A TECHNICAL AND NORMATIVE INVESTIGATION OF SOCIAL BIAS AMPLIFICATION

Anonymous authors
Paper under double-blind review

ABSTRACT

The conversation around the fairness of machine learning models is growing and evolving. In this work, we focus on one particular component, the issue of *bias amplification*: the tendency of models trained from data containing social biases to further amplify these biases. By definition, this problem stems from the algorithm, and cannot be attributed to the dataset. We make three contributions regarding its measurement. First, building off of Zhao et al. (2017), we introduce and analyze a new, decoupled metric for measuring bias amplification, BiasAmp$_{\rightarrow}$, which possesses a number of attractive properties, including the ability to pinpoint the cause of bias amplification. Second, we demonstrate the lack of consistency in values reported by fairness metrics across models that are equally accurate, and encourage the use of confidence intervals when reporting such fairness measures. Finally, we consider what bias amplification means in the context of domains where labels either don't exist at test time, or correspond to uncertain future events. We provide a deeply interrogative look at the technical measurement of bias amplification, guided by our normative ideas of what we want it to encompass.

1 INTRODUCTION

The machine learning community is becoming increasingly cognizant of problems surrounding fairness and bias, and correspondingly a plethora of new algorithms and metrics are being proposed (see e.g., Mehrabi et al. (2019) for a review). The gatekeepers checking the systems to be deployed 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 happens when a model exacerbates biases from the training data at test time. It is purely the result of the algorithm, and unlike other forms of bias, cannot be attributed to the dataset.

To this end, we propose a new way of measuring bias amplification, BiasAmp$_{\rightarrow}$, that builds off a prior metric from Men Also Like Shopping (Zhao et al., 2017), that we will call BiasAmp$_{\text{MALS}}$. 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) take into account the base rates that people of each attribute appear, 2) disentangle the directions of amplification, 3) be agnostic to the classification threshold picked, and 4) generalize beyond binary attributes. Concretely, 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. 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 $T \rightarrow A$. However, we also consider and discuss whether predicting protected attributes is necessary in the first place; by not doing so, we can trivially remove $T \rightarrow A$ amplification.

BiasAmp$_{\rightarrow}$, along with other fairness metrics, suffers from the Rashomon Effect (Breiman, 2001), or multiplicity of good models. In deep neural networks, random seeds have relatively little impact on accuracy; however, that is not the case for fairness, which is more brittle to randomness. We advocate for transparency, and the inclusion of error bars as a norm when reporting fairness measures.

A trait of bias amplification is that it is not at odds with accuracy, unlike many other fairness metrics, because the goal of not amplifying biases and matching task-attribute correlations is aligned with that of accurate predictions. While this is desired in object detection where systems should perfectly predict the labels, we will explore the nuances of what this means in situations where the validity

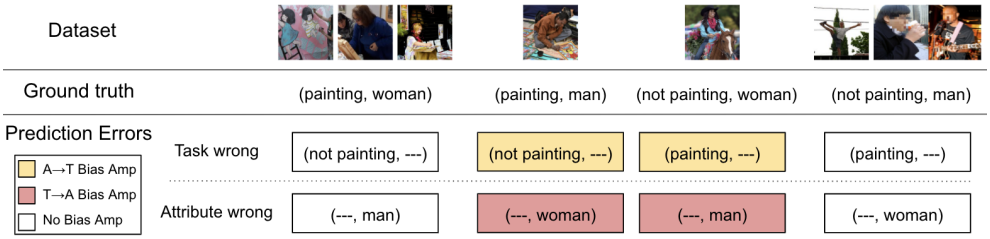

Figure 1: Consider an image recognition dataset where the goal is to classify the task, T, as painting or not painting, and the attribute, A, as woman or man. Note that in this dataset women are correlated with painting, and men with not painting. In this work we are particularly concerned with errors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

of the labels, and thus task-attribute correlations themselves, are up for debate. For example, in the risk prediction task which assesses someone's likelihood of recidivism, the label represents whether someone with a set of input features ended up recidivating, but is not a steadfast indicator of what another person with the same input features will do. Here, we would not want to replicate the task-attribute correlations at test time, and it is important to keep this in mind when deciding what fairness metrics to apply. The notion of amplification also 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 more of a negative connotation than overclassifying men as cooking.[1] Distinguishing between which errors are more harmful than others is a pattern that can often be lifted from the training data.

To ground our work, we first distinguish what bias amplification captures that standard fairness metrics cannot, and then distinguish $\text{BiasAmp}_{\rightarrow}$ from $\text{BiasAmp}_{\text{MALS}}$. Our main contributions are: 1) proposing and thoroughly analyzing a new way to measure bias amplification that allows us to better diagnose where a model goes wrong, 2) calling for the inclusion of error bars on measures of fairness evaluation due to the multiplicity of good models that vary in measured bias, and 3) delineating what bias amplification means when applied to different contexts where the ground-truth of the labels is not a given, and encouraging thoughtfulness with each application.

2 Related Work

Fairness Measurements. Fairness is nebulous and context-dependent, and approaches to quantifying it (Verma & Rubin, 2018; Buolamwini & Gebru, 2018) 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, which is a type of group fairness where correlations are amplified.

Bias Amplification. Bias amplification has been measured by looking at binary classifications (Leino et al., 2019), GANs (Jain et al., 2020; Choi et al., 2020), and correlations (Zhao et al., 2017). Wang et al. (2019) measures this using dataset leakage and model leakage. The difference between these values is the level of bias amplification, but this is not a fair comparison because the attribute classifier gets discrete labels for the former but continuous model outputs for the latter. Jia et al. (2020) looks at output distributions like we do, but with a different formulation.

The Word Embedding Association Test (WEAT) (Caliskan et al., 2017) measures bias amplification in de-contextualized word embeddings, looking at correlations but not causations (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

[1]We use the terms man and woman to refer to binarized socially-perceived gender expression, recognizing these labels are not inclusive, and in vision datasets are often assigned by annotators rather than self-disclosed.

context (Lu et al., 2019; Kuang & Davison, 2016). We use these models to measure the directional aspect of these amplifications, as well as to situate them in the broader world of bias amplification.

Directionality. Directionality of amplification has been observed in computer vision (Stock & Cisse, 2018) and language (Qian et al., 2019). It has also been studied with causality (Bhattacharya & Vogt, 2007; Wooldridge, 2016; Pearl, 2010; Middleton et al., 2016; Steiner & Yongnam, 2016). We take an approach that is more more rigorous than anecdotal, and empirical than theoretical.

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 discoveries and investigate how fairness measurements also differ between equally "good" models.

3 METRIC

In Sec. 3.1 we begin by presenting existing fairness metrics and showing how bias amplification can distinguish errors resulting from under- and overclassification in a way that others cannot; in Sec. 3.2 we discuss some shortcomings of $\text{BiasAmp}_{\text{MALS}}$; in Sec. 3.3 we present our metric, $\text{BiasAmp}_{\rightarrow}$, which addresses the shortcomings noted in the previous section, with experiments in Sec. 3.4; we conclude in Sec. 3.5 with a discussion about predicting sensitive attributes.

3.1 EXISTING FAIRNESS METRICS

We begin with a review of existing fairness metrics in a concrete classification setting. We consider again the example from Fig. 1, where on this dataset women (a_1) are correlated with painting (t_1), and men (a_2) with not painting (t_2), such that there are N images each of (a_1, t_1) and (a_2, t_2) but only $N/2$ images of (a_1, t_2) and (a_2, t_1). A classifier trained to recognize painting on this data is likely to learn this association and over-predict painting on images of women and under-predict painting on images of men; however, algorithmic interventions may counteract this effect and in fact result in the opposite behavior.

Fig. 2 shows the behavior of fairness metrics under varying amounts of learned amplification of the correlation. The four fairness metrics are: *False Positive Rate (FPR) and True Positive Rate (TPR) difference:* the difference in false positive (true positive) rate of predicting the label t_1 on images of a_2 versus on images of a_1 (Chouldechova, 2016; Hardt et al., 2016), *accuracy difference:* difference between the overall task prediction accuracy on images of a_2 versus on images of a_1 (Berk et al., 2017), and *mean accuracy across subgroups:* mean task prediction accuracy across the four image subgroups ($(a_1, t_1), (a_2, t_1), (a_1, t_2)$, and (a_2, t_2)) (Buolamwini & Gebru, 2018). However, these metrics are not designed to account for the training correlations, and are unable to distinguish between cases of increased or decreased learned correlations, as seen in Fig. 2.

Zhao et al. (2017) introduced an alternative to these that explicitly captures the notion of bias amplification. Concretely, they consider $P(A = a|T = t)$ of the training data as the fraction of times a protected attribute $a \in \mathcal{A}$ appears on images corresponding to task $t \in \mathcal{T}$. They then compare this with the test-time predictions made by the model, $P(\hat{A} = a|\hat{T} = t)$, or the number of times attribute a is predicted on images where the task is predicted as t, which allows them to measure bias amplification in the absence of any additional annotations on the hold-out set. Note that in this formulation they are assuming that the model is making predictions for both the task and the attribute. The full bias amplification metric (reformulated in our terms), is computed as:

$$\text{BiasAmp}_{\text{MALS}} = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \sum_{a \in \mathcal{A}} \underbrace{\mathbb{1}\left[P(A = a|T = t) > \frac{1}{|\mathcal{A}|}\right]}_{y(t,a)} \underbrace{\left(P(\hat{A} = a|\hat{T} = t) - P(A = a|T = t)\right)}_{\Delta(t,a)} \tag{1}$$

Fig. 2 empirically demonstrates that this metric is able to capture the level of increasing bias amplification. (For consistency in comparison with prior metrics, we assume the model always correctly predicts the protected attribute A.) However, as we discuss in the next section, there are some properties of bias amplification that this metric is not able to capture: for example, it does not distinguish

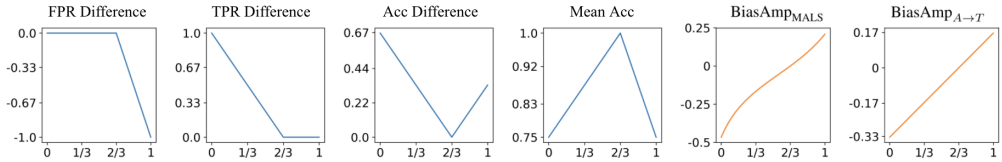

Figure 2: Different fairness metrics vary in how they respond to model errors. In our image dataset (Fig. 1) of predicting someone who is a woman or man to be painting or not, we consider a classifier that always predicts the task correctly for men, but varies for women. The x-axes of the graphs correspond to the fraction of women predicted to be painting, where the ground-truth is $2/3$, and the model does not predict false positives before this point. The first two metrics, FPR and TPR difference, only capture one of under- or overclassification. The next two metrics are symmetric around $2/3$, being unable to differentiate under- or overclassification. Both bias amplification metrics are able to distinguish between under- and overclassification.

between errors in predicting the protected attribute versus errors in predicting the task. Thus we introduce a new metric (last graph of Fig. 2) which maintains the desirable properties from Zhao et al. (2017) while including a number of innovations.

3.2 Shortcomings of BiasAmp$_{\text{MALS}}$

Despite the advantages we just documented about BiasAmp$_{\text{MALS}}$ (Eqn. 1) and its ability to distinguish under- and overclassifications of training correlations, this metric also suffers from a number of shortcomings. To ground our discussion, we will work directly with the model outputs released by Zhao et al. (2017) from their Conditional Random Field (CRF) model on COCO (Lin et al., 2014), which has predictions for gender and objects detected for each image.

3.2.1 Base rates

The first shortcoming of BiasAmp$_{\text{MALS}}$ is the fact that the metric does not take into account the base rates of each attribute. Concretely, when determining in $y(t, a)$ of Eqn. 1 whether the attribute a is correlated with the task t, $P(A = a|T = t)$ is compared to $\frac{1}{|\mathcal{A}|}$. However, this assumes that all a's within $\mathcal{A}$ are evenly distributed, which may not be the case. For example, in COCO there are about 2.5x as many men as women, so it would appear that most objects positively correlate with men simply by nature of there being an overrepresentation of men. Consider the object `oven`; BiasAmp$_{\text{MALS}}$ calculates $P(A = \texttt{man}|T = \texttt{oven}) = 0.56 > \frac{1}{2}$ and thus considers this object to be correlated with men rather than women. However, computing $P(A = \texttt{man}, T = \texttt{oven}) = 0.0031 < 0.0038 = P(A = \texttt{man})P(T = \texttt{oven})$ reveals that men are in fact *not* correlated with oven, and the seeming overrepresentation comes from the fact that men are overrepresented in the dataset more generally. Not surprisingly, the model trained on this data learns to associate women with ovens and underpredicts men with ovens at test time, i.e., $P(\hat{A} = \texttt{man}|\hat{T} = \texttt{oven}) - P(A = \texttt{man}|T = \texttt{oven}) = -0.10$. BiasAmp$_{\text{MALS}}$ erroneously counts this as *inverse* bias amplification.

3.2.2 Entangling directions

Another shortcoming we observe is the inability to distinguish between different *types* of bias amplification. 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 $P(A = \texttt{man}|T = \texttt{keyboard}) = .70$ and $P(\hat{A} = \texttt{man}|\hat{T} = \texttt{keyboard}) = .83$, 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:

$$P(\hat{T} = \texttt{keyboard}|A = \texttt{man}) = 0.0020 < 0.0032 = P(T = \texttt{keyboard}|A = \texttt{man}) \quad (2)$$

$$P(\hat{A} = \texttt{man}|T = \texttt{keyboard}) = 0.78 > 0.70 = P(A = \texttt{man}|T = \texttt{keyboard}) \quad (3)$$

indicating that keyboards are under-predicted on images with men yet men are over-predicted on images with keyboards. Thus the root cause of this amplification appears to be in the gender pre-

dictor rather than the object detector. Such disentangement allows us to properly focus algorithmic intervention efforts. This also highlights the need to consider the ground truth labels on the hold-out set when measuring bias amplification in addition to the predictions (since when considering only the predictions, it is impossible to decouple the different sources of bias).

3.2.3 THRESHOLD

Finally, measuring bias amplification on the binary classification decisions rather than probability outputs leads to a metric that is very sensitive to the threshold picked. Concretely, when using $\text{BiasAmp}_{\text{MALS}}$ on their model predictions, we measure .040, as in Zhao et al. (2017). 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 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 amplification solution proposed in Zhao et al. (2017) of corpus-level constraints, which achieved a drop to only 0.021. Measuring fairness using predictions that are the result of a threshold is common of many fairness metrics, but much like how classification accuracy is often measured with average precision (Zhang & Zhang, 2009) which integrates over all thresholds, measuring across thresholds is a more stable and precise way of directly evaluating the properties of the model.

3.3 OUR METRIC

We address the shortcomings discussed in Sec. 3.2 in $\text{BiasAmp}_{\rightarrow}$. To account for the need to disentangle the two possible directions of bias amplification (Sec. 3.2.2) the metric consists of two values, $\text{BiasAmp}_{A \rightarrow T}$ corresponding to the amplification of bias resulting from the protected attribute influencing the task prediction, and $\text{BiasAmp}_{T \rightarrow A}$, corresponding to the amplification of bias resulting from the task influencing the protected attribute prediction. Concretely, the metric is defined as:

$$\text{BiasAmp}_{\rightarrow} = \frac{1}{|\mathcal{A}||\mathcal{T}|} \int_{thr} \max \left\{ 0, \sum_{t \in \mathcal{T}, a \in \mathcal{A}} \left(y(t,a)\Delta_{thr}(t,a) + (1 - y(t,a))(-\Delta_{thr}(t,a)) \right) \right\} \mathrm{d}thr \tag{4}$$

where $y(t,a) = \mathbb{1}\left[P(T=t, A=a) > P(T=t)P(A=a)\right]$ (5)

$$\Delta_{thr}(t,a) = \begin{cases} P(\hat{T}_{thr} = t | A = a) - P(T = t | A = a) & \text{if measuring } A \rightarrow T \\ P(\hat{A}_{thr} = a | T = t) - P(A = a | T = t) & \text{if measuring } T \rightarrow A \end{cases} \tag{6}$$

Note that $\hat{T}_{thr}$ and $\hat{A}_{thr}$ correspond to the predictions of the task and protected attribute respectively under a classification threshold thr. The metric addresses the shortcomings described in Sec. 3.2 along with a few additional modifications. Concretely, Eqn. 4 generalizes $\text{BiasAmp}_{\text{MALS}}$ to: (1) measure the amplification of both positive and negative correlations between task t and attribute a, depending on $y(t,a)$, when the attributes are non-binary, and (2) integrate the bias amplification over possible thresholds thr to make the metric less sensitive to the threshold picked, as described in Sec. 3.2.3. Eqn. 5 identifies the direction of correlation of attribute a with task t, properly accounting for base rates as described in Sec. 3.2.1. Finally, Eqn. 6 decouples the two possible directions of bias amplification as in Sec. 3.2.2.

3.4 VALIDATING THE $\text{BIASAMP}_{\rightarrow}$ METRIC

Benefits of Thresholding. In Sec. 3.2.3 we discussed how the model outputs from Zhao et al. (2017) ended up exhibiting bias amplification largely due to a thresholding problem. We run $\text{BiasAmp}_{\rightarrow}$ on the same outputs and measure $\text{BiasAmp}_{A \rightarrow T}$=.0009 and $\text{BiasAmp}_{T \rightarrow A}$=.0002. By integrating across all thresholds, our metric was able to correctly discern that there is limited bias amplification.

Disentangling Directions. We now introduce a scenario for validating the decoupling aspect of our metric using a baseline amplification removal idea of masking out parts of the image, such as the person, to remove certain bias (Wang et al., 2019). To do so, we train on the COCO classification task a VGG16 (Simonyan & Zisserman, 2014) model pretrained on ImageNet (Russakovsky et al.,

2015) to predict objects and gender, with a Binary Cross Entropy Loss over all outputs, and measure $\text{BiasAmp}_{T \to A}$ and $\text{BiasAmp}_{A \to T}$; we report error bars for 5 runs of each scenario. We have three hypotheses about what masking should do to bias amplification, and all are substantiated. 1) If we mask the person, $T \to A$ should increase, because once the model only has the objects to base gender prediction off of, it will have to lean into biases from the objects. Indeed, $\text{BiasAmp}_{T \to A}$ increases from $.0481 \pm .0086$ to $.0706 \pm .0071$. 2) Under this same person masking, $A \to T$ should decrease, because we have now removed much of the gender cues that could bias the task. Experiments show $\text{BiasAmp}_{A \to T}$ decreases from $.0028 \pm .0001$ to $.0023 \pm .0004$. 3) If we mask everything but the person, $T \to A$ should decrease, because there are no object visual cues to stereotype gender prediction from. As expected, $\text{BiasAmp}_{T \to A}$ decreases from $.0481 \pm .0086$ to $.0249 \pm .0010$.

3.5 Discussion of T → A Bias Amplification

If we think more deeply about these bias amplifications, we might come to a normative conclusion that $T \to A$, which measures sensitive attribute predictions conditioned on the tasks, should not exist in the first place. There are very few situations in which predicting sensitive attributes makes sense (Scheuerman et al., 2020; Larson, 2017). Thus, before performing prediction of sensitive attributes, we should carefully consider if this is strictly necessary for the target application. For the image domains discussed, the task of gender prediction is poorly-formulated, so by simply removing the notion of predicting gender, we trivially remove all $T \to A$ 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 perhaps by simply replacing all instances of gendered words like "man" and "woman" with "person" in captions, we could remove much of this. 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 $A \to T$ amplification, which is important to discover, as we do not want our prediction task to be biased for or against people with certain attributes.

4 Lack of Consistency in Bias Measurement

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, we noticed that 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 the various local minima, while equal in accuracy, are not necessarily equal in terms of 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).

We look at the facial image domain of CelebA (Liu et al., 2015), defining the two tasks of classifying "big nose" and "young" as our T, and treating the gender labels as our attribute, A. Note that we are not going to classify gender, for reasons raised in Sec. 3.5, so we only measure $A \to T$ amplification. For these tasks, women are more correlated with no big nose and being young, and men with big nose and not being young. 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). Hyperparameters are tuned on the validation set, and the threshold for measuring FPR difference is selected to optimize F1 score, with details in Appendix A.2. In Fig. 3 we see that while the ranking of models across runs is consistent for average precision, the same is not true for the two measures of fairness we display, FPR difference and $\text{BiasAmp}_{A \to T}$.

Given this, we may be left wondering what we can draw conclusions about for various models and algorithms if bias amplification is so fickle across runs. Tuning random seeds certainly does not seem like the way to go. Thus, like previous work (Fisher et al., 2019), we urge transparency, and advocate for the inclusion of confidence intervals for fairness measures. Confidence intervals allow us to determine whether one model is likely to amplify bias more or less than another model.

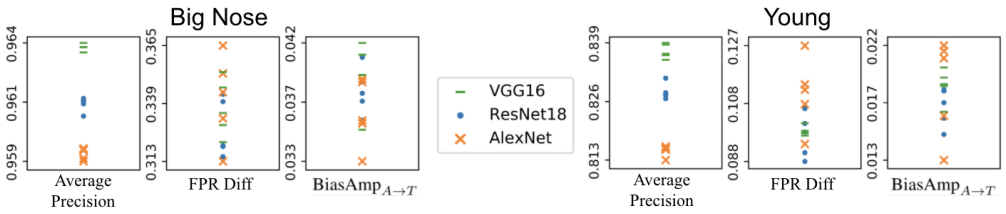

Figure 3: Across 5 runs, the accuracy measure of average precision retains a consistent ranking across models, but two different fairness measures (FPR difference and $A \to T$ bias amplification) have fluctuating rankings between models.

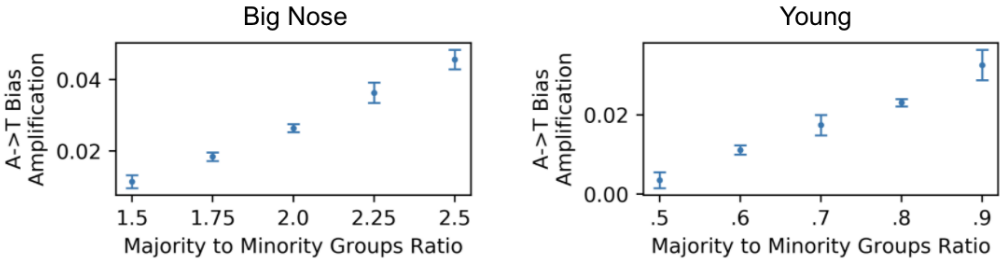

Figure 4: As we adjust the ratio of samples between the minority and majority groups, we see from the confidence intervals that the rankings of bias amplification remain consistent.

We present an example of such a scenario where even though bias amplification fluctuates between identical runs of the same model, we are still able to draw meaningful conclusions. We train new versions of the ResNet18 from before (same trends hold across other architectures) where now we manipulate the ratio of the majority groups (e.g., young women and not young men) to the minority groups (e.g., not young women and young men). While we previously saw that across-model differences in amplification are largely due to a lack of fairness consistency, we can tell from Fig. 4 that across-ratio differences are significant enough to draw conclusions; namely, that as the ratio of samples in the majority to minority group increases, so does bias amplification.

5 Considerations for Bias Amplification

A property of bias amplification is that it does not conflict with having perfect accuracy. However, in turn, such a model with perfect accuracy would exactly replicate the correlations present in the training data. In this section we will discuss two cases that challenge this desire for matching training correlations: in the first, there will be no ground-truth labels from which to lift these correlations, and in the second, our goal is actually to not match the training correlations.

5.1 No Ground Truth

When we don't have labels to derive training correlation rates from, bias amplification becomes harder to measure. 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 equally suitable. Therefore, to measure bias amplification in this setting, we can manually set the base correlations, e.g., $P(T = t|A = a), P(A = a|T = t)$. To see the effect that adjusting base correlations has, we test the bias amplification for 32 different occupations in this task, conditioning on the pronoun and filling in the occupation and vice versa, to measure the two directions of bias amplification. In Table 1, 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. We use: 1) equality amongst the pronouns (using 2 and 3 pronouns), 2) co-occurrence counts from English Wikipedia (one of the datasets BERT was trained on), and 3) WinoBias (Zhao et al., 2018) with additional information supplemented from the 2016 U.S. Labor Force Statistics data. Details are in Appendix A.3. It is

Base Correlation Source	$\text{BiasAmp}_{T \to A}$	$\text{BiasAmp}_{A \to T}$
Uniform (2 pronouns)	.1597 ± .0203	.0097 ± 0.0003
Uniform (3 pronouns)	.1066 ± .0137	.0064 ± .0002
Wikipedia (2 pronouns)	.0892 ±.0060	.0026 ± .0000
2016 U.S. Labor Force (WinoBias) (2 pronouns)	0. ± 0.	0. ± 0.

Table 1: Mean and standard deviation of $\text{BiasAmp}_{\rightarrow}$ for different base correlation sources.

interesting to note that relative to U.S. Labor Force data on these particular occupations, FitBERT actually exhibits no bias amplification. For the occupation of hairdresser, the Labor statistics are biased at 92% women while FitBERT is at 80%, reflecting in fact a reduction in bias amplification. This demonstrates the importance of setting appropriate base correlations, because picking one that already exhibits strong amounts of bias will only flag models that further amplify this. In the next section we discuss another manifestation of this phenomenon, where the training correlation itself would be misleading to compare to, because of the strong amount of bias it contains.

5.2 Future Outcome Prediction

Next, we examine the risk prediction setting, where matching the base correlations may not be our desired goal. The labels here do not represent an objective ground truth because they: 1) suffer from problems like historical and selection bias (Suresh & Guttag, 2019; Olteanu et al., 2019; Green, 2020), and 2) will be used for prediction of a future event for which no one knows the answer. We will put aside the significant problems stemming from the first point for this current discussion, and focus on the latter. In this section, when using the word "predict" we will revert from the machine learning meaning and adopt the colloquial sense of the word, as in the forecast of a future event. What we had called object prediction, we will now call object labeling.

Consider bias amplification in the context of a domain like risk prediction relative to previous domains looked at, such as object detection. The difference is not in tabular versus vision, but rather in prediction versus labeling. The notion of "ground-truth" doesn't quite exist the way we might think about it, because given the input features that define a particular person, one could imagine an individual with these features who does recidivate, and one who does not. The training label is just how someone with these input features once acted, but is not necessarily a rigid indicator of how someone else with these features will act. There is no right answer of something a model could predict that is always going to be correct. On the other hand, in object detection the labels are very much the ground-truth. Thus, in object labeling, bias amplification is a reasonable metric to gauge fairness by, because we know the base correlations are correct and want to match them. However, we do not know this for risk prediction, and thus, matching the training correlations should not be our intended goal. In fact, achieving no bias amplification would probably be the bare minimum first step of fairness for this particular domain, because that would mean the model would still be perpetuating training time biases, which is problematic, but at the very least not amplifying them.

For each application, different metrics of fairness are more or less applicable, and $\text{BiasAmp}_{\rightarrow}$ is no different. It is crucial that we think thoughtfully when deciding how to evaluate a model. In previous applications we imagined the direction of systemic bias to be captured by the training data and thus we lifted base correlations of bias from there. However, one could imagine a similarly manual intervention as was done for NLP on other tasks like risk prediction, where a domain expert verifies the direction of bias determined by the training set, and even sets the base correlations.

6 Conclusion

In this paper, we take a deep dive into the measure of bias amplification. We introduce a new metric, $\text{BiasAmp}_{\rightarrow}$. Through the use of this metric, diagnosing models will provide more actionable insights. Additionally, during such diagnoses and evaluations, we urge the use of error bars when measuring fairness metrics, due to possible fluctuations across runs. And finally, we hope that by discussing how the application of bias amplification varies across different contexts, great care and thoughtfulness will be used when deciding which fairness metrics to use.

REFERENCES

Steve Bearman, Neill Korobov, and Avril Thorne. The fabric of internalized sexism. *Journal of Integrated Social Sciences 1(1): 10-47*, 2009.

Richard Berk, Hoda Heidari, Shahin Jabbari, Michael Kearns, and Aaron Roth. Fairness in criminal justice risk assessments: The state of the art. *Sociological Methods and Research*, 2017.

Jay Bhattacharya and William B. Vogt. Do instrumental variables belong in propensity scores? *NBER Technical Working Papers 0343, National Bureau of Economic Research, Inc.*, 2007.

Tolga Bolukbasi, Kai-Wei Chang, James Zou, Venkatesh Saligrama, and Adam Kalai. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. *Advances in Neural Information Processing Systems (NeurIPS)*, 2016.

Leo Breiman. Statistical modeling: The two cultures. *Statistical Science*, 16:199–231, 2001.

Joy Buolamwini and Timnit Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. *Proceedings of Machine Learning Research*, 81, 2018.

Aylin Caliskan, Joanna J. Bryson, and Arvind Narayanan. Semantics derived automatically from language corpora contain human-like biases. *Science*, 2017.

Kristy Choi, Aditya Grover, Trisha Singh, Rui Shu, and Stefano Ermon. Fair generative modeling via weak supervision. *arXiv:1910.12008*, 2020.

Alexandra Chouldechova. Fair prediction with disparate impact: A study of bias in recidivism prediction instrument. *Big Data*, 2016.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 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.

Jiayun Dong and Cynthia Rudin. Variable importance clouds: A way to explore variable importance for the set of good models. *arXiv:1901.03209*, 2019.

Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through awareness. *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, 2012.

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*, 20, 2019.

Ben Green. The false promise of risk assessments: Epistemic reform and the limits of fairness. *ACM Conference on Fairness, Accountability, and Transparency (ACM FAccT)*, 2020.

Nina Grgic-Hlaca, Muhammad Bilal Zafar, Krishna P Gummadi, and Adrian Weller. The case for process fairness in learning: Feature selection for fair decision making. *NeurIPS Symposium on Machine Learning and the Law*, 2016.

Leif Hancox-Li. Robustness in machine learning explanations: Does it matter? *Conference on Fairness, Accountability, and Transparency (FAccT)*, 2020.

Moritz Hardt, Eric Price, and Nathan Srebro. Equality of opportunity in supervised learning. *arXiv:1610.02413*, 2016.

Sam Havens and Aneta Stal. Use BERT to fill in the blanks, 2019. URL `https://github.com/Qordobacode/fitbert`.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *European Conference on Computer Vision (ECCV)*, 2016.

Lisa Anne Hendricks, Kaylee Burns, Kate Saenko, Trevor Darrell, and Anna Rohrbach. Women also snowboard: Overcoming bias in captioning models. *European Conference on Computer Vision (ECCV)*, 2018.

Niharika Jain, Alberto Olmo, Sailik Sengupta, Lydia Manikonda, and Subbarao Kambhampati. Imperfect imaganation: Implications of gans exacerbating biases on facial data augmentation and snapchat selfie lenses. *arXiv:2001.09528*, 2020.

Shengyu Jia, Tao Meng, Jieyu Zhao, and Kai-Wei Chang. Mitigating gender bias amplification in distribution by posterior regularization. *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2020.

Os Keyes. The misgendering machines: Trans/HCI implications of automatic gender recognition. *Proceedings of the ACM on Human-Computer Interaction*, 2018.

Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 1097–1105, 2012.

Sicong Kuang and B. D. Davison. Semantic and context-aware linguistic model for bias detection. *Proc. of the Natural Language Processing meets Journalism IJCAI-16 Workshop*, 2016.

Matt J. Kusner, Joshua R. Loftus, Chris Russell, and Ricardo Silva. Counterfactual fairness. *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.

Brian N. Larson. Gender as a variable in natural-language processing: Ethical considerations. *Proceedings of the First ACL Workshop on Ethics in Natural Language Processing*, 2017.

Klas Leino, Emily Black, Matt Fredrikson, Shayak Sen, and Anupam Datta. Feature-wise bias amplification. *International Conference on Learning Representations (ICLR)*, 2019.

Paul Pu Liang, Irene Mengze Li, Emily Zheng, Yao Chong Lim, Ruslan Salakhutdinov, and Louis-Philippe Morency. Towards debiasing sentence representations. *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2020.

Tsung-Yi Lin, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, C. Lawrence Zitnick, and Piotr Dollar. Microsoft COCO: Common objects in context. *European Conference on Computer Vision (ECCV)*, 2014.

Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.

Kaiji Lu, Piotr Mardziel, Fangjing Wu, Preetam Amancharla, and Anupam Datta. Gender bias in neural natural language processing. *arXiv:1807.11714*, 2019.

Charles T. Marx, Flavio du Pin Calmon, and Berk Ustun. Predictive multiplicity in classification. *arXiv:1909.06677*, 2020.

Chandler May, Alex Wang, Shikha Bordia, Samuel R. Bowman, and Rachel Rudinger. On measuring social biases in sentence encoders. *Annual Conference of the North American Chapter of the Association for Computational Linguistics (NACCL)*, 2019.

Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. A survey on bias and fairness in machine learning. *arXiv:1908.09635*, 2019.

Joel A. Middleton, Marc A. Scott, Ronli Diakow, and Jennifer L. Hill. Bias amplification and bias unmasking. *Political Analysis*, 3:307–323, 2016.

U.S. Bureau of Labor Statistics. Employed persons by detailed occupation, sex, race, and hispanic or latino ethnicity. 2016. URL `https://www.bls.gov/cps/aa2016/cpsaat11.pdf`.

Alexandra Olteanu, Carlos Castillo, Fernando Diaz, and Emre Kiciman. Social data: Biases, methodological pitfalls, and ethical boundaries. *Frontiers in Big Data*, 2019.

Martin Pawelczyk, Klaus Broelemann, and Gjergji Kasneci. On counterfactual explanations under predictive multiplicity. *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2020.

Judea Pearl. On a class of bias-amplifying variables that endanger effect estimates. *Uncertainty in Artificial Intelligence*, 2010.

Yusu Qian, Urwa Muaz, Ben Zhang, and Jae Won Hyun. Reducing gender bias in word-level language models with a gender-equalizing loss function. *ACL-SRW*, 2019.

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. doi: 10.1007/s11263-015-0816-y.

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. *Proceedings of the ACM on Human-Computer Interaction*, 2020.

Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv:1409.1556*, 2014.

Peter M. Steiner and Kim Yongnam. The mechanics of omitted variable bias: Bias amplification and cancellation of offsetting biases. *Journal of causal inference*, 4, 2016.

Pierre Stock and Moustapha Cisse. ConvNets and ImageNet beyond accuracy: Understanding mistakes and uncovering biases. *European Conference on Computer Vision (ECCV)*, 2018.

Harini Suresh and John V. Guttag. A framework for understanding unintended consequences of machine learning. *arXiv:1901.10002*, 2019.

Ruixiang Tang, Mengnan Du, Yuening Li, Zirui Liu, and Xia Hu. Mitigating gender bias in captioning systems. *arXiv:2006.08315*, 2020.

Sahil Verma and Julia Rubin. Fairness definitions explained. *ACM/IEEE International Workshop on Software Fairness*, 2018.

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. *International Conference on Computer Vision (ICCV)*, 2019.

Jeffrey M. Wooldridge. Should instrumental variables be used as matching variables? *Research in Economics*, 70:232–237, 2016.

Ethan Zhang and Yi Zhang. Average precision. *Encyclopedia of Database Systems*, 2009.

Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. Men also like shopping: Reducing gender bias amplification using corpus-level constraints. *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2017.

Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. Gender bias in coreference resolution: Evaluation and debiasing methods. *North American Chapter of the Association for Computational Linguistics (NAACL)*, 2018.

A Appendix

A.1 Additional Metric Details

We provide additional details here about $\text{BiasAmp}_{\rightarrow}$, as defined in Sec. 3.3.

In practice the indicator variable, $y(t, a)$, is computed over the statistics of the training set, whereas everything else is computed over the test set. The reason behind this, is that the direction of bias

is determined by the existing biases in the training set. Additionally, for computing the integration across all thresholds, we sparsely sample from all of the probabilities outputted to compute an approximation.

Comparisons of the values outputted by BiasAmp$_{\rightarrow}$ should only be done relatively. In particular, within one of the directions at a time, either $A \rightarrow T$ or $T \rightarrow A$, on one dataset. In particular, comparing $A \rightarrow T$ to $T \rightarrow A$ directly is not a signal as to which direction of amplification is stronger.

Because we are only trying to measure the errors that do amplify bias, we clamp negative values to be zero. However, since values may be negative, reporting the aggregated bias amplification value could obscure task-attribute pairs that exhibit strong bias amplification; thus, disaggregated results per pair can be returned for a more detailed diagnosis.

A.2 Details and Experiment from Lack of Consistency in Bias

For the models we trained in Sec. 4, we performed hyperparameter tuning on the validation set, and ended up using the following: ResNet18 had a learning rate of .0001, AlexNet of .0003, and VGG16 of .00014. All models were trained with stochastic gradient descent, a batch size of 64, and 10 epochs.

A.3 Details on Measuring Bias Amplification in BERT

Here we provide additional details behind the numbers presented in Tbl. 1 in Sec. 5.1.

As noted, and done, by Liang et al. (2020), a large and diverse corpus of sentences is needed to sample from the large variety of contexts. However, that is out of scope for this work, where we simply use 2 sentences: "[he/she/(they)] is a(n) [occupation]" or "[he/she/(they)] was a(n) [occupation]" to test

When calculating the amount of bias amplification when the base rates are equal, we picked the direction of bias based on that provided by the WinoBias dataset. In practice, this can be thought of as setting the base correlation, $P(A = a|T = t)$ for a men-biased job like "cook" to be $.5 + \epsilon$ for "he" and $.5 - \epsilon$ for "she" when there are two pronouns, and $.33 + \epsilon$ for "he" and $.33 - \epsilon$ for "she" and "they", where in practice we used $\epsilon = 1\text{e}{-7}$. This ensures that the indicator variable, $y(t, a)$ from Eq. 5, is set in the direction fo the gender bias, but the magnitudes of $\Delta(t, a)$ from Eq. 6 are not affected to a significant degree.

To generate a rough approximation of what training correlation rates could look like in this domain, we look to one of the datasets that BERT was trained on, the Wikipedia dataset. We do so by simply counting the cooccurrences of all the occupations along with gendered words such as "man", "he", "him", etc. There are flaws with this approach because in a sentence like "She went to see the doctor.", the pronoun is in fact not referring to the gender of the person with the occupation. However, we leave a more accurate measurement of this to future work, as our aim for showing these results was more for demonstrative purposes illustrating the manipulation of the correlation rate, rather than in rigorously measuring the training correlation rate.

We use 32 rather than 40 occupations in WinoBias Zhao et al. (2018), because when we went to the 2016 U.S. Labor Force Statistics data (of Labor Statistics, 2016) to collect the actual numbers of each gender and occupation in order to be able to calculate $P(T = t|A = a)$, since WinoBias only had $P(A = a|T = t)$, we found 8 occupations to be too ambiguous to be able to determine the actual numbers. For example, for "attendant", there were many different attendant jobs listed, such as "flight attendants" and "parking lot attendant", so we opted rather to drop these jobs from the list of 40. The 8 from the original WinoBias dataset that we ignored are: supervisor, manager, mechanician, CEO, teacher, assistant, clerk, and attendant. The first four are biased towards men, and the latter four towards women, so that we did not skew the distribution of jobs biased towards each gender.