacts actually address, and in turn, which artifacts may still remain even after applying mitigation techniques.

**Analysis of data collection practices.** Analyzing dataset collection practices is key to better understanding where and how dataset biases arise; see e.g., Paullada et al. [52] for an overview of how collection, annotation, and documentation practices introduce biases into visual datasets. For example, the common practice of scraping images from the internet injects biases inherently present in image search engines, such as gender [38, 48] or geographic biases [19, 66]. In addition, prior work [49] has identified that crowdsourced ground truth labels can be a potential source of bias as annotators may systematically differ in their perceptions, often influenced by their demographic background. In response, fairer dataset collection practices have been suggested, e.g.,

drawing inspiration from other disciplines [35, 36]. One focus has been on improving transparency, including creating public datasheets and tools to guide researcher intervention [10, 26, 32]. There has also been an effort to create more structured processes, including creating checklists and advocating for more institutional oversight [45, 53].

## 3. Setup

### 3.1. Datasets

We focus our analysis on COCO [42] and OpenImages [40] as they are widely used across different computer vision tasks [24, 33, 46], making it particularly important to investigate biases present in these datasets [18]. COCO has been the testbed for previous work on bias mitigation [31, 76, 83] making it central to our analyses. We use OpenImages to validate that our findings on COCO are generalizable. These datasets are also unique in that we are able to derive labels from existing annotations for perceived gender expression to train a gender artifact model.

For COCO, following Zhao et al. [83], we derive gender labels using the COCO 2014 captions as this practice is standard in existing bias mitigation methods [31, 76, 83]. For OpenImages, we use the More Inclusive Annotations for People (MIAP) [63] and exclude images with more than one person of different gender labels. Notably both datasets are skewed male: 69.2% COCO and 62.9% OpenImages.

### 3.2. Model

Our gender artifact model is a ResNet-50 [30] pre-trained on ImageNet [20] (see Appendix C for details). To remove gender artifacts, we train and evaluate on selectively modified versions of images. For a threshold agnostic performance measure, we report area under the ROC curve (AUC) and standard error on the test set. An AUC of 50% suggests the classifier cannot distinguish between male and female gender labels (i.e., there are no image artifacts that allow the classifier to distinguish between these labels).

To note, AUC is not a fairness metric in the way equalized odds [29], equality of opportunity [29], or statistical/prediction parity [6] are. A higher or lower AUC does not indicate that a dataset is more or less fair. Rather, we use AUC to help understand where correlations with gender arise in visual datasets. We follow similar approaches to research within the computational biology community that finds that gender, race, and age can be inferred from various medical imaging techniques (e.g., chest and hand X-rays, mammograms, retinal fundus photography) due to subtle cues in the input space [27, 34, 54]. They similarly predict an attribute from medical images and report AUC.

### 3.3. Ethical considerations

While we only use gender expression classification as a discovery mechanism for visual artifacts, we acknowledge that our setup, which predicts “male” or “female”, reifies the notion of binary gender and grants it legitimacy. Gender classifiers are fundamentally imperfect as gender is reduced to a simplistic binary that can be harmful to individuals from the trans and/or non-binary community who may not fit into these narrow categories [28, 39]. While it is important to understand where gender biases may arise in automated systems, we do not condone the use of automated gender recognition in practice. The purpose of a gender classifier in this project is solely for the study of gender bias propagation. We are aware that by simply calling a gender classifier a “gender artifact model” does not change what it fundamentally is — however, by naming it as such we underscore its purpose as a method of studying gender artifacts *only* and not for its prediction output of gender.

### 4. Resolution and Color

To start, we analyze the higher-level perceptual components of the image that may serve as gender artifacts. We specifically investigate the resolution and color of the images. The scenarios we experiment with are images (originally 224x224) downsampled to 112x112, 56x56, 28x28, 14x14, 7x7. As the images decrease in resolution, it becomes harder to distinguish the contents, such that color becomes one of the more salient artifacts. Thus, we also consider these images of lower resolution in grayscale, to better understand what the model is relying on when the image has been distorted such that individual objects and people are not perceptible. All results are in Fig. 2.

The experiments in manipulating image resolution are motivated by previous work that finds a spatial resolution of 32x32 is sufficient to identify the semantic category of real world scenes [71]. We downsample the images to analyze what happens to a gender artifact model when the input images are beyond a discernible image resolution. The model’s AUC plateaus at around 61.9% for color and 51.9% (near-random) for grayscale COCO images at 28x28 resolution. These results suggest that the shapes in an image are no longer meaningful at that resolution, and the gender artifact model must be predominantly relying on color features.

It is also surprising to see that a gender artifact model trained and evaluated on 7x7 images (a resolution at which almost all image artifacts except color are removed from an image) still achieves an AUC of above random chance (59.2% for COCO and 63.9% for OpenImages).

Exploring this further, we train a logistic regression model<sup>2</sup> on three features: the mean pixel value of the red, green, and blue color channels. Surprisingly, this simple

<sup>2</sup>This model is trained with L2 regularization,  $\lambda = 1$ .

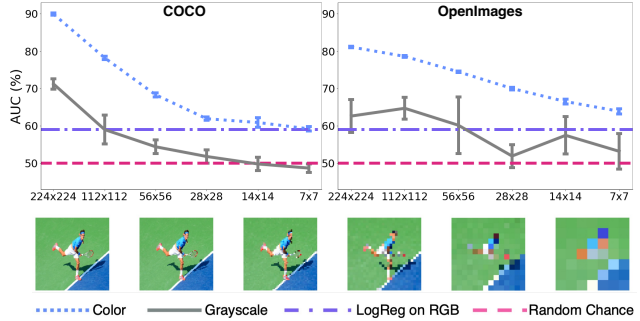


Figure 2. AUC of the gender artifact models on the COCO [42] (top left) and OpenImages [40] (top right) datasets. The models trained and evaluated on color images (blue curves) perform significantly above random chance (pink lines) even as the image resolution gets as low as 7x7. The models trained on grayscale images (gray curves) approach random chance around 28x28 resolution. A sample image at different resolutions is shown on the bottom. Most interestingly, as also highlighted in Fig. 1, the logistic regression classifier trained on just the mean red, green and blue channel values (purple line) performs above random chance on both datasets, underscoring the role of color as a gender artifact and the overall ubiquity of gender artifacts in image datasets.

model achieves an AUC of  $58.0\% \pm 0.4$  on COCO and  $59.1\% \pm 0.4$  on OpenImages. This likely occurs as many images of people labeled male are in grassy sports fields, whereas many images of people labeled female are of the person, generally of lighter skin tone [82], up close. The fact that just three color values are sufficient for a model to differentiate between images of two different genders, emphasizes how ubiquitous gender artifacts are in the dataset.

### 5. Person and Background

Next, we turn to disentangling gender artifacts arising from the person versus the background (i.e., non-person parts of the image). Prior bias mitigation works [31, 76] assume that background information should not contain gender artifacts. For example, Hendricks et al. [31] propose a bias mitigation strategy for image captioning that encourages the model to be “confused” when predicting gendered words if the person is occluded. Similarly, adversarial methods [76] learn to obscure person information from the input image. However, additional work has shown models can achieve non-trivial accuracy in image classification by relying on the background alone [78]. A natural question is whether other aspects of the image (e.g., background, person pose, person size) contain gender artifacts as well.

**Experimental setup.** We use a series of occlusions on COCO and OpenImages. For simplicity, our experiments adhere to the nomenclature shown at the bottom of Fig. 3. We train and evaluate the gender artifact model from Sec. 3 on the manipulated images. All models are trained and



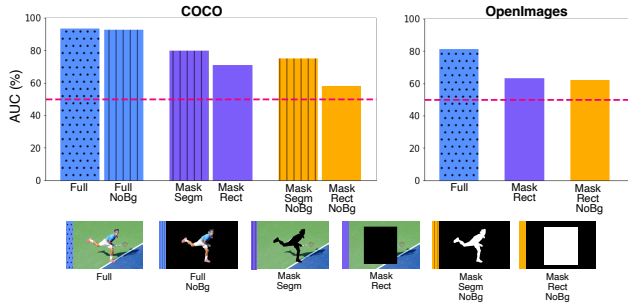


Figure 3. AUC of gender artifact models (top) trained and evaluated on images with different levels of occlusions (shown on bottom). The naming refers to what from the person is shown in an image: all person pixels (Full), segmentation mask (MaskSegm), or bounding box (MaskRect). “NoBg” denotes a masked background. The pink line represents random chance.

tested on the same occlusions (e.g., model trained on images with occluded background are tested on images with occluded background). Additionally, since COCO and OpenImages are skewed by gender, we consider the effect of the imbalanced datasets, replicate these experiments on balanced dataset, and find similar results (see Appendix F).

**Person’s appearance is a gender artifact.** We start by examining the artifacts arising from the person pixels. First, we consider the model trained on Full (i.e., no occlusions) which achieves an AUC of  $93.4\% \pm 0.2$  on COCO. We then compare the performance on Full with on Full NoBg (i.e., only the outline of the person is visible against a black background).<sup>3</sup> The model’s performance does not change considerably even after the background is removed, achieving an AUC of  $92.7\% \pm 0.2$  (see blue bars in Fig. 3).

**A gender artifact model performs above random chance when trained on the person’s shape and location.** Subsequently, we question *what* about the person, namely in their shape and location, is a gender artifact. To isolate artifacts arising from the person’s shape, we occlude the background and then gradually remove information about the person’s appearance (see orange bars in Fig. 3). First, looking at MaskSegm NoBg (i.e., person occluded with white pixels against the black background), we observe that the model performs considerably better than random chance with an AUC of  $74.8\% \pm 0.3$  for COCO. This result may be expected as the segmentation shapes themselves still reveal information about appearance, such as clothing or hair length.

Next, removing even more information about the person, we train our model using only the 17 person keypoint locations from COCO. The model continues to perform above random chance (AUC of  $64.8\% \pm 0.4$ ). Further, to disentangle the role of pose versus other image factors (e.g., size, location), we resize the keypoints to have an area of 4,000

<sup>3</sup>Results when occluding with white pixels are similar; in Appendix E.

pixels and center the keypoints in the image. These images are then used to train a multi-layer perceptron (MLP) for 100 iterations, achieving an AUC of  $58.0\% \pm 0.0$ , suggesting pose is a salient gender artifact. Finally, we inspect the poses with the greatest absolute scores and find the poses predicted to be more likely male are smaller and in action (e.g., playing a sport, jumping) whereas those predicted to be more likely female tended to be larger and standing still.

**There is a learnable difference between the size and location of people of different gender expressions.** Since we observe a difference in the size for poses, we further question whether the model can use only the size and location of people as gender artifacts. We use MaskRect NoBg, (i.e., only the bounding box around the person occluded with white pixels against a black background). Although performance decreases from when pose information was provided, the model continues to find a learnable difference with an AUC of  $58.0\% \pm 0.4$  on COCO. This pattern is also reflected in OpenImages ( $62.2\% \pm 0.4$ ). In fact, in the COCO training set, the size of people labeled as female (normalized by the image area) is  $0.18 \pm 10^{-4}$  and  $0.13 \pm 10^{-5}$  for people labeled as male. This difference is statistically significant with  $p < 0.01$  level. The same pattern holds for OpenImages.

We also verify whether location is predictive. Using a one-hot encoding of the  $32 \times 32$  image where each pixel is set to 0 except the center of the person bounding box, we train an MLP for 100 iterations. The model achieves an AUC of  $51.5\% \pm 0.0$ . Considering *only* the center pixel’s location provides a slight learnable difference.

**Backgrounds are significantly different for images with people of different gender expressions.** Next we occlude information about the person but keep the background intact (see purple bars in Fig. 3). For both MaskSegm ( $79.6\% \pm 0.3$ ) and MaskRect ( $70.8\% \pm 0.3$ ) the model performs better than random chance on COCO. The same is true for MaskRect for OpenImages ( $63.1\% \pm 0.4$ ). Here, our occlusions match those from proposed mitigation strategies [31, 76]. Once again, occluding only the person’s appearance is insufficient to remove all the gender artifacts in the image. Thus, any attempt to remove person-related gender artifacts shifts the source of artifact information from the person to the background, thus not satisfying the goal of removing these artifacts altogether.

**Qualitative analysis validates our findings.** To better understand the impact of background as a gender artifact, we qualitatively inspect images with the highest and lowest scores from our gender artifact model (Fig. 4). In (c) MaskSegm, the person’s location, together with background, is used as a gender artifact; the model associates city backgrounds with “female” and athletic scenes with “male.” Also, from (d) MaskSegm NoBg, we corroborate

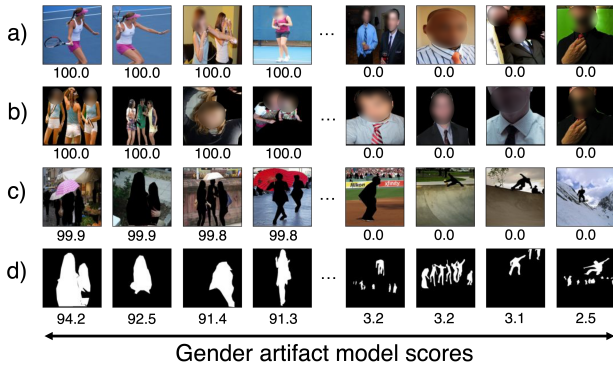


Figure 4. Images with highest and lowest scores for models trained on the following occlusions: a) Full, b) Full NoBg, c) MaskSegm, and d) MaskSegm NoBg. 100 corresponds to the model predicting “female ” and 0 to “male.” Face pixelation is not in the COCO image but included to partially preserve privacy.

Table 1. Out-of-distribution evaluation between the two datasets (e.g., training on COCO and evaluating on OpenImages) and the AUC for the three occlusion settings. We report the 95% confidence interval from bootstrapping.

Train / Test	Full	MaskRect	MaskRect NoBg
COCO / COCO	93.4 ±0.1	70.7 ±0.1	58.3 ±0.2
COCO / OI	91.2 ±0.2	60.7 ±0.4	63.2 ±0.4
OI / COCO	76.4 ±0.3	57.9 ±0.4	54.4 ±0.4
OI / OI	81.2 ±0.2	63.3 ±0.4	59.7 ±0.7

our finding that the model learns to associate larger, static poses with “female” and smaller, active poses with “male.” However, we do note that when the person pixels are included, the model is able to make more confident predictions for images with non-stereotypical settings (e.g., person labeled female playing tennis, person labeled male in a static, close-up shot) as seen in rows (a) and (b).

**Gender artifacts generalize across datasets.** Finally, we examine whether gender artifacts are dataset-specific, i.e., resulting from a dataset’s specific collection process and distribution. We conduct out-of-distribution evaluation (e.g., training on COCO and evaluating on the same occlusion setting on OpenImages). Across occlusion settings, the models perform above random chance even when evaluated on an out-of-distribution test set (Tbl. 1). This suggests that these gender artifacts may be common across visual datasets. To verify that our findings are not restricted to models pre-trained on ImageNet, for which there are known gender biases [10, 79], we also use a model pre-trained on other datasets and obtain similar results (see Appendix C).

## 6. Contextual Objects

In the previous section, we discovered that gender artifacts are present in both person and background pixels. In this section, we work to better understand these gender

artifacts in the background, or non-person parts of the image, and distinguish between two components of the background: the objects<sup>4</sup> and the “scenes” (i.e., the rest).

Prior works [31, 76] often assume gender artifacts in the background (i.e., objects or scene) are forms of “bias” and aim to mitigate them by studying object-gender co-occurrences. Additionally, prior works such as Singh et al. [68] point out the dangers of bias from contextual objects, particularly in visual recognition datasets, and propose mitigation techniques to recognize an object or attribute in the absence of its typical context and account for imbalances in object co-occurrences.

In this section, we describe how we identify the relative gender artifact contributions of the objects, as well as which objects contribute to the model’s ability to differentiate images of people of different genders.

**Top objects contributing to gender prediction.** We analyze the role of specific objects in gender prediction by training a logistic regression model using a binary presence-absence vector as input, where each value of the vector corresponds to the presence of an annotated object (1 if the object is present and 0 otherwise).<sup>5</sup> COCO contains 80 objects chosen by the dataset creators as a representative set of all categories, relevant to practical applications, and occurring with high enough frequency [42]. OpenImages contains 599 objects, but explicitly gendered objects (i.e., woman, man, girl, boy) are removed before training the classifier [40]. The model achieves an AUC of 75.4% on COCO (Fig. 5) and 63.3% on OpenImages. The performance of the OpenImages classifier is significantly lower than that of COCO; this may be attributable to the difference in annotated objects in the train and test distribution (the OpenImages test set has an average of 7.7 object instances annotated per image while the training set has an average of 3.7).<sup>6</sup>

This classifier’s weights loosely correspond to the object’s role in gender classification. We use the weights to identify the most relevant objects that contribute to gender prediction sorted by importance. For COCO, the *bottom* 10 classifier weights correspond to the top 10 objects most useful to classify male (e.g., skateboard, tie, snowboard). The *top* 10 classifier weights correspond to the top 10 objects most useful to classify female (e.g., hair drier, handbag, teddy bear) (see Fig. 5).

To further understand the role of contextual objects, we iteratively remove objects from the logistic regression classifier (by decreasing the input vector dimension) to “break” the classifier. That is, we want to see how many objects need to be removed before the classifier performs at random

<sup>4</sup>Albeit limited to the objects labeled by COCO and Open Images.

<sup>5</sup>Trained with L1 regularization, reg. strength of 1, and liblinear solver.

<sup>6</sup>OpenImages dataset creators generated denser labels for the validation and test splits than for the train split [40].




Figure 5. The change in AUC (%) as objects are iteratively removed from a logistic regression classifier trained on COCO (left). Our goal is to see how many objects must be removed before the classifier performs at random chance. When removing random objects, we report the mean and confidence interval from bootstrapping and train five separate random classifiers (see Appendix D). We also show the 10 most relevant objects in descending order, as identified by the weights of the logistic regression classifier (right).

chance. We remove the “most gendered objects” as defined by the object-based logistic regression model’s weights as shown in the table in Fig. 5. In the first iteration, the objects corresponding to the highest and lowest classifier weights are removed and the model is retrained. For the second iteration, the objects corresponding to the top 2 highest and top 2 lowest weights are removed, etc. Interestingly, while removing the gendered objects does force the AUC to decrease faster, we find that this classifier still performs above 50% AUC and can achieve an AUC of 72.1% when given 14 of the most gendered objects (as noted by iteratively removing the “least gendered objects” or objects corresponding to the lowest classifier weights). As shown in Fig. 5, we also experiment with removing randomly selected objects and find that even when given a small number of random objects, the gender artifact model is able to perform above random chance with an AUC of 66.8% trained only the one-hot encodings of 30 randomly chosen objects among the 80 annotated objects and 58.0% with 10 random objects.

The analysis of contextual object underscores the main takeaway: gender artifacts are everywhere. Even when only presented with one-hot encodings of objects, a gender artifact model can still perform above random chance.

## 7. Fairness through Blindness Methods

One common approach to fairness in computer vision is that of “fairness through blindness” where the classifier explicitly does not use the sensitive attribute e.g., gender. Such bias mitigation methods will remove gender artifacts and occlude gender-related features [32, 72, 78] (e.g., face, person segmentation mask, person bounding box [78]) to encourage the model to learn representations not reliant on gendered information. In this section, we use Wang et al.’s [76] adversarial approach as a case study and investi-

gate what gender artifacts are being removed or debiased.

**Method.** As in [76], we train a U-Net [58] on COCO training images to do gender debiasing by removing image pixels corresponding to salient gender artifacts. We then identify and analyze the pixels in COCO validation images (with  $\geq 1$  person object) that were removed in the debiasing process (see Appendix G for details).

**Results.** Intuitively, we would assume that fairness-through-blindness approaches remove gender artifacts by removing the person; however, in practice, we find the method removes artifacts that overwhelmingly come from the background, not the person. Of the pixels that are debiased (i.e., removed) 23.2% are in the person and 76.8% are in the background. Notably, this results in  $20.4\% \pm 15.7\%$  of all the background pixels are masked out (see Appendix G for more results). While fairness-through-blindness methods may purport to be removing gender cues primarily from the person, we find that much of the removed information comes from important parts of the image that may be necessary for the downstream task of object recognition.

## 8. Discussion

The results of most of our experiments are that the labelled gender expression is consistently discernible, i.e., classifiable at above random chance. This result holds even when an image is reduced to its average color (i.e. RGB values). Removal of such artifacts can be harmful if they lead to downstream allocational or representational harms (e.g., stereotyping, differences in performance) [4, 11].

### 8.1. Implications for bias mitigation

Fairness-through-blindness approaches are still common in computer vision; the community has not fully shifted to the fairness-aware approach [23, 43] this analysis recommends. Prior work in computer vision fairness [9, 31, 76] has operated under the notion that gender bias in datasets solely comes from a person’s appearance and not differences in the background or the average color of the image. Thus, these mitigation efforts attempt to remove gender bias by directly removing gender artifacts from the person. We engage with these prior works by showing that visual datasets have distinctly different distributions for images containing people of different genders. Similar to how proxies for protected attributes exist in tabular data [16, 17, 25] thus rendering efforts to naively remove these attributes ineffective, removing gender artifacts in computer vision bias mitigation techniques may also be ineffective.

Overall, our results have three main implications for bias mitigation methods. First, gender cannot be removed from a visual dataset without removing information that is relevant for the downstream task (e.g., removing objects, scene). Thus, fairness-through-blindness approaches in-

evitably remove information that affects the model’s performance. Second, while some methods assume removing part of a person’s appearance will remove the parts of the image that reveal gender [76], we find in practice that these methods heavily rely on artifacts outside of the person. Finally, we encourage researchers and practitioners to shift away from fairness-through-blindness approaches towards those considering *fairness through awareness* [22], as gender artifacts are simply too ubiquitous to consider trying to remain blind to. Existing fairness-aware methods first explicitly encode domain information and then explicitly mitigate the domain information [22]. One form of fairness-aware methods can include constructing decoupled classifiers [23] in which a separate classifier is trained for different sub-groups of a sensitive attribute such as gender. These algorithms do not try to erase gender, rather to accept its existence, and adjust algorithms accordingly.

## 8.2. Implications for dataset construction

The solution is not to accept different distributions of images for people of different genders as an inevitable conclusion and make no efforts during dataset construction to consider gender representation. Rather, we point out these potentially harmful correlations and encourage practitioners to decide what kinds of gender artifacts they deem permissible or not in each specific context.

The gender artifacts we discover and the differences in input distribution more broadly are not necessarily harmful. These differences can often represent historically important distinctions between groups of people. Moreover, gender, socially constructed as it is, is important to many social concerns about equal treatment [44, 56, 57]. A participatory and iterative process can help to elucidate which artifacts should be removed (e.g., offensive stereotypes) and which may be appropriate to persist depending on the context.

It is also important to consider the downstream task to determine what kinds of gender artifacts are permissible. One way of determining what kinds of artifacts may be more permissible is performing *disaggregated measures of evaluation*. Of course, there are many nuances here to consider as well [5], and awareness about prominent gender artifacts can help guide the selection of axes along which to perform disaggregated analysis (beyond just demographic groups). For example, if scene is known to be a prominent gender artifact and the downstream task is scene classification, evaluation could be done across demographic group as well as scene type, to understand if specific scenes are always misclassified for people of a particular gender. Another way can be through engaging with communities affected by these downstream technologies [8, 11, 37], as some gender effects such as those encoding harmful stereotypes, may be inherently harmful.

Finally, it is important to note we are only stating the

*presence* of such artifacts and not claiming these artifacts are necessarily harmful. While our analysis reveals these artifacts do exist in datasets, it does not guarantee the artifacts will be captured by a task-specific downstream problem or, even if they are captured, whether they lead to any harms. Nonetheless, it can be useful for practitioners to identify artifacts, following a similar process to that used in this work, and provide documentation for dataset users.

## 8.3. Incoherence of gender predictions

Finally, our analysis of gender artifacts elucidates another takeaway: we might consider that what our gender artifact model (i.e., a gender classifier) predicts in these contexts to be a nearly incoherent concept. The fact that the outline of a person, or even a pixelated 7x7 square, leads to high predictive performance of the notion of gender we have operationalized, should lead us to wonder about what exactly the model is predicting. Perhaps what the model is actually outputting is “the gender expression of the person most likely to be holding an umbrella,” or “the gender expression of the person most likely to be in a running pose.” All of these are notably different from predicting the person’s gender.

While normative arguments against gender prediction [39, 41, 62] are sufficient to hinder its usage, our results suggest that a model’s prediction of gender has little to do with the notion of gender the model designer is interested in. Although explicit gender prediction is now less common due to consciousness-raising on the topic, it is still prevalent in image captioning [7, 9]. As prior work [61] has shown, different gender classification systems frequently produce conflicting gender annotations. The concept of “gender” the model learns is as likely to be the most salient visual attribute correlated with gender in the dataset, as it is any kind of gender performance we deem socially meaningful [14]. We must be careful should we ever find ourselves imbuing any meaning or value to the “gender” prediction of a model.

## 9. Conclusion

Given the importance of datasets in computer vision, works critically examining these datasets provide insight into their limitations and have implications for the many methods evaluated on them. We go beyond prior works that analyze only annotated attributes and conduct a more comprehensive analysis of gender artifacts by studying what a “gender classifier” relies on to make predictions. We analyze gender artifacts by occluding various combinations of color, person and background, and contextual objects. Gender is consistently discernible despite different manipulations to the input image. This analysis highlights the infeasibility of removing gender artifacts as a means for bias mitigation. Further, our work points to a general incoherence

in the outputs of gender prediction models, as they appear to rely on any number of spurious correlations.

The permeation of gender artifacts throughout image datasets reveals it is futile to control for all possible differences. The aim for researchers should not be to remove these artifacts, but rather critically consider which gender artifacts are permissible and further to propose methods that are robust to these distribution differences.

**Acknowledgements.** This material is based upon work supported by the National Science Foundation under Grant No. 1763642, Grant No. 2112562, Grant No. 2145198, and the Graduate Research Fellowship to A.W., as well as by the Princeton Engineering Project X Fund. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

## References

- [1] Mohsan S. Alvi, Andrew Zisserman, and Christoffer Nellåker. Turning a blind eye: Explicit removal of biases and variation from deep neural network embeddings. In *European Conference on Computer Vision Workshops (ECCVW)*, 2018.
- [2] Yuki M. Asano, Christian Rupprecht, Andrew Zisserman, and Andrea Vedaldi. Pass: An imagenet replacement for self-supervised pretraining without humans. In *Neural Information Processing Systems (NeurIPS)*, 2021.
- [3] Guha Balakrishnan, Yuanjun Xiong, Wei Xia, and Pietro Perona. Towards causal benchmarking of bias in face analysis algorithms. In *European Conference on Computer Vision (ECCV)*, 2020.
- [4] Solon Barocas, Kate Crawford, Aaron Shapiro, and Hanna Wallach. The problem with bias: Allocative versus representational harms in machine learning. In *Annual Conference of the Special Interest Group for Computing, Information and Society (SIGCIS)*, 2017.
- [5] Solon Barocas, Anhong Guo, Ece Kamar, Jacquelyn Kronen, Meredith Ringel Morris, Jennifer Wortman Vaughan, Duncan Wadsworth, and Hanna Wallach. Designing disaggregated evaluations of AI systems: Choices, considerations, and tradeoffs. In *AAAI/ACM Conference on AI, Ethics, and Society (AIES)*, 2021.
- [6] Solon Barocas, Moritz Hardt, and Arvind Narayanan. Fairness and machine learning limitations and opportunities. 2018.
- [7] Cynthia L. Bennett, Cole Gleason, Morgan Klaus Scheuerman, Jeffrey P. Bigham, Anhong Guo, and Alexandra To. “It’s complicated”: Negotiating accessibility and (mis)representation in image descriptions of race, gender, and disability. *ACM Conference on Human Factors in Computing Systems (CHI)*, 2021.
- [8] Cynthia L Bennett and Os Keyes. What is the point of fairness? Disability, AI and the complexity of justice. *ACM SIGACCESS Accessibility and Computing*, (125):1–1, 2020.
- [9] Shruti Bhargava and David Forsyth. Exposing and correcting the gender bias in image captioning datasets and models. *arXiv:1912.00578*, 2019.
- [10] Abeba Birhane and Vinay Uday Prabhu. Large image datasets: A pyrrhic win for computer vision? In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2021.
- [11] Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. Language (technology) is power: A critical survey of “bias” in NLP. In *Association for Computational Linguistics (ACL)*, 2020.
- [12] Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. Man is to computer programmer as woman is to homemaker? Debiasing word embeddings. In *Neural Information Processing Systems (NeurIPS)*, 2016.
- [13] Joy Buolamwini and Timnit Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *ACM Conference on Fairness, Accountability and Transparency (FAccT)*, 2018.
- [14] Judith Butler. *Gender trouble*. Routledge, 2002.
- [15] Aylin Caliskan, Joanna J Bryson, and Arvind Narayanan. Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334):183–186, 2017.
- [16] Simon Caton and Christian Haas. Fairness in machine learning: A survey. *arXiv preprint arXiv:2010.04053*, 2020.
- [17] Sam Corbett-Davies and Sharad Goel. The measure and mismeasure of fairness: A critical review of fair machine learning. *arXiv preprint arXiv:1808.00023*, 2018.
- [18] Kathleen Creel and Deborah Hellman. The algorithmic leviathan: Arbitrariness, fairness, and opportunity in algorithmic decision making systems. In *ACM Conference on Fairness, Accountability, and Transparency (FAccT)*, 2021.
- [19] Terrance De Vries, Ishan Misra, Changhan Wang, and Laurens Van der Maaten. Does object recognition work for everyone? In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2019.
- [20] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009.
- [21] Emily Denton, Ben Hutchinson, Margaret Mitchell, and Timnit Gebru. Detecting bias with generative counterfactual face attribute augmentation. *arXiv e-prints*, pages arXiv–1906, 2019.
- [22] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Rich Zemel. Fairness through awareness. In *Innovations in Theoretical Computer Science Conference (ITCS)*, 2012.
- [23] Cynthia Dwork, Nicole Immorlica, Adam Tauman Kalai, and Max Leiserson. Decoupled classifiers for group-fair and efficient machine learning. In *ACM Conference on Fairness, Accountability, and Transparency (FAccT)*, 2018.
- [24] Francis Ferraro, Nasrin Mostafazadeh, Lucy Vanderwende, Jacob Devlin, Michel Galley, Margaret Mitchell, et al. A



- survey of current datasets for vision and language research. *arXiv preprint arXiv:1506.06833*, 2015.
- [25] Aaron Fisher, Cynthia Rudin, and Francesca Dominici. All models are wrong, but many are useful: Learning a variable’s importance by studying an entire class of prediction models simultaneously. *Journal of Machine Learning Research (JMLR)*, 20(177):1–81, 2019.
- [26] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé Iii, and Kate Crawford. Datasheets for datasets. *Communications of the ACM*, 64(12):86–92, 2021.
- [27] Judy Wawira Gichoya, Imon Banerjee, Ananth Reddy Bhimireddy, John L Burns, Leo Anthony Celi, Li-Ching Chen, Ramon Correa, Natalie Dullerud, Marzyeh Ghassemi, Shih-Cheng Huang, Po-Chih Kuo, Matthew P Lungren, Lyle J Palmer, Brandon J Price, Saptarshi Purkayastha, Ayis T Pyros, Lauren Oakden-Rayner, Chima Okechukwu, Laleh Seyyed-Kalantari, Hari Trivedi, Ryan Wang, Zachary Zaiman, and Haoran Zhang. AI recognition of patient race in medical imaging: a modelling study. *The Lancet Digital Health*, 4(6), 2022.
- [28] Foad Hamidi, Morgan Klaus Scheuerman, and Stacy M Branham. Gender recognition or gender reductionism? the social implications of embedded gender recognition systems. 2018.
- [29] Moritz Hardt, Eric Price, Eric Price, and Nati Srebro. Equality of opportunity in supervised learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2016.
- [30] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [31] Lisa Anne Hendricks, Kaylee Burns, Kate Saenko, Trevor Darrell, and Anna Rohrbach. Women also snowboard: Overcoming bias in captioning models. In *European Conference on Computer Vision (ECCV)*, 2018.
- [32] Sarah Holland, Ahmed Hosny, Sarah Newman, Joshua Joseph, and Kasia Chmielinski. The dataset nutrition label: A framework to drive higher data quality standards. *arXiv preprint arXiv:1805.03677*, 2018.
- [33] MD Zakir Hossain, Ferdous Sohel, Mohd Fairuz Shiratuddin, and Hamid Laga. A comprehensive survey of deep learning for image captioning. *ACM Computing Surveys (CSUR)*, 51(6), 2019.
- [34] Frederick Howard, James Dolezal, Sara Kochanny, Jeffrey Schulte, Heather Chen, Lara Heij, Dezheng Huo, Rita Nanda, Olufunmilayo Olopade, Jakob Kather, Nicole Cipriani, Robert Grossman, and Alexander Pearson. The impact of site-specific digital histology signatures on deep learning model accuracy and bias. *Nature Communications*, 12, 2021.
- [35] Ben Hutchinson, Andrew Smart, Alex Hanna, Emily Denton, Christina Greer, Oddur Kjartansson, Parker Barnes, and Margaret Mitchell. Towards accountability for machine learning datasets: Practices from software engineering and infrastructure. In *ACM Conference on Fairness, Accountability, and Transparency (FAccT)*, 2021.
- [36] Eun Seo Jo and Timnit Gebru. Lessons from archives: Strategies for collecting sociocultural data in machine learning. In *ACM Conference on Fairness, Accountability, and Transparency (FAccT)*, 2020.
- [37] Michael Katell, Meg Young, Dharma Dailey, Bernease Herman, Vivian Guetler, Aaron Tam, Corinne Bintz, Daniella Raz, and PM Krafft. Toward situated interventions for algorithmic equity: lessons from the field. In *ACM Conference on Fairness, Accountability, and Transparency (FAccT)*, 2020.
- [38] Matthew Kay, Cynthia Matuszek, and Sean A Munson. Unequal representation and gender stereotypes in image search results for occupations. In *Conference on Human Factors in Computing Systems (CHI)*, 2015.
- [39] Os Keyes. The misgendering machines: Trans/HCI implications of automatic gender recognition. *ACM Conference on Computer Supported Cooperative Work (CSCW)*, 2018.
- [40] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, et al. The Open Images Dataset V4. *International Journal of Computer Vision*, 128(7):1956–1981, 2020.
- [41] Brian Larson. Gender as a variable in natural-language processing: Ethical considerations. In *ACL Workshop on Ethics in Natural Language Processing*, 2017.
- [42] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: Common objects in context. In *European Conference on Computer Vision (ECCV)*, 2014.
- [43] Xiaofeng Lin, Seungbae Kim, and Jungseock Joo. Fairgrape: Fairness-aware gradient pruning method for face attribute classification. In *European Conference on Computer Vision (ECCV)*, 2022.
- [44] Judith Lorber. Gender inequality. *Los Angeles, CA: Roxbury*, 2001.
- [45] Michael A Madaio, Luke Stark, Jennifer Wortman Vaughan, and Hanna Wallach. Co-designing checklists to understand organizational challenges and opportunities around fairness in ai. In *Conference on Human Factors in Computing Systems (CHI)*, 2020.
- [46] Shervin Minaee, Yuri Y Boykov, Fatih Porikli, Antonio J Plaza, Nasser Kehtarnavaz, and Demetri Terzopoulos. Image segmentation using deep learning: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2021.
- [47] Vidya Muthukumar, Tejaswini Pedapati, Nalini Ratha, Prasanna Sattigeri, Chai-Wah Wu, Brian Kingsbury, Abhishek Kumar, Samuel Thomas, Aleksandra Mojsilović, and Kush R. Varshney. Color-theoretic experiments to understand unequal gender classification accuracy from face images. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2019.
- [48] Safiya Umoja Noble. *Algorithms of Oppression: How Search Engines Reinforce Racism*. NYU Press, 2018.
- [49] Jahna Otterbacher, Pinar Barlas, Styliani Kleanthous, and Kyriakos Kyriakou. How do we talk about other people? group (un) fairness in natural language image descriptions. In *AAAI Conference on Human Computation and Crowdsourcing (HCOMP)*, 2019.



- [50] Jahna Otterbacher, Jo Bates, and Paul Clough. Competent men and warm women: Gender stereotypes and backlash in image search results. In *Conference on Human Factors in Computing Systems (CHI)*, 2017.
- [51] Joon Sung Park, Michael S Bernstein, Robin N Brewer, Ece Kamar, and Meredith Ringel Morris. Understanding the representation and representativeness of age in ai data sets. In *AAAI/ACM Conference on AI, Ethics, and Society (AIES)*, 2021.
- [52] Amandalynne Paullada, Inioluwa Deborah Raji, Emily M Bender, Emily Denton, and Alex Hanna. Data and its (dis) contents: A survey of dataset development and use in machine learning research. *Patterns*, 2(11), 2021.
- [53] Kenneth L Peng, Arunesh Mathur, and Arvind Narayanan. Mitigating dataset harms requires stewardship: Lessons from 1000 papers. In *Neural Information Processing Systems (NeurIPS)*, 2021.
- [54] Ryan Poplin, Avinash Vaidyanathan Varadarajan, Katy Blumer, Yun Liu, Mike McConnell, Greg Corrado, Lily Peng, and Dale Webster. Predicting cardiovascular risk factors in retinal fundus photographs using deep learning. *Nature Biomedical Engineering*, 2018.
- [55] Vikram V. Ramaswamy, Sunnie S. Y. Kim, and Olga Russakovsky. Fair attribute classification through latent space de-biasing. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [56] Cecilia L Ridgeway. *Framed by gender: How gender inequality persists in the modern world*. Oxford University Press, 2011.
- [57] Barbara J Risman. Gender as a social structure. In *Handbook of the Sociology of Gender*, pages 19–43. Springer, 2018.
- [58] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015.
- [59] 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. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015.
- [60] Prasanna Sattigeri, Samuel C Hoffman, Vijil Chenthamarakshan, and Kush R Varshney. Fairness GAN: Generating datasets with fairness properties using a generative adversarial network. *IBM Journal of Research and Development*, 63(4/5):3–1, 2019.
- [61] Morgan Klaus Scheuerman, Jacob M Paul, and Jed R Brubaker. How computers see gender: An evaluation of gender classification in commercial facial analysis services. In *ACM Conference on Computer Supported Cooperative Work (CSCW)*, 2019.
- [62] Morgan Klaus Scheuerman, Kandrea Wade, Caitlin Lustig, and Jed R. Brubaker. How we’ve taught algorithms to see identity: Constructing race and gender in image databases for facial analysis. In *ACM Conference on Computer Supported Cooperative Work (CSCW)*, 2020.
- [63] Candice Schumann, Susanna Ricco, Utsav Prabhu, Vittorio Ferrari, and Caroline Pantofaru. A step toward more inclusive people annotations for fairness. In *AAAI/ACM Conference on AI, Ethics, and Society (AIES)*, 2021.
- [64] Carsten Schwemmer, Carly Knight, Emily D Bello-Pardo, Stan Oklobdzija, Martijn Schoonvelde, and Jeffrey W Lockhart. Diagnosing gender bias in image recognition systems. *Socius*, 6:2378023120967171, 2020.
- [65] Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-CAM: Visual explanations from deep networks via gradient-based localization. In *International Conference on Computer Vision (ICCV)*, 2017.
- [66] Shreya Shankar, Yoni Halpern, Eric Breck, James Atwood, Jimbo Wilson, and D Sculley. No classification without representation: Assessing geodiversity issues in open data sets for the developing world. In *Neural Information Processing Systems Workshops (NeurIPSW)*, 2017.
- [67] Viktoriia Sharmanska, Lisa Anne Hendricks, Trevor Darrell, and Novi Quadrianto. Contrastive examples for addressing the tyranny of the majority. *arXiv preprint arXiv:2004.06524*, 2020.
- [68] Krishna Kumar Singh, Dhruv Mahajan, Kristen Grauman, Yong Jae Lee, Matt Feiszli, and Deepti Ghadiyaram. Don’t judge an object by its context: Learning to overcome contextual bias. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [69] Pierre Stock and Moustapha Cisse. Convnets and imagenet beyond accuracy: Understanding mistakes and uncovering biases. In *European Conference on Computer Vision (ECCV)*, 2018.
- [70] Ruixiang Tang, Mengnan Du, Yuening Li, Zirui Liu, Na Zou, and Xia Hu. Mitigating gender bias in captioning systems. In *The Web Conference*, 2021.
- [71] Antonio Torralba. How many pixels make an image? *Visual Neuroscience*, 26(1):123–131, 2009.
- [72] Antonio Torralba and Alexei A Efros. Unbiased look at dataset bias. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011.
- [73] Emiel van Miltenburg. Stereotyping and bias in the Flickr30K dataset. In *Multimodal Corpora: Computer Vision and Language Processing (MMC)*, 2016.
- [74] Angelina Wang, Arvind Narayanan, and Olga Russakovsky. REVISE: A tool for measuring and mitigating bias in visual datasets. In *European Conference on Computer Vision (ECCV)*, 2020.
- [75] Jialu Wang, Yang Liu, and Xin Eric Wang. Are gender-neutral queries really gender-neutral? mitigating gender bias in image search. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2021.
- [76] Tianlu Wang, Jieyu Zhao, Mark Yatskar, Kai-Wei Chang, and Vicente Ordonez. Balanced datasets are not enough: Estimating and mitigating gender bias in deep image representations. In *International Conference on Computer Vision (ICCV)*, 2019.

- [77] Zeyu Wang, Klint Qinami, Ioannis Karakozis, Kyle Genova, Prem Nair, Kenji Hata, and Olga Russakovsky. Towards fairness in visual recognition: Effective strategies for bias mitigation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [78] Kai Yuanqing Xiao, Logan Engstrom, Andrew Ilyas, and Aleksander Madry. Noise or signal: The role of image backgrounds in object recognition. In *International Conference on Learning Representations (ICLR)*, 2021.
- [79] Kaiyu Yang, Klint Qinami, Li Fei-Fei, Jia Deng, and Olga Russakovsky. Towards fairer datasets: Filtering and balancing the distribution of the people subtree in the ImageNet hierarchy. In *ACM Conference on Fairness, Accountability and Transparency (FAT\*)*, 2020.
- [80] Rich Zemel, Yu Wu, Kevin Swersky, Toni Pitassi, and Cynthia Dwork. Learning fair representations. In *International Conference on Machine Learning (ICML)*, 2013.
- [81] Kaipeng Zhang, Zhanpeng Zhang, Zhifeng Li, and Yu Qiao. Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE Signal Processing Letters*, 23(10), 2016.
- [82] Dora Zhao, Angelina Wang, and Olga Russakovsky. Understanding and evaluating racial biases in image captioning. In *International Conference on Computer Vision (ICCV)*, 2021.
- [83] Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. Men also like shopping: Reducing gender bias amplification using corpus-level constraints. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2017.
- [84] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. In *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2017.