

---

# Directional Bias Amplification

---

Angelina Wang and Olga Russakovsky  
Princeton University

## Abstract

Mitigating bias in machine learning systems requires refining our understanding of bias propagation pathways: from societal structures to large-scale data to trained models to impact on society. In this work, we focus on one aspect of the problem, namely *bias amplification*: the tendency of models to amplify the biases present in the data they are trained on. A metric for measuring bias amplification was introduced in the seminal work by Zhao et al. (2017); however, as we demonstrate, this metric suffers from a number of shortcomings including conflating different types of bias amplification and failing to account for varying base rates of protected attributes. We introduce and analyze a new, decoupled metric for measuring bias amplification, BiasAmp<sub>→</sub> (Directional Bias Amplification). We thoroughly analyze and discuss both the technical assumptions and normative implications of this metric. We provide suggestions about its measurement by cautioning against predicting sensitive attributes, encouraging the use of confidence intervals due to fluctuations in the fairness of models across runs, and discussing the limitations of what this metric captures. Throughout this paper, we work to provide an interrogative look at the technical measurement of bias amplification, guided by our normative ideas of what we want it to encompass. Code is located at <https://github.com/princetonvisualai/directional-bias-amp>.

## 1. Introduction

The machine learning community is becoming increasingly cognizant of problems surrounding fairness and bias, and

Correspondence to: Angelina Wang <angelina.wang@princeton.edu>.

*Proceedings of the 38<sup>th</sup> International Conference on Machine Learning*, PMLR 139, 2021. Copyright 2021 by the author(s).

correspondingly, a plethora of new algorithms and metrics are being proposed (see e.g., Mehrabi et al. (2019) for a survey). The analytic gatekeepers of the systems often take the form of fairness evaluation metrics, and it is vital that these be deeply investigated both technically and normatively. In this paper, we endeavor to do this for bias amplification.

Bias amplification occurs when a model exacerbates biases from the training data at test time. It is the result of the algorithm (Foulds et al., 2018), and unlike some other forms of bias, cannot be solely attributed to the dataset.

**Directional bias amplification metric.** We propose a new way of measuring bias amplification, BiasAmp<sub>→</sub> (Directional Bias Amplification),<sup>1</sup> that builds off a prior metric from “Men Also Like Shopping: Reducing Gender Bias Amplification using Corpus-level Constraints” (Zhao et al., 2017), that we will call BiasAmp<sub>MALS</sub>. Our metric’s technical composition aligns with the real-world qualities we want it to encompass, addressing a number of the previous metric’s shortcomings by being able to: 1) focus on both positive and negative correlations, 2) take into account the base rates of each protected attribute, and most importantly 3) disentangle the directions of amplification.

As an example, consider a visual dataset (Fig. 1) where each image has a label for the task  $T$ , which is painting or not painting, and further is associated with a protected attribute  $A$ , which is woman or man.<sup>2</sup> If the gender of the person biases the prediction of the task, we consider this  $A \rightarrow T$  bias amplification; if the reverse happens, then  $T \rightarrow A$ . Bias amplification as it is currently being measured merges together these two different paths which have different normative implications and therefore demand different remedies. This speaks to a larger problem of imprecision when discussing problems of bias (Blodgett et al., 2020). For example, “gender bias” can be vague; it is unclear if the system is assigning gender in a biased way, or if there is a disparity in model performance between different genders. Both are harmful in different ways, but the conflation of

<sup>1</sup>The arrow is meant to signify the direction that bias amplification is flowing, and not intended to be a claim about causality.

<sup>2</sup>We use the terms *man* and *woman* to refer to binarized socially-perceived (frequently annotator-inferred) gender expression, recognizing these labels are not inclusive and may be inaccurate.





Dataset				
Ground truth	(painting, woman)	(painting, man)	(not painting, woman)	(not painting, man)
Prediction Errors				
<div><div>A→T Bias Amp</div><div>T→A Bias Amp</div><div>No Bias Amp</div></div>	Task wrong (not painting, ---)	(not painting, ---)	(painting, ---)	(painting, ---)
Attribute wrong	(---, man)	(---, woman)	(---, man)	(---, woman)

Figure 1. Consider an image dataset where the goal is to classify the task,  $T$ , as painting or not painting, and the attribute,  $A$ , as woman or man. Women are correlated with painting, and men with not painting. In this work we are particularly concerned with errors contributing to the amplification of bias, i.e., existing training correlations (yellow and red in the figure). We further disentangle these errors into those that amplify the attribute to task correlation (i.e., incorrectly predict the task based on the person’s attribute; shown in yellow) versus those that amplify the task to attribute (shown in red).

these biases can lead to misdirected solutions.

**Bias amplification analysis.** The notion of bias amplification allows us to encapsulate the idea that systemic harms and biases can be more harmful than errors made without such a history (Bearman et al., 2009). For example, in images, overclassifying women as cooking carries a more negative connotation than overclassifying men as cooking. The distinction of which errors are more harmful can often be determined by lifting the patterns from the training data.

In our analysis and normative discussion, we look into this and other implications through a series of experiments. We consider whether predicting protected attributes is necessary in the first place; by not doing so, we can trivially remove  $T \rightarrow A$  amplification. We also encourage the use of confidence intervals because  $\text{BiasAmp}_{\rightarrow}$ , along with other fairness metrics, suffers from the Rashomon Effect (Breiman, 2001), or multiplicity of good models. In other words, in supervised machine learning, random seeds have relatively little impact on *accuracy*; in contrast, they appear to have a greater impact on *fairness*.

Notably, a trait of bias amplification is that it is not at odds with accuracy. Bias amplification measures the model’s errors, so a model with perfect accuracy will have perfect (zero) bias amplification. (Note nevertheless that the metrics are not always correlated.) This differs from many other fairness metrics, because the goal of not amplifying biases and thus matching task-attribute correlations is aligned with that of accurate predictions. For example, satisfying fairness metrics like demographic parity (Dwork et al., 2012) are incompatible with perfect accuracy when parity is not met in the ground-truth. For the same reason bias amplification permits a classifier with perfect accuracy, it also comes with a set of limitations that are associated with treating data correlations as the desired ground-truth, and thus make it less appropriate for social applications where other metrics are better suited for measuring a fair allocation of resources.

**Outline.** To ground our work, we first distinguish what bias amplification captures that standard fairness metrics cannot, then distinguish  $\text{BiasAmp}_{\rightarrow}$  from  $\text{BiasAmp}_{\text{MALS}}$ . Our key contributions are: 1) proposing a new way to measure bias amplification, addressing multiple shortcomings of prior work and allowing us to better diagnose models, and 2) providing a comprehensive technical analysis and normative discussion around the use of this measure in diverse settings, encouraging thoughtfulness with each application.

## 2. Related Work

**Fairness Measurements.** Fairness is nebulous and context-dependent, and approaches to quantifying it include equalized odds (Hardt et al., 2016), equal opportunity (Hardt et al., 2016), demographic parity (Dwork et al., 2012; Kusner et al., 2017), fairness through awareness (Dwork et al., 2012; Kusner et al., 2017), fairness through unawareness (Grgic-Hlaca et al., 2016; Kusner et al., 2017), and treatment equality (Berk et al., 2017). We examine bias amplification, a type of group fairness where correlations are amplified.

As an example of what differentiates bias amplification, we present a scenario based on Fig. 1. We want to classify a person whose attribute is man or woman with the task of painting or not. The majority groups “(painting, woman)” and “(not painting, man)” each have 30 examples, and the minority groups “(not painting, woman)” and “(painting, man)” each have 10. A classifier trained to recognize painting on this data is likely to learn these associations and over-predict painting on images of women and under-predict painting on images of men; however, algorithmic interventions may counteract this and result in the opposite behavior. In Fig. 2 we show how four standard fairness metrics (in blue) vary under different amounts of learned amplification: FPR difference, TPR difference (Chouldechova, 2016; Hardt et al., 2016), accuracy difference in task prediction (Berk et al., 2017), and average mean-per-class accuracy across sub-

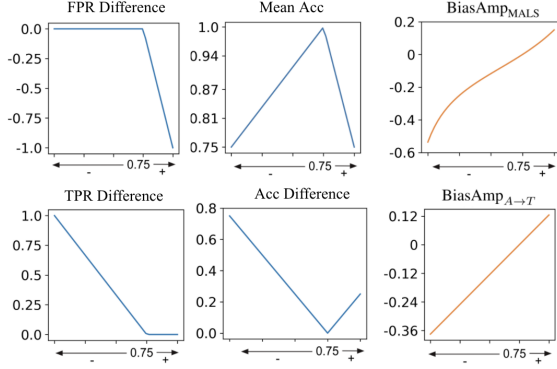


Figure 2. Fairness metrics vary in how they respond to model errors. In our image dataset (Fig. 1) of predicting someone who is a woman or man to be painting or not, we consider a painting classifier that always predicts the task correctly for men, but varies for women. The x-axes correspond to the percentage of women predicted to be painting, where the ground-truth is 0.75. Below that the model is under-predicting women to be painting, and above it the model is over-predicting. The two metrics in the first column, FPR and TPR difference, only capture one of under- or over-prediction. The next two metrics in the second column, accuracy difference between attribute subgroups and average mean-per-class accuracy across attribute subgroups, are symmetric around 0.75, being unable to differentiate. Thus, the bias amplification metrics in the last column are needed to distinguish between under- and over-prediction ( $\text{BiasAmp}_{\text{MALS}}$  from Zhao et al. (2017) in Sec. 3.2, and our proposed  $\text{BiasAmp}_{A \rightarrow T}$  in Sec. 4).  $\text{BiasAmp}_{\text{MALS}}$  requires attribute predictions, so we assume perfect attribute prediction here to make the comparison.

groups (Buolamwini & Gebru, 2018). However, these four metrics are not designed to account for the training correlations, and unable to distinguish between cases of increased or decreased learned correlations, motivating a need for a measurement that can: bias amplification.

**Bias Amplification.** Bias amplification has been measured by looking at binary classifications without attributes (Leino et al., 2019), GANs (Jain et al., 2020; Choi et al., 2020), and correlations (Zhao et al., 2017; Jia et al., 2020). We consider attributes in our formulation, which is a classification setting, and thus differs from GANs. We dissect in detail the correlation work, and propose measuring *conditional* correlations, which we term “directional.” Wang et al. (2019) measures amplification by predicting the sensitive attribute from the model outputs, thus relying on multiple target labels simultaneously; we propose a decomposable metric to allow for more precise model diagnosis.

The Word Embedding Association Test (WEAT) (Caliskan et al., 2017) measures bias amplification in de-contextualized word embeddings, specifically, non-conditional correlations (Bolukbasi et al., 2016). However, with newer models like BERT and ELMo that have

contextualized embeddings, WEAT does not work (May et al., 2019), so new techniques have been proposed incorporating context (Lu et al., 2019; Kuang & Davison, 2016). We use these models to measure the directional aspect of amplifications, as well as to situate them in the broader world of bias amplification. Directionality of amplification has been observed (Stock & Cisse, 2018; Qian et al., 2019), but we take a more systematic approach.

**Causality.** Bias amplification is also studied in the causal statistics literature (Bhattacharya & Vogt, 2007; Wooldridge, 2016; Pearl, 2010, 2011; Middleton et al., 2016). However, despite the same terminology, the definitions and implications are largely distinct. Our work follows the machine learning bias amplification literature discussed in the previous section and focuses on the amplification of socially-relevant correlations in the training data.

**Predictive Multiplicity.** The Rashomon Effect (Breiman, 2001), or multiplicity of good models, has been studied in various contexts. The variables investigated that differ across good models include explanations (Hancox-Li, 2020), individual treatments (Marx et al., 2020; Pawelczyk et al., 2020), and variable importance (Fisher et al., 2019; Dong & Rudin, 2019). We build on these and investigate how fairness also differs between equally “good” models.

### 3. Existing Bias Amplification Metric

We describe the existing metric (Zhao et al., 2017) and highlight shortcomings that we address in Sec. 4.

#### 3.1. Notation

Let  $\mathcal{A}$  be the set of protected demographic groups: for example,  $\mathcal{A} = \{\text{woman}, \text{man}\}$  in Fig. 1.  $A_a$  for  $a \in \mathcal{A}$  is the binary random variable corresponding to the presence of the group  $a$ ; thus  $P(A_{\text{woman}} = 1)$  can be empirically estimated as the fraction of images in the dataset containing women. Note that this formulation is generic enough to allow for multiple protected attributes and intersecting protected groups. Let  $T_t$  with  $t \in \mathcal{T}$  similarly correspond to binary target tasks, e.g.,  $\mathcal{T} = \{\text{painting}\}$  in Fig. 1.

#### 3.2. Formulation and shortcomings

Using this notation, Zhao et al. (2017)’s metric is:

$$\text{BiasAmp}_{\text{MALS}} = \frac{1}{|\mathcal{T}|} \sum_{a \in \mathcal{A}, t \in \mathcal{T}} y_{at} \Delta_{at} \quad (1)$$

$$\text{with } y_{at} = \mathbb{1} \left[ P(A_a = 1 | T_t = 1) > \frac{1}{|\mathcal{A}|} \right]$$

$$\Delta_{at} = P(\hat{A}_a = 1 | \hat{T}_t = 1) - P(A_a = 1 | T_t = 1)$$

where  $\hat{A}_a$  and  $\hat{T}_t$  denote model predictions for the protected group  $a$  and the target task  $t$  respectively. One attractive

property of this metric is that it doesn't require any ground truth test labels: assuming the training and test distributions are the same,  $P(A_a = 1|T_t = 1)$  can be estimated on the training set, and  $P(\hat{A}_a = 1|\hat{T}_t = 1)$  relies only on the *predicted* test labels. However, it also has a number of shortcomings.

**Shortcoming 1: The metric focuses only on positive correlations.** This may lead to numerical inconsistencies, especially in cases with multiple protected groups.

To illustrate, consider a scenario with 3 protected groups  $A_1$ ,  $A_2$ , and  $A_3$  (disjoint; every person belongs to one), one binary task  $T$ , and the following dataset<sup>3</sup>:

- When  $A_1 = 1$ : **10** exs. of  $T = 0$  and **40** exs. of  $T = 1$
- When  $A_2 = 1$ : **40** exs. of  $T = 0$  and **10** exs. of  $T = 1$
- When  $A_3 = 1$ : **10** exs. of  $T = 0$  and **20** exs. of  $T = 1$

From Eq. 1, we see  $y_1 = 1$ ,  $y_2 = 0$ ,  $y_3 = 0$ . Now consider a model that always makes correct predictions of the protected attribute  $\hat{A}_a$ , always correctly predicts the target task when  $A_1 = 1$ , but predicts  $\hat{T} = 0$  whenever  $A_2 = 1$  and  $\hat{T} = 1$  whenever  $A_3 = 1$ . Intuitively, this would correspond to a case of overall learned bias amplification. However, Eqn. 1 would measure bias amplification as 0 since the strongest positive correlation (in the  $A_1 = 1$  group) is not amplified.

Note that this issue doesn't arise as prominently when there are only 2 disjoint protected groups (binary attributes), which was the case implicitly considered in Zhao et al. (2017). However, even with two groups there are miscalibration concerns. For example, consider the dataset above but only with the  $A_1 = 1$  and  $A_2 = 1$  examples. A model that correctly predicts the protected attribute  $\hat{A}_a$ , correctly predicts the task on  $A_1 = 1$ , yet predicts  $\hat{T} = 0$  whenever  $A_2 = 1$  would have bias amplification value of  $\Delta_1 = \frac{40}{40} - \frac{40}{50} = 0.2$ . However, a similar model that now correctly predicts the task on  $A_2 = 1$  but always predicts  $\hat{T} = 1$  on  $A_1 = 1$  would have a much smaller bias amplification value of  $\Delta_1 = \frac{50}{60} - \frac{40}{50} = 0.033$ , although intuitively the amount of bias amplification is the same.

**Shortcoming 2: The chosen protected group may not be correct due to imbalance between groups.** To illustrate, consider a scenario with 2 disjoint protected groups:

- When  $A_1 = 1$ : **60** exs. of  $T = 0$  and **30** exs. of  $T = 1$
- When  $A_2 = 1$ : **10** exs. of  $T = 0$  and **20** exs. of  $T = 1$

We calculate  $y_1 = \mathbb{1}[\frac{30}{50} > \frac{1}{2}] = 1$  and  $y_2 = 0$ , even though the correlation is actually the reverse. Now a model, which always predicts  $\hat{A}_a$  correctly, but intuitively amplifies bias by predicting  $\hat{T} = 0$  whenever  $A_1 = 1$  and predicting

$\hat{T} = 1$  whenever  $A_2 = 1$  would actually get a *negative* bias amplification score of  $\frac{0}{30} - \frac{30}{50} = -0.6$ .

BiasAmp<sub>MALS</sub> erroneously focuses on the protected group with the most examples ( $A_1 = 1$ ) rather than on the protected group that is actually correlated with  $T = 1$  ( $A_2 = 1$ ). This situation manifests when  $\min\left(\frac{1}{|A|}, P(A_a = 1)\right) < P(A_a = 1|T_t = 1) < \max\left(\frac{1}{|A|}, P(A_a = 1)\right)$ , which is more likely to arise as the the distribution of attribute  $A_a = 1$  becomes more skewed.

**Shortcoming 3: The metric entangles directions of bias amplification.** By considering only the predictions rather than the ground truth labels at test time, we are unable to distinguish between errors stemming from  $\hat{A}_a$  and those from  $\hat{T}$ . For example, looking at just the test predictions we may know that the prediction pair  $\hat{T} = 1, \hat{A}_1 = 1$  is overrepresented, but do not know whether this is due to overpredicting  $\hat{T} = 1$  on images with  $A_1 = 1$  or vice versa.

### 3.3. Experimental analysis

To verify that the above shortcomings manifest in practical settings, we revisit the analysis of Zhao et al. (2017) on the COCO (Lin et al., 2014) image dataset with two disjoint protected groups  $A_{\text{woman}}$  and  $A_{\text{man}}$ , and 66 binary target tasks,  $T_t$ , corresponding to the presence of 66 objects in the images. We directly use the released model predictions of  $\hat{A}_a$  and  $\hat{T}_t$  from Zhao et al. (2017).

First, we observe that in COCO there are about 2.5x as many men as women, leading to shortcoming 2 above. Consider the object `oven`; BiasAmp<sub>MALS</sub> calculates  $P(A_{\text{man}} = 1|T_{\text{oven}} = 1) = 0.56 > \frac{1}{2}$  and thus considers this to be correlated with men rather than women. However, computing  $P(A_{\text{man}} = 1, T_{\text{oven}} = 1) = 0.0103 < 0.0129 = P(A_{\text{man}} = 1)P(T_{\text{oven}} = 1)$  reveals that men are in fact *not* correlated with oven, and this result stems from the fact that men are overrepresented in the dataset generally. Not surprisingly, the model trained on this data associates women with ovens and underpredicts men with ovens at test time, i.e.,  $P(\hat{A}_{\text{man}} = 1|\hat{T}_{\text{oven}} = 1) - P(A_{\text{man}} = 1|T_{\text{oven}} = 1) = 0.10$ , erroneously measuring *negative* bias amplification.

In terms of directions of bias amplification, we recall that Zhao et al. (2017) discovers that ‘‘Technology oriented categories initially biased toward men such as `keyboard`... have each increased their bias toward males by over 0.100.’’ Concretely, from Eqn. 1,  $P(A_{\text{man}} = 1|T_{\text{keyboard}} = 1) = .70$  and  $P(\hat{A}_{\text{man}} = 1|\hat{T}_{\text{keyboard}} = 1) = .83$ , demonstrating an amplification of bias. However, the *direction* or cause of bias amplification remains unclear: is the presence of man in the image increasing the probability of predicting a keyboard, or vice versa? Looking more closely at the model's disentangled predictions, we see that when conditioning on

<sup>3</sup>For the rest of this subsection, for simplicity since we have only one task, we drop the subscript  $t$  so that  $T_t$ ,  $y_{at}$  and  $\Delta_{at}$  become  $T$ ,  $y_a$  and  $\Delta_a$  respectively. Further, assume the training and test datasets have the same number of examples (exs.).



the attribute,  $P(\hat{T}_{\text{keyboard}} = 1 | A_{\text{man}} = 1) = 0.0020 < 0.0032 = P(T_{\text{keyboard}} = 1 | A_{\text{man}} = 1)$ , and when conditioning on the task,  $P(\hat{A}_{\text{man}} = 1 | T_{\text{keyboard}} = 1) = 0.78 > 0.70 = P(A_{\text{man}} = 1 | T_{\text{keyboard}} = 1)$ , indicating that while keyboards are under-predicted on images with men, men are over-predicted on images with keyboards. Thus the root cause of this amplification appears to be in the gender predictor rather than the object detector. Such disentanglement allows us to properly focus algorithmic intervention efforts.

Finally, we make one last observation regarding the results of Zhao et al. (2017). The overall bias amplification is measured to be .040. However, we observe that “man” is being predicted at a higher rate (75.6%) than is actually present (71.2%). With this insight, we tune the decision threshold on the validation set such that the gender predictor is well-calibrated to be predicting the same percentage of images to have men as the dataset actually has. When we calculate  $\text{BiasAmp}_{\text{MALS}}$  on these newly thresholded predictions for the test set, we see bias amplification drop from 0.040 to 0.001 just as a result of this threshold change, outperforming even the solution proposed in Zhao et al. (2017) of corpus-level constraints, which achieved a drop to only 0.021. Fairness can be quite sensitive to the threshold chosen (Chen & Wu, 2020), so careful threshold selection should be done, rather than using a default of 0.5. When a threshold is needed in our experiments, we pick it to be well-calibrated on the validation set. In other words, we estimate the expected proportion  $p$  of positive labels from the training set and choose a threshold such that on  $N$  validation examples, the  $Np$  highest-scoring are predicted positive. Although we do not take this approach, because at deployment time it is often the case that discrete predictions are required, one could also imagine integrating bias amplification across threshold to have a threshold-agnostic measure of bias amplification, similar to what is proposed by Chen & Wu (2020).

#### 4. BiasAmp $_{\rightarrow}$ (Directional Bias Amplification)

Now we present our metric,  $\text{BiasAmp}_{\rightarrow}$ , which retains the desirable properties of  $\text{BiasAmp}_{\text{MALS}}$ , while addressing the shortcomings noted in Section 3.2. To account for the need to disentangle the two possible directions of bias amplification (shortcoming 3) the metric consists of two values:  $\text{BiasAmp}_{A \rightarrow T}$  corresponds to the amplification of bias resulting from the protected attribute influencing the task prediction, and  $\text{BiasAmp}_{T \rightarrow A}$  corresponds to the amplification of bias resulting from the task influencing the protected attribute prediction. Concretely, our directional bias amplifi-




Image Condition			
$\text{BiasAmp}_{\text{MALS}}$	.0101 $\pm$ .0040	.0141 $\pm$ .0080	.0193 $\pm$ .0055
$\text{BiasAmp}_{A \rightarrow T}$	.0009 $\pm$ .0017	-.0042 $\pm$ .0014	-.0051 $\pm$ .0019
$\text{BiasAmp}_{T \rightarrow A}$	.0267 $\pm$ .0127	.0425 $\pm$ .0089	.0491 $\pm$ .0092

Table 1. Bias amplification, as measured on the test set, changes across three image conditions: original, noisy person mask, full person mask.  $\text{BiasAmp}_{\text{MALS}}$  misleadingly makes it appear as if bias amplification *increases* as the gender cues are removed. In reality,  $A \rightarrow T$  decreases with less visual attribute cues to bias the task prediction, while it is  $T \rightarrow A$  that increases as the model relies on visual cues from the task to predict the attribute.

cation metric is:

$$\text{BiasAmp}_{\rightarrow} = \frac{1}{|A||T|} \sum_{a \in A, t \in T} y_{at} \Delta_{at} + (1 - y_{at})(-\Delta_{at})$$

$$y_{at} = \mathbb{1}[P(A_a = 1, T_t = 1) > P(A_a = 1)P(T_t = 1)]$$

$$\Delta_{at} = \begin{cases} P(\hat{T}_t = 1 | A_a = 1) & P(T_t = 1 | A_a = 1) \\ \text{if measuring } A \rightarrow T \\ P(\hat{A}_a = 1 | T_t = 1) & P(A_a = 1 | T_t = 1) \\ \text{if measuring } T \rightarrow A \end{cases} \quad (2)$$

The first line generalizes  $\text{BiasAmp}_{\text{MALS}}$  to include all attributes  $A_a$  and measure the amplification of their positive or negative correlations with task  $T_t$  (shortcoming 1). The new  $y_{at}$  identifies the direction of correlation of  $A_a$  with  $T_t$ , properly accounting for base rates (shortcoming 2). Finally,  $\Delta_{at}$  decouples the two possible directions of bias amplification (shortcoming 3). Since values may be negative, reporting the aggregated bias amplification value could obscure attribute-task pairs that exhibit strong bias amplification; thus, disaggregated results per pair can be returned.

##### 4.1. Experimental analysis

We verify that our metric successfully resolves the empirical inconsistencies of Sec. 3.2. As expected,  $\text{BiasAmp}_{A \rightarrow T}$  is positive at .1778 in shortcoming 1 and .3333 in 2;  $\text{BiasAmp}_{T \rightarrow A}$  is 0 in both. We further introduce a scenario for empirically validating the decoupling aspect of our metric. We use a baseline amplification removal idea of applying segmentation masks (noisy or full) over the people in an image to mitigate bias stemming from human attributes (Wang et al., 2019). We train on the COCO classification task, with the same 66 objects from Zhao et al. (2017), a VGG16 (Simonyan & Zisserman, 2014) model pre-trained on ImageNet (Russakovsky et al., 2015) to predict objects and gender, with a Binary Cross Entropy Loss over all outputs, and measure  $\text{BiasAmp}_{T \rightarrow A}$  and  $\text{BiasAmp}_{A \rightarrow T}$ ;

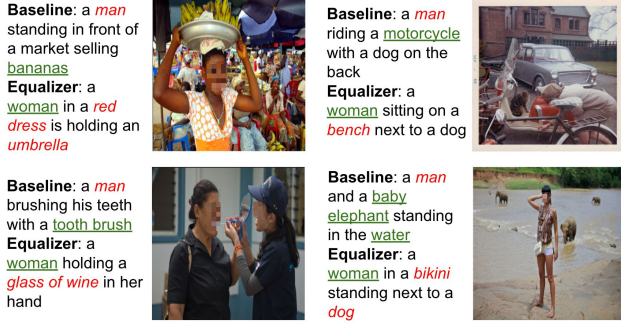


Figure 3. Illustrative captions from the Equalizer model (Hendricks et al., 2018), which in these captions decreases  $T \rightarrow A$  bias amplification from the Baseline, but inadvertently increases  $A \rightarrow T$ . Green underlined words are correct, and red italicized words are incorrect. In the images, the Equalizer improves on the Baseline for the gendered word, but introduces biased errors in the captions.

we report 95% confidence intervals for 5 runs of each scenario. In Tbl. 1 we see that, misleadingly, BiasAmp<sub>MALS</sub> reports increased amplification as gender cues are removed. However what is actually happening is, as expected, that as less of the person is visible,  $A \rightarrow T$  decreases because there are less human attribute visual cues to bias the task prediction. It is  $T \rightarrow A$  that increases because the model must lean into task biases to predict the person’s attribute. However, we can also see from the overlapping confidence intervals that the difference between noisy and full masks does not appear to be particularly robust; we continue a discussion of this phenomenon in Sec. 5.2.<sup>4</sup>

## 5. Analysis and Discussion

We now discuss some of the normative issues surrounding bias amplification: in Sec. 5.1 with the existence of  $T \rightarrow A$  bias amplification, which implies the prediction of sensitive attributes; in Sec. 5.2 about the need for confidence intervals to make robust conclusions; and in Sec. 5.3 about scenarios in which the original formulation of bias amplification as a desire to match base correlations may not be the intention.

### 5.1. Considerations of $T \rightarrow A$ Bias Amplification

If we think more deeply about these bias amplifications, we might come to a normative conclusion that  $T \rightarrow A$ , which measures sensitive attribute predictions conditioned on the tasks, should not exist in the first place. There are

<sup>4</sup>This simultaneously serves as inspiration for an intervention approach to mitigating bias amplification. In Appendix A.4 we provide a more granular analysis of this experiment, and how it can help to inform mitigation. Further mitigation techniques are outside of our scope, but we look to works like Singh et al. (2020); Wang et al. (2019); Agarwal et al. (2020).

very few situations in which predicting sensitive attributes makes sense (Scheuerman et al., 2020; Larson, 2017), so we should carefully consider if this is strictly necessary for target applications. For the image domains discussed, by simply removing the notion of predicting gender, we trivially remove all  $T \rightarrow A$  bias amplification. In a similar vein, there has been great work done on reducing gender bias in image captions (Hendricks et al., 2018; Tang et al., 2020), but it is often focused on targeting  $T \rightarrow A$  rather than  $A \rightarrow T$  amplification. When disentangling the directions of bias, we find that the Equalizer model (Hendricks et al., 2018), which was trained with the intention of increasing the quality of gender-specific words in captions, inadvertently increases  $A \rightarrow T$  bias amplification for certain tasks. We treat gender as the attribute and the objects as different tasks. In Fig. 3 we see examples where the content of the Equalizer’s caption exhibits bias coming from the person’s attribute. Even though the Equalizer model reduces  $T \rightarrow A$  bias amplification in these images, it inadvertently increases  $A \rightarrow T$ . It is important to disentangle the two directions of bias and notice that while one direction is becoming more fair, another is actually becoming more biased. Although this may not always be the case, depending on the downstream application (Bennett et al., 2021), perhaps we could consider simply replacing all instances of gendered words like “man” and “woman” in the captions with “person” to trivially eliminate  $T \rightarrow A$ , and focus on  $A \rightarrow T$  bias amplification. Specifically when gender is the sensitive attribute, Keyes (2018) thoroughly explains how we should carefully think about why we might implement Automatic Gender Recognition (AGR), and avoid doing so.

On the other hand, sensitive attribute labels, ideally from self-disclosure, can be very useful. For example, these labels are necessary to measure  $A \rightarrow T$  amplification, which is important to discover, as we do not want our prediction task to be biased for or against people with certain attributes.

### 5.2. Variance in Estimator Bias

Evaluation metrics, ours included, are specific to each model on each dataset. Under common loss functions such as cross entropy loss, some evaluation metrics like average precision are not very sensitive to random seed. However, bias amplification, along with other fairness metrics like FPR difference, often fluctuates greatly across runs. Because the loss functions that machine learning practitioners tend to default to using are proxies for accuracy, it makes sense that various local minima, while equal in accuracy, are not necessarily equal for other measurements. The phenomena of differences between equally predictive models has been termed the Rashomon Effect (Breiman, 2001), or predictive multiplicity (Marx et al., 2020).

Thus, like previous work (Fisher et al., 2019), we urge

transparency, and advocate for the inclusion of confidence intervals. To illustrate the need for this, we look at the facial image domain of CelebA (Liu et al., 2015), defining two different scenarios of the classification of `big_nose` or `young` as our task, and treating the gender labels as our attribute. Note that we do not classify gender, for reasons raised in Sec. 5.1, so we only measure  $A \rightarrow T$  amplification. For these tasks, women are correlated with no big nose and being young, and men with big nose and not being young. We examine two different scenarios, one where our independent variable is model architecture, and another where it is the ratio between number of images of the majority groups (e.g., young women and not young men) and minority groups (e.g., not young women and young men). By looking at the confidence intervals, we can determine which condition allows us to draw reliable conclusions about the impact of that variable on bias amplification.

For model architecture, we train 3 models pretrained on ImageNet (Russakovsky et al., 2015) across 5 runs: ResNet18 (He et al., 2016), AlexNet (Krizhevsky et al., 2012), and VGG16 (Simonyan & Zisserman, 2014). Training details are in Appendix A.2. In Fig. 4 we see from the confidence intervals that while model architecture does not result in differing enough of bias amplification to conclude anything about the relative fairness of these models, across-ratio differences are significant enough to draw conclusions about the impact of this ratio on bias amplification. We encourage researchers to include confidence intervals so that findings are more robust to random fluctuations. Concurrent work covers this multiplicity phenomenon in detail (D’Amour et al., 2020), and calls for more application-specific specifications that would constrain the model space. However, that may not always be feasible, so for now our proposal of error bars is more general and immediately implementable. In a survey of recently published fairness papers from prominent machine learning conferences, we found that 25 of 48 (52%) reported results of a fairness metric without error bars (details in Appendix A.2. Even if the model itself is deterministic, error bars could be generated through bootstrapping (Efron, 1992) to account for the fact that the test set itself is but a sample of the population, or varying the train-test splits (Friedler et al., 2019).

### 5.3. Limitations of Bias Amplification

An implicit assumption that motivates bias amplification metrics, including ours, is that the ground truth exists and is known. Further, a perfectly accurate model can be considered perfectly fair, despite the presence of task-attribute correlations in the training data. This allows us to treat the disparity between the correlations in the input vs correlations in the output as a fairness measure.

It follows that bias amplification would *not* be a good way to

measure fairness when the ground truth is either unknown or does not correspond to desired classification. In this section, we discuss two types of applications where bias amplification should not necessarily be used out-of-the-box as a fairness metric.

**Sentence completion: no ground truth.** Consider the fill-in-the-blank NLP task, where there is no ground truth for how to fill in a sentence. Given “The [blank] went on a walk”, a variety of words could be suitable. Therefore, to measure bias amplification in this setting, we need to subjectively set the base correlations, i.e.,  $P(T_t = 1|A_a = 1), P(A_a = 1|T_t = 1)$ .

To see the effect of adjusting base correlations, we test the bias amplification between occupations and gender pronouns, conditioning on the pronoun and filling in the occupation and vice versa. In Tbl. 2, we report our measured bias amplification results on the FitBERT (Fill in the blanks BERT) (Havens & Stal, 2019; Devlin et al., 2019) model using various sources as our base correlation of bias from which amplification is measured. The same outputs from the model are used for each set of pronouns, and the independent variable we manipulate is the source of base correlations: 1) equality amongst the pronouns, using two pronouns (he/she) 2) equality amongst the pronouns, using three pronouns (he/she/they) 3) co-occurrence counts from English Wikipedia (one of the datasets BERT was trained on), and 4) WinoBias (Zhao et al., 2018) with additional information supplemented from the 2016 U.S. Labor Force Statistics data. Additional details are in Appendix A.3.

We find that relative to U.S. Labor Force data on these particular occupations, FitBERT actually exhibits no bias amplification. Yet it would be simplistic to conclude that FitBERT presents no fairness concerns with respect to gender and occupation. For one, it is evident from Fig. 5 that there is an overall bias towards “he” (this translates to a bias amplification for some occupations and a bias reduction for others; the effects roughly cancel out in our bias amplification metric when aggregated). More importantly, whether U.S. labor statistics are the right source of base correlations depends on the specific application of the model and the cultural context in which it is deployed. This is clear when noticing that the measured  $\text{BiasAmp}_{T \rightarrow A}$  is much stronger when the gender distribution is expected to be uniform, instead of gender-biased Labor statistics.

**Risk prediction: future outcomes unknown.** Next, we examine the criminal risk prediction setting. A common statistical task in this setting is predicting the likelihood a defendant will commit a crime if released pending trial. This setting has two important differences compared to computer vision detection tasks: 1) The training labels typically come from arrest records and suffer from problems like historical and selection bias (Suresh & Guttag, 2019; Olteanu et al.,

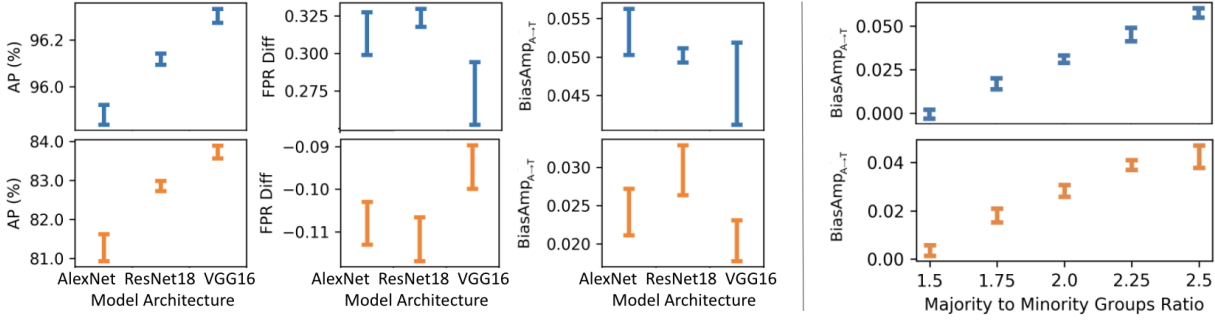


Figure 4. We investigate the consistency of various metrics by looking at 95% confidence intervals as we manipulate two independent variables: model architecture (left three graphs), and majority to minority groups ratio (right graph). The top row (blue) is for the attribute of *big nose*, and bottom row (orange) is for *young*. For model architecture, across 5 runs, the accuracy measure of average precision retains a consistent ranking across models, but two different fairness measures (FPR difference and  $A \rightarrow T$  bias amplification) have overlapping intervals. This does not allow us to draw conclusions about the differing fairness of these models. However, across-ratio differences in bias amplification are significant enough to allow us to draw conclusions about the differing levels of fairness.




Figure 5. Each point represents an occupation’s probability at being associated with the pronoun for a man. FitBERT perpetuates gender-occupation biases seen in the U.S. Labor Force, and additionally over-favors the pronoun for men.

2019; Green, 2020), and 2) the task is to predict future events and thus the outcome is not knowable at prediction time. Further, the risk of recidivism is not a static, immutable trait of a person. Given the input features that are used to represent individuals, one could imagine an individual with a set of features who does recidivate, and one who does not. In contrast, for a task like image classification, two instances with the same pixel values will always have the same labels (if the ground truth labels are accurate).

As a result of these setting differences, risk prediction tools may be considered unfair even if they exhibit no bias amplification. Indeed, one might argue that a model that shows no bias amplification is necessarily unfair as it perpetuates past biases reflected in the training data. Further, modeling risk as immutable misses the opportunity for intervention to change the risk (Barabas et al., 2018). Thus, matching the training correlations should not be the intended goal (Wick et al., 2019; Hebert-Johnson et al., 2018).

To make this more concrete, in Fig. 6 we show the metrics of  $BiasAmp_{A \rightarrow T}$  and False Positive Rate (FPR) disparity

Base Correlation Source (# pronouns)	$BiasAmp_{T \rightarrow A}$	$BiasAmp_{A \rightarrow T}$
Uniform (2)	$.1368 \pm .0226$	$.0084 \pm .0054$
Uniform (3)	$.0914 \pm .0151$	$.0056 \pm .0036$
Wikipedia (2)	$.0372 \pm .0307$	$-.0002 \pm .0043$
2016 U.S. Labor Force (WinoBias) (2)	$-.1254 \pm .0026$	$-.0089 \pm .0054$

Table 2.  $BiasAmp_{\rightarrow}$  for different base correlation sources. The value-laden choice of base correlation source depends on the downstream application.

measured on COMPAS predictions (Angwin et al., 2016), only looking at two racial groups, for various values of the risk threshold. A false positive occurs when a defendant classified as high risk but does not recidivate; FPR disparity has been interpreted as measuring how unequally different groups suffer the costs of the model’s errors (Hardt et al., 2016). The figure shows that bias amplification is close to 0 for almost all thresholds. This is no surprise since the model was designed to be calibrated by group (Flores et al., 2016). However, for all realistic values of the threshold, there is a large FPR disparity. Thus, risk prediction is a setting where a lack of bias amplification should not be used to conclude that a model is fair.

Like any fairness metric, ours captures only one perspective, which is that of not amplifying already present biases. It does not require a correction for these biases. Settings that bias amplification are more suited for include those with a known truth in the labels, where matching them would be desired. For example, applicable contexts include certain social media bot detection tasks where the sensitive attribute is the region of origin, as bot detection methods may be biased against names from certain areas. More broadly, it






Figure 6. COMPAS risk predictions exhibit FPR disparities, but little bias amplification. Bias amplification measures only whether the model matches the (biased) training data, not the bias of the overall system.

is crucial that we pick fairness metrics thoughtfully when deciding how to evaluate a model.

## 6. Conclusion

In this paper, we take a deep dive into the measure of bias amplification. We introduce a new metric,  $\text{BiasAmp}_{A \rightarrow T}$ , that disentangles the directions of bias to provide more actionable insights when diagnosing models. Additionally, we analyze and discuss normative considerations to encourage exercising care when determining which fairness metrics are applicable, and what assumptions they are encoding. The mission of this paper is not to tout bias amplification as the optimal fairness metric, but rather to give a comprehensive and critical study as to how it should be measured.

## Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant No. 1763642. We thank Sunnie S. Y. Kim, Karthik Narasimhan, Vikram Ramaswamy, Brandon Stewart, and Felix Yu for feedback. We especially thank Arvind Narayanan for significant comments and advice. We also thank the authors of Men Also Like Shopping (Zhao et al., 2017) and Women Also Snowboard (Hendricks et al., 2018) for uploading their model outputs and code online in a way that made it easily reproducible, and for being prompt and helpful in response to clarifications.

## References

- Agarwal, V., Shetty, R., and Fritz, M. Towards causal VQA: Revealing and reducing spurious correlations by invariant and covariant semantic editing. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- Angwin, J., Larson, J., Mattu, S., and Kirchner, L. Machine bias. *Propublica*, 2016.
- Barabas, C., Dinakar, K., Ito, J., Virza, M., and Zittrain, J. Interventions over predictions: Reframing the ethical debate for actuarial risk assessment. *Fairness, Accountability and Transparency in Machine Learning*, 2018.
- Bearman, S., Korobov, N., and Thorne, A. The fabric of internalized sexism. *Journal of Integrated Social Sciences 1(1)*: 10-47, 2009.
- Bennett, C. L., Gleason, C., Scheuerman, M. K., Bigham, J. P., Guo, A., and To, A. “it’s complicated”: Negotiating accessibility and (mis)representation in image descriptions of race, gender, and disability. *Conference on Human Factors in Computing Systems (CHI)*, 2021.
- Berk, R., Heidari, H., Jabbari, S., Kearns, M., and Roth, A. Fairness in criminal justice risk assessments: The state of the art. *Sociological Methods and Research*, 2017.
- Bhattacharya, J. and Vogt, W. B. Do instrumental variables belong in propensity scores? *NBER Technical Working Papers 0343*, National Bureau of Economic Research, Inc., 2007.
- Blodgett, S. L., Barocas, S., III, H. D., and Wallach, H. Language (technology) is power: A critical survey of “bias” in nlp. *Association for Computational Linguistics (ACL)*, 2020.
- Bolukbasi, T., Chang, K.-W., Zou, J., Saligrama, V., and Kalai, A. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. *Advances in Neural Information Processing Systems (NeurIPS)*, 2016.
- Breiman, L. Statistical modeling: The two cultures. *Statistical Science*, 16:199–231, 2001.
- Buolamwini, J. and Gebru, T. Gender shades: Intersectional accuracy disparities in commercial gender classification. *Proceedings of Machine Learning Research*, 81, 2018.
- Caliskan, A., Bryson, J. J., and Narayanan, A. Semantics derived automatically from language corpora contain human-like biases. *Science*, 2017.
- Chen, M. and Wu, M. Towards threshold invariant fair classification. *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2020.
- Choi, K., Grover, A., Singh, T., Shu, R., and Ermon, S. Fair generative modeling via weak supervision. *arXiv:1910.12008*, 2020.
- Chouldechova, A. Fair prediction with disparate impact: A study of bias in recidivism prediction instrument. *Big Data*, 2016.

- D’Amour, A., Heller, K., Moldovan, D., Adlam, B., Alipanahi, B., Beutel, A., Chen, C., Deaton, J., Eisenstein, J., Hoffman, M. D., Hormozdiari, F., Houlisby, N., Hou, S., Jerfel, G., Karthikesalingam, A., Lucic, M., Ma, Y., McLean, C., Mincu, D., Mitani, A., Montanari, A., Nado, Z., Natarajan, V., Nielson, C., Osborne, T. F., Raman, R., Ramasamy, K., Sayres, R., Schrouff, J., Seneviratne, M., Sequeira, S., Suresh, H., Veitch, V., Vladymyrov, M., Wang, X., Webster, K., Yadlowsky, S., Yun, T., Zhai, X., and Sculley, D. Underspecification presents challenges for credibility in modern machine learning. *arXiv:2011.03395*, 2020.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. BERT: Pre-training of deep bidirectional transformers for language understanding. *Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, 2019.
- Dong, J. and Rudin, C. Variable importance clouds: A way to explore variable importance for the set of good models. *arXiv:1901.03209*, 2019.
- Dwork, C., Hardt, M., Pitassi, T., Reingold, O., and Zemel, R. Fairness through awareness. *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, 2012.
- Efron, B. Bootstrap methods: another look at the jackknife. In *Breakthroughs in statistics*, pp. 569–593. Springer, 1992.
- Fisher, A., Rudin, C., and Dominici, F. 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*, 20, 2019.
- Flores, A. W., Bechtel, K., and Lowenkamp, C. T. False positives, false negatives, and false analyses: A rejoinder to ”machine bias: There’s software used across the country to predict future criminals. and it’s biased against blacks.”. *Federal Probation Journal*, 80, 2016.
- Foulds, J., Islam, R., Keya, K. N., and Pan, S. An intersectional definition of fairness. *arXiv:1807.08362*, 2018.
- Friedler, S. A., Scheidegger, C., Venkatasubramanian, S., Choudhary, S., Hamilton, E. P., and Roth, D. A comparative study of fairness-enhancing interventions in machine learning. *Conference on Fairness, Accountability, and Transparency (FAccT)*, 2019.
- Green, B. The false promise of risk assessments: Epistemic reform and the limits of fairness. *ACM Conference on Fairness, Accountability, and Transparency (ACM FAccT)*, 2020.
- Grgic-Hlaca, N., Zafar, M. B., Gummadi, K. P., and Weller, A. The case for process fairness in learning: Feature selection for fair decision making. *NeurIPS Symposium on Machine Learning and the Law*, 2016.
- Hancox-Li, L. Robustness in machine learning explanations: Does it matter? *Conference on Fairness, Accountability, and Transparency (FAccT)*, 2020.
- Hardt, M., Price, E., and Srebro, N. Equality of opportunity in supervised learning. *arXiv:1610.02413*, 2016.
- Havens, S. and Stal, A. Use BERT to fill in the blanks, 2019. URL <https://github.com/Qordobacode/fitbert>.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. *European Conference on Computer Vision (ECCV)*, 2016.
- Hebert-Johnson, U., Kim, M. P., Reingold, O., and Rothblum, G. N. Multicalibration: Calibration for the (computationally-identifiable) masses. *International Conference on Machine Learning (ICML)*, 2018.
- Hendricks, L. A., Burns, K., Saenko, K., Darrell, T., and Rohrbach, A. Women also snowboard: Overcoming bias in captioning models. *European Conference on Computer Vision (ECCV)*, 2018.
- Jain, N., Olmo, A., Sengupta, S., Manikonda, L., and Kambhampati, S. Imperfect imagination: Implications of gans exacerbating biases on facial data augmentation and snapchat selfie lenses. *arXiv:2001.09528*, 2020.
- Jia, S., Meng, T., Zhao, J., and Chang, K.-W. Mitigating gender bias amplification in distribution by posterior regularization. *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2020.
- Keyes, O. The misgendering machines: Trans/HCI implications of automatic gender recognition. *Proceedings of the ACM on Human-Computer Interaction*, 2018.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. ImageNet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 1097–1105, 2012.
- Kuang, S. and Davison, B. D. Semantic and context-aware linguistic model for bias detection. *Proc. of the Natural Language Processing meets Journalism IJCAI-16 Workshop*, 2016.
- Kusner, M. J., Loftus, J. R., Russell, C., and Silva, R. Counterfactual fairness. *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.

- Larson, B. N. Gender as a variable in natural-language processing: Ethical considerations. *Proceedings of the First ACL Workshop on Ethics in Natural Language Processing*, 2017.
- Leino, K., Black, E., Fredrikson, M., Sen, S., and Datta, A. Feature-wise bias amplification. *International Conference on Learning Representations (ICLR)*, 2019.
- Lin, T.-Y., Maire, M., Belongie, S., Bourdev, L., Girshick, R., Hays, J., Perona, P., Ramanan, D., Zitnick, C. L., and Dollar, P. Microsoft COCO: Common objects in context. *European Conference on Computer Vision (ECCV)*, 2014.
- Liu, Z., Luo, P., Wang, X., and Tang, X. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
- Lu, K., Mardziel, P., Wu, F., Amancharla, P., and Datta, A. Gender bias in neural natural language processing. *arXiv:1807.11714*, 2019.
- Marx, C. T., du Pin Calmon, F., and Ustun, B. Predictive multiplicity in classification. *arXiv:1909.06677*, 2020.
- May, C., Wang, A., Bordia, S., Bowman, S. R., and Rudinger, R. On measuring social biases in sentence encoders. *Annual Conference of the North American Chapter of the Association for Computational Linguistics (NACCL)*, 2019.
- Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., and Galstyan, A. A survey on bias and fairness in machine learning. *arXiv:1908.09635*, 2019.
- Middleton, J. A., Scott, M. A., Diakow, R., and Hill, J. L. Bias amplification and bias unmasking. *Political Analysis*, 3:307–323, 2016.
- Olteanu, A., Castillo, C., Diaz, F., and Kiciman, E. Social data: Biases, methodological pitfalls, and ethical boundaries. *Frontiers in Big Data*, 2019.
- Pawelczyk, M., Broelemann, K., and Kasneci, G. On counterfactual explanations under predictive multiplicity. *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2020.
- Pearl, J. On a class of bias-amplifying variables that endanger effect estimates. *Uncertainty in Artificial Intelligence*, 2010.
- Pearl, J. Invited commentary: Understanding bias amplification. *American Journal of Epidemiology*, 174, 2011.
- Qian, Y., Muaz, U., Zhang, B., and Hyun, J. W. Reducing gender bias in word-level language models with a gender-equalizing loss function. *ACL-SRW*, 2019.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A. C., and Fei-Fei, L. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. doi: 10.1007/s11263-015-0816-y.
- Scheuerman, M. K., Wade, K., Lustig, C., and Brubaker, J. R. How we’ve taught algorithms to see identity: Constructing race and gender in image databases for facial analysis. *Proceedings of the ACM on Human-Computer Interaction*, 2020.
- Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv:1409.1556*, 2014.
- Singh, K. K., Mahajan, D., Grauman, K., Lee, Y. J., Feiszli, M., and Ghadiyaram, D. Don’t judge an object by its context: Learning to overcome contextual bias. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- Stock, P. and Cisse, M. ConvNets and ImageNet beyond accuracy: Understanding mistakes and uncovering biases. *European Conference on Computer Vision (ECCV)*, 2018.
- Suresh, H. and Gutttag, J. V. A framework for understanding unintended consequences of machine learning. *arXiv:1901.10002*, 2019.
- Tang, R., Du, M., Li, Y., Liu, Z., and Hu, X. Mitigating gender bias in captioning systems. *arXiv:2006.08315*, 2020.
- Wang, T., Zhao, J., Yatskar, M., Chang, K.-W., and Ordonez, V. Balanced datasets are not enough: Estimating and mitigating gender bias in deep image representations. *International Conference on Computer Vision (ICCV)*, 2019.
- Wick, M., Panda, S., and Tristan, J.-B. Unlocking fairness: a trade-off revisited. *Conference on Neural Information Processing Systems (NeurIPS)*, 2019.
- Wooldridge, J. M. Should instrumental variables be used as matching variables? *Research in Economics*, 70:232–237, 2016.
- Zhao, J., Wang, T., Yatskar, M., Ordonez, V., and Chang, K.-W. Men also like shopping: Reducing gender bias amplification using corpus-level constraints. *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2017.
- Zhao, J., Wang, T., Yatskar, M., Ordonez, V., and Chang, K.-W. Gender bias in coreference resolution: Evaluation and debiasing methods. *North American Chapter of*

*the Association for Computational Linguistics (NAACL),*  
2018.