

Measuring Machine Learning Harms from Stereotypes Requires Understanding Who Is Harmed by Which Errors in What Ways

Angelina Wang
Stanford University, Cornell Tech
angelina.wang@cornell.edu

Solon Barocas
Microsoft Research

Xuechunzi Bai
University of Chicago

Su Lin Blodgett
Microsoft Research

Abstract

Despite a proliferation of research on the ways that machine learning models can propagate harmful stereotypes, very little of this work is grounded in the psychological experiences of people exposed to such stereotypes. We use a case study of gender stereotypes in image search to examine how people react to machine learning errors. First, we use surveys to show that not all machine learning errors reflect stereotypes nor are equally harmful. Then, in experimental studies we randomly expose participants to stereotype-reinforcing, -violating, and -neutral machine learning errors. We find stereotype-reinforcing errors induce more experiential harm, while having minimal impact on participants' cognitive beliefs, attitudes, or behaviors. This experiential harm impacts participants who are women more than those who are men. However, certain stereotype-violating errors are more experientially harmful for men, potentially due to perceived threats to masculinity. We conclude by proposing a more nuanced perspective on the harms of machine learning errors—one that depends on who is experiencing what harm and why.

CCS Concepts

• **Social and professional topics** → **Gender**; • **Computing methodologies** → **Object recognition**.

Keywords

stereotypes, machine learning fairness, harms from biases

ACM Reference Format:

Angelina Wang, Xuechunzi Bai, Solon Barocas, and Su Lin Blodgett. 2025. Measuring Machine Learning Harms from Stereotypes Requires Understanding Who Is Harmed by Which Errors in What Ways. In *The 2025 ACM Conference on Fairness, Accountability, and Transparency (FAccT '25)*, June 23–26, 2025, Athens, Greece. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3715275.3732046>

1 Introduction

Over the past decade, researchers have demonstrated that machine learning models run the risk of learning stereotypical associations. For example, natural language processing models have been shown to associate women with homemakers and men with programmers

[11], while computer vision models have been shown to associate women with shopping and men with driving [118]. These associations can cause machine learning models to make systematic errors, mistakenly reporting, for instance, that female doctors are nurses and that male nurses are doctors [94]. A rich literature has developed offering many more such examples, spurring calls to address the risk that machine learning models might make mistakes that perpetuate harmful stereotypes.

Unfortunately, portions of this literature have suffered from three main limitations that impede effective intervention. First, while some research has asked crowd workers to annotate when errors invoke stereotypes or has drawn on pre-existing inventories of stereotypes [9, 11, 15, 16, 98, 106], machine learning researchers often rely on their own moral intuitions to determine which categories of associations to investigate (e.g., associations between gender and occupation) and to judge which specific associations (e.g., the association between women and nursing) are stereotypes. As a result, certain categories of associations that broader populations might perceive as stereotypical have not been investigated and many discovered associations within these categories are assumed to be stereotypical, even if broader populations would not perceive them to be so. For example, researchers identified an object recognition model as reproducing stereotypes because it amplifies the degree to which labels for kitchen items like “knife,” “fork,” and “spoon” are incorrectly assigned to photos featuring women, and labels for technology-related items like “keyboard” and “mouse” are incorrectly assigned to photos featuring men [118]. However, these intuitions might not always be shared by the broader population; indeed, as we’ll show later in this paper, some of these associations are not seen as invoking stereotypes among a sample of people in the United States.

Second, prior work has tended to ignore whether errors reinforce or violate stereotypes, treating each type of error as equally harmful. Researchers often resort to treating all spurious correlations between objects and specific social groups as harmful, even if correlations violate rather than reinforce stereotypes. As a result, it remains unclear the degree to which the harmfulness of errors depends on whether and how these errors invoke stereotypes. And in fact, as we’ll show, stereotype-reinforcing errors and stereotype-violating errors have disparate, sometimes contrasting, impacts on those harmed.

Third, while stereotypes are frequently invoked to explain why some machine learning errors are more harmful than others [1, 4, 8, 107], little has been done to establish the actual impact that these stereotype-aligned errors have on people in practice. Rather

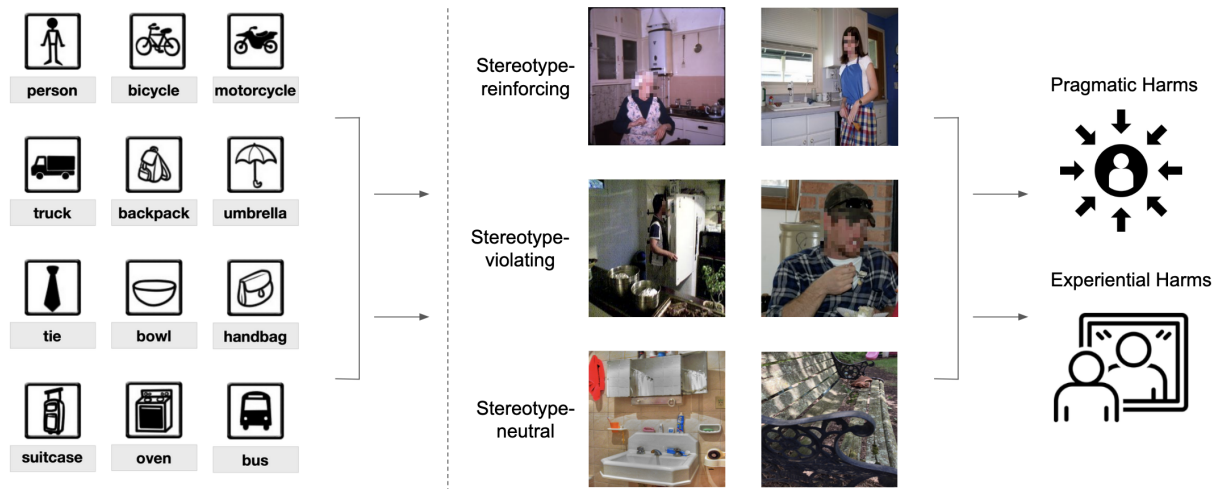


Figure 1: Summary of our studies. The left side represents studies 1 and 2, where we ask human participants to mark which of the relevant objects in our application are stereotypically associated with which gender groups, as well as to qualitatively explain why that is and why it is harmful or not. The right side represents studies 3 and 4 where we randomly expose participants to machine learning errors which are stereotype-reinforcing, stereotype-violating, or stereotype-neutral—as determined by the annotations from study 1. Then, we measure two forms of harm we introduce: pragmatic (measurable changes in someone’s cognitive beliefs, attitudes, or behaviors toward the group being stereotyped) and experiential (self-reports of negative affect). The images shown are examples of misclassifications of oven where stereotype-reinforcing errors are when it is falsely predicted on a woman, stereotype-violating when on a man, and stereotype-neutral when no gender stereotypes are invoked. These images are shown to participants on a search result page of oven.

than measuring the harms that stereotypes bring about and identifying the particular mechanisms by which they do so, much of the literature simply presumes that machine learning stereotypes have negative impacts on the people so stereotyped or on society more generally. While some prior work has found that exposure to gender-biased image search results can lead both to more biased estimations of the representation of different gender groups in certain occupations and to a decreased sense of belonging [65, 73], there remains a paucity of empirical work that seeks to measure the psychological and practical effects of exposure to such errors and to characterize the nature of these harms. In fact, not all stereotypes are harmful [6, 7]—we hold stereotypes that children are bad drivers, yet generally do not see this as a harmful stereotype. Recent FAccT work has pointed out how measures of representational harm like stereotyping have tended to be overly focused on behaviorist definitions, neglecting those which affect cognitive and affective states [17]. In our work, we study these overlooked dimensions based on the psychological notion that stereotypes themselves are cognitive beliefs which may manifest in different ways.

In this paper, we draw on the psychological literature on social stereotypes to overcome each of these limitations, performing a series of empirical studies in which human participants are exposed to machine learning errors (Fig. 1). As a concrete application to ground these studies, we focus on gender stereotypes in the machine learning task of object recognition, which is now a common feature in photo management software like Apple Photos and Google Photos.

To overcome the first limitation of which associations people view as stereotypical, we draw on the psychological literature which

suggests that the space of stereotypical associations is potentially much broader than what has been examined in the machine learning literature [35, 38, 54, 78]. Hence, we make no a priori decisions about the categories of association worth investigating. We expand our scope of analysis beyond commonly focused-on categories—for example, occupations or activities—to include all objects that might appear in an image. We do so via two studies. In the first (Study 1), we present participants with images from popular computer vision datasets that include a range of objects (COCO [68] and Open-Images [67] datasets (Fig. 2)), and ask whether these objects are stereotypically associated with different gender groups. To better understand why certain objects are seen as possessing stereotypical associations, we conduct a qualitative analysis (Study 2) on participants’ open-ended responses to the questions in Study 1.

Then, to overcome the second and third limitations of imprecision about which kind of stereotypical errors may cause what kind of harm, we causally assess the effect that exposure to these stereotypes have on people in practice. We focus on two kinds of potential harm: pragmatic and experiential. Research in psychology describes stereotypes as cognitive beliefs in people’s minds, which can have an influence on attitudes (i.e., prejudice) and behaviors (i.e., discrimination) [2, 52, 54, 63, 69]. Beyond these beliefs, attitudes, and behaviors, members of the stereotyped social group may also feel disrespected, demeaned, or otherwise discounted. Such experiences can be thought of as dignitary harms that bring about negative affect in members of the group so stereotyped [64]. While dignitary harms are frequently treated as less important than the concrete effects of changes in beliefs, attitudes, and behaviors,

they may impose a substantial emotional toll, akin to microaggressions [87]. We therefore introduce a distinction between *pragmatic harms*, which involve measurable changes in someone’s cognitive beliefs, attitudes, or behaviors toward the group being stereotyped, and *experiential harms*, which involve self-reports of negative affect.

Additionally, we differentiate between errors that are stereotype-reinforcing, stereotype-violating, or stereotype-neutral. This is because errors that invoke stereotypes may do so in different ways and may therefore have different effects. For example, an error associating women with homemakers reinforces a gender stereotype and may make women feel their agency is devalued. In contrast, an error associating men with homemakers violates a gender stereotype and may have little effect on men’s self-perceived agency and instead threaten their masculinity [96].

With the above considerations in place, we examine the extent to which participants are harmed by different types of stereotypical errors in an image recognition task. Concretely, we randomly assign participants to synthesized image search result pages which contain different kinds of errors, then measure pragmatic harms through monitoring changes in cognitive, attitude, and behavior measures, and experiential harms by collecting negative affective states reported by participants.

We find little immediate causal effect of pragmatic harms, but sizable evidence that stereotype-reinforcing errors are experientially harmful—a finding that is more pronounced among participants who identify as women compared to those who identify as men (Study 3). We find that while the stereotyped group (e.g., women) generally finds it more experientially harmful for the error to reinforce rather than violate stereotypes, this is not true when it comes to clothing-related items typically associated with women (e.g., cosmetics, necklaces) being misclassified on men. Here, we see a backlash towards violations of the norms around gender presentation: Men tend to find these misclassifications of, e.g., cosmetics, more harmful on men rather than women. This last observation calls into question the idea that it is always normatively desirable to reduce errors perceived as more harmful due to their relationship to stereotypes (Study 4).

By bringing greater clarity to different types of machine learning errors based on their relationship to a stereotype and embracing the rich psychological experiences of such errors, we urge researchers and practitioners to more carefully consider different kinds of machine learning errors, potential harms, and the relevant relationships between them. Investigating the psychological experiences that people have when encountering machine learning errors is critical to understanding the potential harm of a system, and in turn, mitigating it.

The paper is organized as follows: We first discuss related work, then give context for our case study on gender stereotypes in photo search. Next, we report four human studies, and conclude with a discussion of our results.

2 Related Work

Stereotyping is a common (often implicit) operationalization of “bias” that is studied in the literature on fairness in machine learning [4, 8]. Abbasi et al. [1] mathematically formalize stereotyping


and show how models trained on stereotyped data can lead to unfair outcomes. Bias amplification is a different statistical notion that rests on the idea that any amplification of an existing bias is undesirable [51, 108, 110, 118]. We show in our work that these approximations of the type of error that is more harmful is often poorly correlated with human judgment, and provide a detailed comparison in Appendix H. Work in natural language processing considers stereotypes reflected in word embeddings [11, 15] and across many tasks, such as language modeling [16, 59, 77], machine translation [91], and question answering [84], with critiques pointing out that benchmarks aimed at measuring stereotyping may fail to meaningfully conceptualize and operationalize stereotyping [10]. Though researchers have proposed different operationalizations of stereotypes [20, 26, 42], there is generally a large reliance on occupation data from the American Bureau of Labor Statistics—e.g., WinoBias [119]—as a data source for stereotypes. One large limitation of this usage (in addition to it only representing American occupation data) is that it only captures descriptive stereotypes, e.g., overrepresentations of groups in an occupation, and misses prescriptive stereotypes, e.g., beliefs about what occupations people of different groups should be in; these two types can often differ in practice [13, 70]. More recent work has addressed some of these issues by proposing a community-driven approach to collecting stereotypes that are uniquely present in India [23, 58].

Psychology of Stereotypes and their Consequences. Psychology research describes stereotypes as cognitive beliefs, which can have an influence on attitudes (i.e., prejudice) and behaviors (i.e., discrimination) [2, 52, 54, 63, 69]. For example, people may have *cognitive beliefs* that women are more warm but less competent than men, and thus *express* protective attitudes and pity for women [19, 44]. People then *behave* in ways that maintain women’s warmth and discount their competence, such as being less likely to promote women to leadership positions [35, 38]. Therefore, stereotypes of certain social groups can prompt shifts in attitudes and behaviors that can ultimately harm the stereotyped group. We draw on these distinctions to formulate our measure of pragmatic harm.

Image Search. We concretize our task of object recognition by constructing human studies involving the application of image search. There is a large body of work on biases in image search engines [50, 65, 73, 82, 83, 105], as well as search algorithms and media systems more broadly [29, 47, 56, 80].

Most findings on image search engines have focused on auditing the search engine itself for biased results [50, 65, 73, 82, 83, 105], leaving more speculative the implications of these results for feeding back into the stereotypes people hold. However, some components of this work also attempt to understand the influence that biased search results can have on users. Kay et al. [65] and Metaxa et al. [73] both find that exposure to gender biased image search results can influence an individual’s opinion on the gender representation of that occupation (accounting for 7% of a person’s opinion in the former and roughly 5% increases, with great variance, when representation shifts from 10% to 50% to 90% in the latter). Metaxa et al. [73] also find an effect of .146 (95% CI [-.16, .45]) on a 7-point Likert scale for the impact of gender representation on feelings of an occupation’s inclusivity, with no effect for interest in an occupation. Notably, these feelings of inclusivity depend on the gender of the participant. Vlasceanu and Amodio [105] take a different approach

Common Objects in Context (COCO)



man; person; surfboard

Open Images



woman; person; tree; wheel; tire; dog

Figure 2: COCO and Open Images object recognition datasets. We use two commonly used image recognition datasets to represent the application of a photo search engine. Both datasets contain annotations for perceived binary gender expression of the people in the images and the objects present in each image. The left panel shows one example figure from COCO annotated with surfboard. The right panel shows one example figure from Open Images annotated with objects such as tree and dog.

by studying occupations (e.g., peruker, lapidary) for which there are very few preconceived stereotypes. Our work focuses on the reinforcement of existing stereotypes, rather than trying to induce new ones. Recent work by Guilbeault et al. [50] conducts a large-scale study of social categories including occupations and roles, and measures psychological effects such as through the Implicit Association Test. Unlike much of this existing work, we do not audit an existing image search engine, instead focusing on downstream concerns that may arise from biases in the search results. We also expand the scope of what is a stereotype (e.g., objects).

Crowdsourcing Opinions on Harm. We solicit our labels of stereotypes and harm from crowdworkers, similar to how prior work has elicited stereotypes [11, 98] and fairness notions in different contexts [49, 100, 116]. Woodruff et al. [114] specifically focus on the opinions of individuals from marginalized communities who are likely to be adversely affected by algorithms. Unlike this existing work, however, we further attempt to measure the psychological effect of exposure to machine learning errors, rather than relying exclusively on self-reported opinions. In particular, we measure changes in participants’ beliefs, attitudes, and behaviors, as well as the extent to which they feel that they have been *personally* harmed by these errors.

3 Case Study: Gender Stereotypes in Photo Search

We explore a popular task in machine learning known as object recognition (i.e., classifying the objects present in an image). To make it concrete for our human studies, we use it in the context of a smart phone’s photo search engine, and examine gender stereotypes. Specifically, we consider one type of machine learning error called a false positive: when an object is predicted to be present in an image when it is in fact not there. This causes the image with a false

positive to be wrongly surfaced on an image search results page.¹ We study the *effect* of the misclassification, not why the model may have made the mistake, or what the participant thinks is the reason the model made the mistake. We use a mixture of qualitative and regression analyses to report our findings. For our within-subjects surveys, we regress with a mixed-effect model whose parameter estimations are adjusted by the group random effects for each individual. We report the coefficients from our regression analyses, which represent the effect size of that independent variable. Further details on our methods and analysis are in Appendix A, and code is at <https://tinyurl.com/mlstereotypes-code>. To protect participants’ privacy, data is available upon request. All studies are approved by our institution IRB. Studies 1 (<https://tinyurl.com/mlstereotypes-study1>), 3 (<https://tinyurl.com/mlstereotypes-study3pt1>, <https://tinyurl.com/mlstereotypes-study3pt2>), and part of Study 4 (<https://tinyurl.com/mlstereotypes-study4>) are pre-registered on OSF, while Study 2 is exploratory.

4 Study 1: Distinguishing which machine learning errors reflect social stereotypes

4.1 Method

To understand the social stereotypes held by American society relevant to our machine learning task, we first elicit human judgments ($N = 80$) on Common Objects in Context (COCO) [68]. COCO has 80 objects and perceived binary gender expression of pictured people annotated across the images [117]. In the study, we ask the participants whether each object (e.g., keyboard, zebra) is stereotypically associated with men, women, or neither. We recruited U.S.-based participants on Cloud Research with filters for those

¹We note that false negatives are subsumed in this setting because enough false positives will crowd out the results page and ultimately have a similar effect as false negatives on images of the gender that does not have false positives.

who had at least 50 HITs approved, and a HIT approval rate of 98%. We paid a rate of \$15/hour and explicitly recruited equal numbers of participants who identified as men and as women.

4.2 Results

As expected, not all objects reflect gender stereotypes. This is already in contrast to a somewhat common assumption in ML fairness research that *any* difference between groups is an amplification of a stereotype [10]. Among 80 objects, 13 objects are marked as stereotypes by more than half of the participants (Figs. 3). Some examples of stereotypically gendered objects are handbag with women, wine glass with women, tie with men, and truck with men. Among the remaining objects, 18 objects (e.g., keyboard, carrot, traffic light) are marked by zero participants as stereotypes with any gender group, challenging prior assumptions on what is seen as a stereotype [118]. If an object was marked to be a stereotype, we also asked participants whether they believed it was harmful in the abstract. Complete results are in Appendix C, but we use these initial findings to select experimental stimuli in subsequent studies. In Study 3a the stereotype-reinforcing condition includes women and oven (marked to be most harmful), women and hair dryer (marked to be least harmful), and the associated control conditions include women and bowl, women and toothbrush. In Study 3b we also include in the stereotype-reinforcing conditions of men and baseball glove (marked to be more harmful) and men and necktie (marked to be less harmful) with the control conditions of men and bench and men and cup.

5 Study 2: Plurality of stereotypes and harms with image recognition objects

5.1 Method

Next, we report qualitative analyses on open-ended responses from participants' annotations, where they explain why certain objects are seen as stereotypes and harmful or not. While prior work in gender stereotypes has often focused on social roles and traits [33, 44], our data provides insights as to how objects (e.g., oven, hair dryer) can also be associated with stereotypes. This is an important departure because it expands the scope of machine learning tasks for which stereotypes are relevant beyond its current more narrow framing. Specifically, when a participant from Study 1 responds that an object is a stereotype, we follow up and ask: "Please describe in 1-2 sentences a) why you marked the above as a stereotype, and b) why you found it to be harmful or not."

5.2 Results

One of the authors coded the responses for why an object is a stereotype into six categories. The coding was done inductively, then thematically clustered to form the six categories. The most prevalent reasons were: descriptive (45%), e.g., for handbag and women: "women are often seen wearing handbags and buying them"; occupation/role (22%), e.g., for oven and women: "women are stereotyped to always be in the kitchen cooking while the men go out and work"; trait (11%), e.g., for chair and men: "sometimes men would be seen as coming home and just being lazy and lounging in their chair." The full analysis is in Appendix D. It is interesting

to note that an object's association to a stereotype is frequently mediated by its connection to a role or trait, which are the more common sites of inquiry when it comes to stereotypes [33, 35, 38]. We also found that associations between a group and an object can exist through a number of paths. For example, explanations for stereotypical associations between cats and women include: "cat lady," "women are called *kitten*," "women like cats more than dogs," "cats are a feminine animal," and "women are called *cougars*."

When asked why a stereotype was harmful or not, many respondents simply reiterated that the object was a stereotype. Dropping those responses, one of the authors coded the free responses of why a stereotype was marked to be harmful into seven categories, with the top three being: proscriptive (40%), e.g., for dining table and women: "it makes it looked down upon if a man cooks dinner"; prescriptive (26%), e.g., for dining table and women: "I think it puts women in a box that says they must prepare dinner"; negative trait (13%), e.g., for handbag and women: "it is harmful because it implies that women cares more about looks and their appearance." The remaining response categories are in Appendix D. There seems to be a disparity in responses based on the participant's gender regarding whom they believe is harmed. When women specify which of the men group or women group are harmed, they say it is the women group 79% (95% CI [.67, .88]) of the time, while men say it is the women group only 67% (95% CI [.51, .80]) of the time.

Building on Study 1's finding that participants do not even all agree on whether an object is a stereotype or not (and if it is, whether it is harmful), this analysis further shows that even when participants are in agreement that an object is a stereotype, they are not necessarily in agreement about why. The same holds true for whether a stereotype is harmful. One potential implication of this is considering whether different reasoning should lead to different bias mitigation. For example, if the reason an object is a stereotype is descriptive, then mitigation should aim to change the cognitive representations of people. To change these descriptive statistics, while we can work to alter the model outputs, we should also work to change society, the burden of which falls on a much larger group than just machine learning practitioners, e.g., policymakers. On the other hand, if particular stereotypes are deemed harmful because they are proscriptive and seem to restrict people from various avenues, we can consider ways to break free of gender norms, including by representing groups in ways that depart from those proscriptive norms.

6 Study 3a: Errors that reinforce stereotypes may not lead to pragmatic harm

Next, we test the pragmatic harm to users. Following the tripartite model of attitudes [2, 52, 54, 63, 69?] which differentiates between cognitive beliefs, emotions, and behaviors, we define pragmatic harm as negative changes in these three dimensions.

6.1 Method

To test pragmatic harm in stereotype-reinforcing errors, we conduct a between-subject survey experiment, using the stereotype-violating and neutral errors as control conditions (Figure 1). The cover story instructs participants to look at our synthesized search result page, imagining it is their personal phone photo album, and

Stereotyped with Women								Stereotyped with Men							
handbag 21/23	hair drier 17/20	wine glass 8/13	cat 11/19	potted plant 9/17	oven 12/23	cake 8/19	vase 7/18	tie 11/14	truck 14/18	motorcycle 8/11	baseball bat 15/21	baseball glove 11/18	sports ball 12/21	skateboard 10/19	fire hydrant 9/25
dining table 5/21	tennis racket 3/14	cow 3/15	sink 4/20	teddy bear 4/20	horse 5/26	bird 4/22	person 4/26	bear 6/17	snowboard 6/18	remote 7/22	car 6/23	suitcase 6/25	surfboard 5/21	sandwich 5/23	microwave 2/12
sandwich 3/23	umbrella 3/25	parking meter 2/19	toaster 2/21	refrigerator 2/23	sheep 2/24	suitcase 2/25	backpack 2/26	donut 3/19	person 4/26	boat 2/14	tennis racket 2/14	bicycle 2/20	bottle 2/20	tv 2/21	dog 2/21
bench 1/14	apple 1/17	snowboard 1/18	frisbee 1/18	scissors 1/18	baseball glove 1/18	bicycle 1/20	bottle 1/20	couch 2/22	knife 2/22	sheep 2/24	backpack 2/26	bench 1/14	hot dog 1/14	vase 1/18	frisbee 1/18
dog 1/21	book 1/21	knife 1/22	remote 1/22	cup 1/22	mouse 1/22	toothbrush 1/24	chair 1/24	pizza 1/18	cake 1/19	hair drier 1/20	teddy bear 1/20	train 1/20	elephant 1/20	toaster 1/21	stop sign 1/21
orange 1/25	fork 1/26	cell phone 1/30	boat 0/14	bowl 0/18	bus 0/16	skis 0/25	toilet 0/15	chair 1/24	toothbrush 1/24	umbrella 1/25	skis 1/25	horse 1/26	cell phone 1/30	potted plant 0/17	cup 0/22
stop sign 0/21	tie 0/14	keyboard 0/15	sports ball 0/21	baseball bat 0/21	spoon 0/17	carrot 0/25	donut 0/19	scissors 0/18	traffic light 0/17	dining table 0/21	fork 0/26	book 0/21	orange 0/25	toilet 0/15	mouse 0/22
couch 0/22	train 0/20	kite 0/17	clock 0/24	giraffe 0/19	pizza 0/18	zebra 0/19	truck 0/18	cat 0/19	refrigerator 0/23	spoon 0/17	bus 0/16	oven 0/23	zebra 0/19	carrot 0/25	giraffe 0/19
traffic light 0/17	motorcycle 0/11	skateboard 0/19	microwave 0/12	car 0/23	bed 0/15	laptop 0/24	elephant 0/20	bed 0/15	laptop 0/24	sink 0/20	broccoli 0/14	banana 0/26	clock 0/24	bowl 0/18	wine glass 0/13
broccoli 0/14	bear 0/17	banana 0/26	hot dog 0/14	surfboard 0/21	fire hydrant 0/25	airplane 0/18	tv 0/21	parking meter 0/19	airplane 0/18	kite 0/17	handbag 0/23	cow 0/15	keyboard 0/15	bird 0/22	apple 0/17

Figure 3: Study 1 Object Results. Detailed participant responses for each of the 80 objects in COCO dataset. Fraction indicates number of participants asked about each object who marked it as stereotypically related to the gender group of women or men.

find a picture they had taken of someone they saw with a particular object. The search result page looks different for each randomized condition. We randomly assign participants to one of the three conditions ($N = 600$): the stereotype-reinforcing condition exposes an image search result page with stereotype-reinforcing errors, e.g., false positive of oven on images of women; the stereotype-violating condition contains the same for stereotype-violating errors, e.g., false positive of oven on images of men; the stereotype-neutral condition contains neutral errors, e.g., false positive of bowl on images of women. The errors are actual misclassifications from a trained computer vision model to make them as closely aligned to realistic errors a user would encounter. We then measure participants' cognitive beliefs, attitudes, and behaviors [38] to see if there are any changes because of such exposure. For *cognitive beliefs*, we ask three sets of questions which span the spectrum of stereotype-specific to more generically about gendered beliefs. Concretely, we ask about: estimations of who uses ovens and bowls more between men and women; estimations of who tends to be the homemaker more between men and women; and perceived levels of warmth and competence [38] of women. To assess *attitude*, we ask two sets of questions. The first is about how participants feel about women in terms of four emotional components that are believed to mediate interactions between cognitive beliefs and behaviors: a) respect or admiration, b) pity or sympathy, c) disgust or sickening, and d) jealousy or envy [19, 37, 93]. The second asks about sexist attitudes via a shortened scale focused on benevolent sexism [44, 45, 88].²

²We ask questions from the Ambivalent Sexism Inventory [44] about benevolent sexism, as opposed to hostile sexism, because the latter is believed to suffer heavily from social desirability bias.

Finally, for *behavioral* measures, we ask participants to undertake a realistic task they are liable to encounter which can cause harm: data labeling [103]. We chose this behavior measure because online participants are often the source of training labels in large-scale machine learning datasets. We ask participants to perform two common types of labeling on image data: tagging and captioning. In the tagging task, we ask participants to label the top three most relevant tags in an image which contains both the stereotype object (e.g., oven) and neutral object (e.g., bowl). We alter the perceived gender of the person to assess whether this changes what is tagged in the image. For the captioning task we show two people, one who looks masculine and another feminine, and swap whether there is a bowl or oven present in the image. This is to understand if the annotators will differently describe who is interacting with the object depending on whether it is stereotypically associated with women or not. All images are generated and/or manipulated by DALL-E 2. Additional details are in Appendix A. If stereotype-reinforcing errors have an influence on participants' cognitive representations, attitudes, and tagging or captioning behaviors, we should expect to see a statistically significant difference between participants who are exposed to search results with oven-women and those who are exposed to search results with oven-men or bowl-women.

6.2 Results

Contrary to what we had originally hypothesized, after adjusting for multiple comparisons we do not find hypothesized statistically significant differences. We run an Ordinary-Least-Square (OLS) regression with the control condition coded as 0 and the experimental condition coded as 1, composite scores for beliefs, attitudes,

and behaviors respectively as the dependent variables. Results are shown in Fig. 4. In other words, participants' cognitive beliefs, attitudes, and behaviors towards gender groups, although stereotypical, remain stable before and after the treatment effect.

Prior work looking at a subset of these pragmatic harms has found very small effects in terms of cognitive belief changes about the representation of gendered occupations [65, 73]. Another line of work that finds a cognitive effect takes a different approach by studying novel occupations for which there are very few preconceived stereotypes [105]. In our work, we focus on the activation of existing stereotypes, rather than the inducement of novel stereotypes. Overall we find that the pragmatic harms are not measurable after exposure in the current survey experiment, which serves as a conservative test of harm, likely because the effects of these harms are too diffuse and long-term, impacted by all of the facets of society we encounter in our lives [80].

7 Study 3b: Errors that reinforce stereotype lead to experiential harm

Finally, we measure users' experiential harm in response to different machine learning errors. This includes self-reports of negative affect based on the Positive and Negative Affect Schedule (PANAS; [18, 112]). This is used as a measure of dignitary harms [64]. In conceptualizing the experiential harm of machine learning errors which may seem individually minor, we draw a parallel to the concept of microaggressions. Microaggressions are "small act[s] of insult or indignity, relating to a person's membership in a socially oppressed group, which seems minor on its own but plays a part in significant systemic harm" [87]. Just like how a machine learning model's classification error (e.g., of an oven on an image of a woman) may seem small on its own, and are "easily interpretable as inadvertent errors rather than as malevolent actions," their negative effects on the target are real and should not be neglected [87].

7.1 Method

In terms of experiential harm, we design a within-subjects survey experiment ($N = 100$). We operationalize experiential harm by explicitly asking participants to rate how personally harmful they find different kinds of errors (which are stereotype-reinforcing, stereotype-violating, or neutral), on a scale from 0 (not at all) to 9 (extremely). This experience of error is analogous to situations where one reads in the news about the types of errors that artificial intelligence systems make [97], notices such a pattern of errors themselves, or is informed by a friend.

7.2 Results

In line with what we hypothesized, when we compare stereotype-reinforcing against neutral errors, an OLS regression shows participants rate stereotype-reinforcing errors to be more harmful than neutral ones ($b = .62$, 95% CI [.32, .91], $p < .001$). However, when disaggregating by gender this effect is only present among women participants (women: $b = 1.06$, 95% CI [.64, 1.47], $p < .001$; men: $b = .18$, 95% CI [-.23, .59], $p = .393$). When we use the stereotype-violating error as the control condition rather than the neutral error, we again find participants rate stereotype-reinforcing errors to be more harmful, though to a smaller degree, ($b = .28$, 95% CI [-.01,

.58], $p = .062$), once again an effect only for women participants (women: $b = .73$, 95% CI [.31, 1.14], $p = .001$; men: $b = -.16$, 95% CI [-.58, .26], $p = .453$). See Fig. 4.

In short, while we find little immediate evidence of pragmatic harms, we do find the existence of experiential harms resulting from stereotype-reinforcing errors, compared to both stereotype-violating and neutral errors. However, this pattern is present only among woman participants, and not men participants.

8 Study 4: Errors that violate stereotypes can be perceived as harmful too

8.1 Method

We first test the generalizability of the previous findings by replicating the experiment and analysis with another popular dataset in object recognition tasks which is much larger: OpenImages [67].

Then, to test whether participants experience more experiential harm when they are exposed to stereotype-reinforcing (e.g., skirt on women), stereotype-violating (e.g., skirt on men), and neutral (e.g., toothbrush on women) errors, we use a similar procedure as in Study 3b. Rather than asking simply about "personal harm" as we did in Study 3b, here we draw from the Positive and Negative Affect Schedule (PANAS; [18, 112]) and provide more details by asking about if they experience harm such as feeling upset, irritated, ashamed, or distressed. We conduct a within-subjects study and ask participants ($N = 300$) to report their subjective experiences on a Likert scale from 0 to 9 for errors which vary in their relationship to a stereotype. The analysis uses a mixed-effects regression with experimental conditions as the independent variable, a composite score of experiential harm as the dependent variable, participants' gender as the covariate variable, and error terms clustered at the individual level.

Then, we explore a new hypothesis that is not pre-registered about gender presentation-aligned objects, e.g., clothing, to dive deeper into the 40 stereotypical objects to understand why stereotype-violating errors are sometimes perceived to be more experientially harmful than stereotype-reinforcing ones. According to the gender trouble framework, costume (i.e., body and appearance) and script (i.e., behavior, traits, and preferences) are two aspects of gender performance, and reactions to androgynous or conventionally contradictory components can differ depending on which of the two it manifests in [14, 46, 74, 75, 102]. We thus hypothesize that conventionally contradictory costume objects may evoke more negative reactions compared to conventionally contradictory script objects [89]. To test this hypothesis, we explore an additional independent variable we call "wearable." We determined the value of this variable by manually marking 13 of the 40 stereotypical objects to be conventionally wearable by a person. These include objects like football helmet and lipstick, and exclude those like truck or wine glass.

8.2 Results

OpenImages has 600 objects, annotated with perceived binary genders of people present in the image if applicable [92]. Following the same procedure as in the COCO dataset with new online participants ($N = 120$), we find 249 of the 600 objects are marked as stereotypes by more than half of the participants, replicating




Figure 4: Study 3, 4 Results. The effect sizes and 95% confidence intervals are reported for 10 of our 11 measures of pragmatic harm (for the behavior measure of captioning, we provide a descriptive analysis), experiential harm on COCO, and experiential harm on our larger dataset of OpenImages. Deviations from zero indicate that exposure to the stereotype-reinforcing stimulus resulted in our measured harm compared to exposure to the control condition.

the finding that not all objects are perceived as stereotypes. We then compile a list of 40 stereotypical objects (20 about men: e.g., football, tool; 20 about women: e.g., doll, lipstick), and 20 neutral objects (e.g., balloon, goldfish) for this study.

Replicating Study 3b, participants find stereotype-reinforcing errors to be more harmful than neutral ones ($b = .50$, 95% CI [.42, .59], $p < .001$). Again, this pattern is more pronounced among women participants ($b = .67$, 95% CI [.55, .79], $p < .001$), with now a small effect among men participants ($b = .33$, 95% CI [.21, .45], $p < .001$). Unlike Study 3b, we do not see differences in experiential harm between stereotype-reinforcing and stereotype-violating conditions ($b = -.04$, 95% CI [-.13, .05], $p = .338$). The effect is canceled out by the opposite effects for women ($b = .25$, 95% CI [.13, .38], $p < .001$) and men ($b = -.34$, 95% CI [-.46, -.22], $p < .001$) participants. In other words, while women participants feel upset, irritated, ashamed, and distressed when they see stereotype-reinforcing errors (e.g., skirt on women), men participants feel that way when they see stereotype-violating errors (e.g., skirt on men). Results are in Fig. 4.

To better understand this finding, with the “wearable” distinction, we find that participants do rate stereotype-reinforcing errors to be more harmful than stereotype-violating ones ($b = .23$, 95% CI [.12, .34], $p < .001$), though again this effect exists in women participants ($b = .49$, 95% CI [.34, .64], $p < .001$) rather than men participants ($b = -.03$, 95% CI [-.18, .12], $p = .726$). Notably, for the interaction effect of a “wearable” object with the condition type, we find that wearable stereotype-violating errors have higher experiential harm than wearable stereotype-reinforcing errors ($b = .80$, 95% CI [.62,

.99], $p < .001$), which is higher for men ($b = .94$, 95% CI [.67, 1.12], $p < .001$) than women participants ($b = .69$, 95% CI [.43, .94], $p < .001$). In other words, men participants find it more harmful than women participants when lipstick is misclassified on a man than on a woman.

In sum, stereotype-violating errors seem to cause harm too, possibly through different mechanisms. In addition to this result being a consequence of backlash effects [90], we raise two more possible mechanisms. First, it could be seen as an expression of precarious manhood; a concept that suggests manhood is precarious and needs continuous social validation such that threats to traditional masculinity can provoke anxiety in men, thus resulting in higher reports of harm [104]. Second, these results may reflect elements of transphobia, which involves a negative reaction to the apparent incongruity between a person’s perceived gender and a wearable gender presentation item [14, 75]. The divergent effect between men and women participants aligns with research indicating that transphobia is higher amongst cisgender men when judging transgender women due to the perceived threat to masculinity [71, 76]. This analysis pushes us to reevaluate how we should think about reducing experiential harm, as it may encompass intolerances we do not wish to support.

9 Discussion

Taken together, our studies highlight the *complexity* of machine learning errors. We offer three key contributions: (1) In contrast to prior assumptions which rely on a restricted set of stereotypes based

on researcher intuitions, we present a much broader set of stereotypes that are defined by participants. (2) Compared to general notions of stereotypes or harms, we distinguish between machine learning errors that are stereotype-reinforcing, stereotype-violating, or stereotype-neutral, and formulate harm as pragmatic (changes in cognition, attitude, and behaviors) or experiential (self-report negative affect). (3) Our empirical experimental studies suggest that stereotype-reinforcing errors cause experiential, but not pragmatic, harm. Such experiential harm is unequally distributed, impacting more women participants than men. We speculate that long-term observational studies, compared to lab studies, are likely better suited to capture pragmatic harms [39, 41, 57].

Formulating concrete notions of harm as we have done has implications beyond just machine learning: legal documents like the European AI Act are beginning to incorporate notions of psychological harm but lack definitions in which to ground regulation [5, 85]. We also find stereotype-violating errors to be experientially harmful, especially when these errors pertain to wearable items associated with gender presentation. This effect is stronger for participants who identify as men compared to as women. This point warrants an especially nuanced discussion, as it questions whether we should take people's words at face value when they indicate something is personally harmful. To navigate this complexity, we can turn to the notions of epistemic injustice [40] and standpoint epistemology [36, 81, 115]. If we interpret the negative reactions to misclassifications of stereotypically feminine clothing items on men as a manifestation of precarious manhood [104] or transphobia [14], then we should down weight these concerns in designing mitigation approaches. Respecting people's experiential harms may not be as simple as accepting them at face value for direct measurement, but rather involves understanding which groups are likely to be harmed by what kinds of errors and why. This underscores the importance of work about diversity in annotators' perspectives [21, 22, 31, 60, 79, 111], and how much complexity is reduced by the use of discrete labels. Qualitative follow-up questions supplemented our annotations, where a lack of consensus is not a weakness or artifact to be averaged out, but rather a point for deeper inquiry on how to prioritize differential experiences.

Our findings call for reconsidering fairness measurement in supervised machine learning tasks. This involves considering how we can leverage human-driven insights to inform model training and evaluation [12]. Traditionally, fairness evaluations tend to focus on stereotypes only in relation to occupations or traits. However our work expands this idea by showing that labels such as objects can also give rise to such harms. Additionally, most prior work has only considered the implications of errors that reinforce stereotypes, which is relatively more intuitive to think of as harmful. However, both practically and normatively, it is important to understand the implications of stereotype-violating errors. Practically, strategies aimed at mitigating stereotype-reinforcing errors which act upon the target label will inevitably impact the occurrence of stereotype-violating errors as well. And normatively, there are also questions about whether stereotype-violating errors may even play a role in reducing stereotypical associations by counteracting them.

This finding that not only are certain labels more liable to cause harm than others, but that it matters for *which* demographic group that label is misclassified, suggests that generic approaches like

having a higher threshold for the classification of certain labels are insufficient. Instead, more nuanced fairness-through-awareness approaches [32] will need to be taken. While adopting a cost-sensitive framework [66] (e.g., different costs are associated with false positives and false negatives) is a simplified interpretation of our findings, it could be a starting point as one grapples with the questions of whose levels of harms we would prioritize reducing in a bias mitigation framework.

Our findings are limited by the methodological choices we made. For instance, we focused on gender stereotypes and do not know to what extent this finding generalizes to other groups such as race and age. Second, we recruited online participants in America who identify as men and women, missing other populations such as non-binary individuals or those from other cultural contexts. Given that stereotypes are culture-sensitive, and our work also shows that the harm perception is identity-sensitive, future work needs to study the interaction between participants' identity, culture, and harm perceptions. Third, the survey experiment does not capture harms beyond the two we measure (e.g., stereotype-threat [99, 101]), nor the longitudinal effects of machine learning effects. Future work needs to capture not only the plurality of harms from machine learning errors but also how their effects over time.

Overall, drawing on research and methods from psychology, our work offers a rigorous empirical study connecting machine learning outputs to concrete harms by understanding the impact of stereotypical misclassifications. Rather than gesturing at harm as a justification for fairness measurement, we are very concrete in our analysis of the effects on people. Our finding that stereotype-reinforcing errors are experientially harmful for women underscores the importance for machine learning fairness interventions to be more rooted in social contexts, moving beyond objectives like just achieving equal prediction performance across groups. The diversity of responses we've presented, each influenced by participants' unique rationales, suggests the need for greater exploration of human psychological experiences in understanding how machine learning can cause harm.

Ethical considerations

As this is a study that involves human subjects, we sought and obtained IRB approval from our institution. Our study ran the risk of inducing negative feelings among our participants, given that we were investigating whether certain conditions caused subjective harm. That said, our study aimed to simulate conditions that participants likely encounter in their everyday lives already and thus tried to avoid exposing them to greater discomfort than they might experience normally. We also recruited only participants who identified as men or women, and asked questions about stereotypes related to binary gender, which may have caused its own kind of subjective harm to any non-binary participants that we did not enlist in the study.

Researcher positionality

All of us identify as cisgender and thus do not have the direct lived experience of the potentially transphobic harms we find in our work.

Adverse impact statement

While our study does not find statistically significant differences in beliefs, attitudes, and behaviors immediately following exposure to stereotype-associated machine learning errors, this does not mean that such errors do not give rise to cumulative, long-term pragmatic harms. There is a risk that our result will be misinterpreted to support the view that such errors are not harmful or worth addressing, or that such pragmatic harms are too difficult to meaningfully measure. Moreover, the experiential harms that we do detect might be dismissed as purely subjective and therefore relatively insignificant, despite these taking a significant emotional toll. Finally, we emphasize that our takeaway from this finding is not to take self-reports of harm always at face value, but rather to think critically about why people self-report harm and what harms they are actually experiencing.

Acknowledgments

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship to Angelina Wang. We are grateful to funding from the Data-Driven Social Science Initiative at Princeton University. We thank Orly Bareket, Molly Crockett, Sunnie S. Y. Kim, Anne Kohlbrenner, Danaë Metaxa, Vikram V. Ramaswamy, Olga Russakovsky, Hanna Wallach, and members of the Visual AI Lab at Princeton, Fiske Lab at Princeton, and Perception and Judgment Lab at the University of Chicago for feedback.

References

- [1] Mohsen Abbasi, Sorelle A. Friedler, Carlos Scheidegger, and Suresh Venkatasubramanian. 2019. Fairness in representation: quantifying stereotyping as a representational harm. *Siam International Conference on Data Mining* (2019).
- [2] Gordon Willard Allport, Kenneth Clark, and Thomas Pettigrew. 1954. The nature of prejudice. *Addison-wesley Reading, MA* (1954).
- [3] Lisa P. Argyle, Ethan C. Busby, Nancy Fulda, Joshua Gubler, Christopher Rytting, and David Wingate. 2023. Out of One, Many: Using Language Models to Simulate Human Samples. *Political Analysis* (2023).
- [4] Pinar Barlas, Kyriakos Kyriakou, Olivia Guest, Styliani Kleanthous, and Jahna Otterbacher. 2021. To "See" is to Stereotype: Image Tagging Algorithms, Gender Recognition, and the Accuracy-Fairness Trade-off. *Proceedings of the ACM on Human-Computer Interaction (CSCW)* (2021).
- [5] Rachel Bayefsky. 2016. Psychological Harm and Constitutional Standing. *Brooklyn Law Review* (2016).
- [6] Erin Beeghly. 2015. What is a Stereotype? What is Stereotyping? *Hypatia* (2015).
- [7] Erin Beeghly. 2025. What's Wrong with Stereotyping? *Oxford University Press* (2025).
- [8] Jayadev Bhaskaran and Isha Bhallamudi. 2019. Good Secretaries, Bad Truck Drivers? Occupational Gender Stereotypes in Sentiment Analysis. *Proceedings of the First Workshop on Gender Bias in Natural Language Processing* (2019).
- [9] Federico Bianchi, Pratyusha Kalluri, Esin Durmus, Faisal Ladhak, Myra Cheng, Debora Nozza, Tatsunori B Hashimoto, Dan Saul Jurafsky, James Zou, and Aylin Caliskan. 2023. Easily Accessible Text-to-Image Generation Amplifies Demographic Stereotypes at Large Scale. *ACM Conference on Fairness, Accountability, and Transparency (FAccT)* (2023).
- [10] Su Lin Blodgett, Gilsinia Lopez, Alexandra Olteanu, Robert Sim, and Hanna Wallach. 2021. Stereotyping Norwegian Salmon: An Inventory of Pitfalls in Fairness Benchmark Datasets. *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing* (2021).
- [11] Tolga Bolukbasi, Kai-Wei Chang, James Zou, Venkatesh Saligrama, and Adam Kalai. 2016. Man is to Computer Programmer as Woman is to Homemaker? Debiasing Word Embeddings. *Conference on Neural Information Processing Systems (NeurIPS)* (2016).
- [12] C. Malik Boykin, Sophia T. Dasch, Vincent Rice Jr., Venkat R. Lakshminarayanan, Taiwo A. Togun, and Sarah M. Brown. 2021. Opportunities for a More Interdisciplinary Approach to Measuring Perceptions of Fairness in Machine Learning. *Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO)* (2021).
- [13] Diana Burgess and Eugene Borgida. 1999. Who women are, who women should be: descriptive and prescriptive gender stereotyping in sex discrimination. *Psychology, Public Policy, and Law* 5 (1999). Issue 3.
- [14] Judith Butler. 1990. *Gender Trouble: Feminism and the Subversion of Identity*. Routledge (1990).
- [15] Aylin Caliskan, Joanna J. Bryson, and Arvind Narayanan. 2017. Semantics derived automatically from language corpora contain human-like biases. *Science* (2017).
- [16] Yang Cao, Anna Sotnikova, Hal Daumé III, Rachel Rudinger, and Linda Zou. 2022. Theory-Grounded Measurement of U.S. Social Stereotypes in English Language Models. *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (2022).
- [17] Jennifer Chien and David Danks. 2024. Beyond Behaviorist Representational Harms: A Plan for Measurement and Mitigation. *ACM Conference on Fairness, Accountability, and Transparency (FAccT)* (2024).
- [18] John R. Crawford and Julie D. Henry. 2004. The positive and negative affect schedule (PANAS): construct validity, measurement properties and normative data in a large non-clinical sample. *British Journal of Clinical Psychology* (2004).
- [19] Amy J. C. Cuddy, Susan T. Fiske, and Peter Glick. 2007. The BIAS map: behaviors from intergroup affect and stereotypes. *Journal of Personality and Social Psychology* 92 (2007). Issue 4.
- [20] Paula Czarnowska, Yogarshi Vyas, and Kashif Shah. 2021. Quantifying Social Biases in NLP: A Generalization and Empirical Comparison of Extrinsic Fairness Metrics. *Transactions of the Association for Computational Linguistics* (2021).
- [21] Aida Mostafazadeh Davani, Mark Diaz, and Vinodkumar Prabhakaran. 2022. Dealing with Disagreements: Looking Beyond the Majority Vote in Subjective Annotations. *Transactions of the Association for Computational Linguistics* (2022).
- [22] Emily Denton, Mark Diaz, Ian Kivlichen, Vinodkumar Prabhakaran, and Rachel Rosen. 2021. Whose Ground Truth? Accounting for Individual and Collective Identities Underlying Dataset Annotation. *NeurIPS 2021 Workshop on Data-Centric AI* (2021).
- [23] Sunipa Dev, Jaya Goyal, Dinesh Tewari, Shachi Dave, and Vinodkumar Prabhakaran. 2023. Building Socio-culturally Inclusive Stereotype Resources with Community Engagement. *Advances in Neural Information Processing Systems (NeurIPS) Datasets and Benchmarks Track* (2023).
- [24] Sunipa Dev, Tao Li, Jeff Phillips, and Vivek Srikumar. 2020. On Measuring and Mitigating Biased Inferences of Word Embeddings. *AAAI Technical Track: Natural Language Processing* (2020).
- [25] Sunipa Dev and Jeff Phillips. 2019. Attenuating Bias in Word Vectors. *International Conference on Artificial Intelligence and Statistics* (2019).
- [26] Sunipa Dev, Emily Sheng, Jieyu Zhao, Aubrie Amstutz, Jiao Sun, Yu Hou, Mattie Sanseverino, Jjin Kim, Akihiro Nishi, Nanyun Peng, and Kai-Wei Chang. 2022. On Measures of Biases and Harms in NLP. *Findings of the Association for Computational Linguistics: ACL-IJCNLP* (2022).
- [27] Patricia G. Devine and Andrew J. Elliot. 1995. Are Racial Stereotypes Really Fading? The Princeton Trilogy Revisited. *Personality and Social Psychology Bulletin* 21 (1995). Issue 11.
- [28] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *Proceedings of NAACL-HLT* (2019).
- [29] Alejandro M. Diaz. 2008. Through the Google Goggles: Sociopolitical Bias in Search Engine Design. *Information Science and Knowledge Management* (2008).
- [30] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. *International Conference on Learning Representations (ICLR)* (2021).
- [31] Anca Dumitrache, Lora Aroyo, and Chris Welty. 2018. Capturing Ambiguity in Crowdsourcing Frame Disambiguation. *AAAI Conference on Human Computation and Crowdsourcing (HCOMP)* (2018).
- [32] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Rich Zemel. 2012. Fairness Through Awareness. *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference* (2012).
- [33] Alice H. Eagly. 1987. Sex differences in social behavior: A social-role interpretation. *Lawrence Erlbaum Associates, Inc.* (1987).
- [34] Alice H. Eagly, Christa Nater, David I. Miller, Michèle Kaufmann, and Sabine Sczesny. 2020. Gender stereotypes have changed: A cross-temporal meta-analysis of U.S. public opinion polls from 1946 to 2018. *American Psychologist* (2020).
- [35] Naomi Ellemers et al. 2018. Gender stereotypes. *Annual review of psychology* 69 (2018), 275–298.
- [36] Saba Fatima. 2020. I Know What Happened to Me: The Epistemic Harms of Microaggression. *Microaggressions and Philosophy* (2020).
- [37] Susan T. Fiske, Amy J. C. Cuddy, and Peter Glick. 2002. Emotions Up and Down: Intergroup Emotions Result from Status and Competition. *Prejudice to Intergroup Emotions: Differentiated Reactions to Social Groups* (2002).
- [38] Susan T. Fiske, Amy J. C. Cuddy, Peter Glick, and Jun Xu. 2002. A model of (often mixed) stereotype content: Competence and warmth respectively follow from

- perceived status and competition. *Journal of Personality and Social Psychology* 82 (2002). Issue 6.
- [39] Thomas E. Ford. 1997. Effects of Stereotypical Television Portrayals of African-Americans on Person Perception. *Social Psychology Quarterly* (1997).
 - [40] Miranda Fricker. 2009. Epistemic Injustice: Power and the Ethics of Knowing. *Oxford University Press* (2009).
 - [41] Yuki Fujioka. 1999. Television Portrayals and African-American Stereotypes: Examination of Television Effects when Direct Contact is Lacking. *Journalism and Mass Communication Quarterly* (1999).
 - [42] Isabel O. Gallegos, Ryan A. Rossi, Joe Barrow, Md Mehrab Tanjim, Sungchul Kim, Franck Dernoncourt, Tong Yu, Ruiyi Zhang, and Nesreen K. Ahmed. 2023. Bias and Fairness in Large Language Models: A Survey. *arXiv:2309.00770* (2023).
 - [43] Nikhil Garg, Londa Schiebinger, Dan Jurafsky, and James Zhou. 2018. Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences of the United States of America* (PNAS) (2018).
 - [44] Peter Glick and Susan T. Fiske. 1996. The Ambivalent Sexism Inventory: Differentiating hostile and benevolent sexism. *Journal of Personality and Social Psychology* 70 (1996). Issue 3.
 - [45] Peter Glick and Jessica Whitehead. 2010. Hostility Toward Men and the Perceived Stability of Male Dominance. *Social Psychology* 41 (2010). Issue 3.
 - [46] Erving Goffman. 1959. The Presentation of Self in Everyday Life. *Doubleday* (1959).
 - [47] Sherryl Browne Graves. 1999. Television and prejudice reduction: When does television as a vicarious experience make a difference? *Journal of Social Issues* 55 (1999). Issue 4.
 - [48] Anthony G. Greenwald, Debbie E. McGhee, and Jordan L. K. Schwartz. 1998. Measuring individual differences in implicit cognition: the implicit association test. *Journal of Personality and Social Psychology* (1998).
 - [49] Nina Grgić-Hlača, Elissa M. Redmiles, Krishna P. Gummadi, and Adrian Weller. 2018. Human Perceptions of Fairness in Algorithmic Decision Making: A Case Study of Criminal Risk Prediction. *Proceedings of the Web Conference (WWW)* (2018).
 - [50] Douglas Guilbeault, Solène Delecourt, Tasker Hull, Bhargav Srinivasa Desikan, Mark Chu, and Ethan Nadler. 2024. Online images amplify gender bias. *Nature* (2024).
 - [51] Melissa Hall, Laurens van der Maaten, Laura Gustafson, Maxwell Jones, and Aaron Adcock. 2022. A Systematic Study of Bias Amplification. *arXiv:2201.11706* (2022).
 - [52] David L. Hamilton and Jeffrey W. Sherman. 2014. Stereotypes. In *Handbook of social cognition*. Psychology Press, 17–84.
 - [53] Tanja Hentschel, Madeline E. Heilman, and Claudia V. Peus. 2019. The Multiple Dimensions of Gender Stereotypes: A Current Look at Men's and Women's Characterizations of Others and Themselves. *Frontiers in Psychology* (2019).
 - [54] James L. Hilton and William Von Hippel. 1996. Stereotypes. *Annual review of psychology* 47, 1 (1996), 237–271.
 - [55] Perttu Hämäläinen, Mikke Tavast, and Anton Kunnari. 2023. Evaluating Large Language Models in Generating Synthetic HCI Research Data: a Case Study. *Conference on Human Factors in Computing Systems (CHI)* (2023).
 - [56] Lucas D. Introna and Helen Nissenbaum. 2000. Shaping the Web: Why the Politics of Search Engines Matters. *The Information Society* (2000).
 - [57] Joyce Jennings-Walstedt, Florence L. Geis, and Virginia Brown. 1980. Influence of television commercials on women's self-confidence and independent judgment. *Journal of Personality and Social Psychology* (1980).
 - [58] Akshita Jha, Aida Davani, Chandan K. Reddy, Shachi Dave, Vinodkumar Prabhakaran, and Sunipa Dev. 2023. SeeGULL: A Stereotype Benchmark with Broad Geo-Cultural Coverage Leveraging Generative Models. *arXiv:2305.11840* (2023).
 - [59] Akshita Jha, Aida Mostafazadeh Davani, Chandan K Reddy, Shachi Dave, Vinodkumar Prabhakaran, and Sunipa Dev. 2023. SeeGULL: A Stereotype Benchmark with Broad Geo-Cultural Coverage Leveraging Generative Models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (Eds.). Association for Computational Linguistics, Toronto, Canada, 9851–9870. <https://doi.org/10.18653/v1/2023.acl-long.548>
 - [60] Sanjay Kairam and Jeffrey Heer. 2016. Parting Crowds: Characterizing Divergent Interpretations in Crowdsourced Annotation Tasks. *ACM Conference On Computer-Supported Cooperative Work And Social Computing (CSCW)* (2016).
 - [61] Masahiro Kaneko and Danushka Bollegala. 2019. Gender-preserving Debiasing for Pre-trained Word Embeddings. *Annual Conference of the Association for Computational Linguistics (ACL)* (2019).
 - [62] Saket Karve, Lyle Ungar, and João Sedoc. 2019. Conceptor Debiasing of Word Representations Evaluated on WEAT. *arXiv:1906.05993* (2019).
 - [63] Daniel Katz and Kenneth Braly. 1933. Racial Stereotypes of One Hundred College Students. *The Journal of Abnormal and Social Psychology* 28 (1933). Issue 3.
 - [64] Jared Katzman, Angelina Wang, Morgan Scheuerman, Su Lin Blodgett, Kristen Laird, Hanna Wallach, and Solon Barocas. 2023. Taxonomizing and Measuring Representational Harms: A Look at Image Tagging. *AAAI Conference on Artificial Intelligence* (2023).
 - [65] Matthew Kay, Cynthia Matuszek, and Sean A. Munson. 2015. Unequal Representation and Gender Stereotypes in Image Search Results for Occupations. *Conference on Human Factors in Computing Systems (CHI)* (2015).
 - [66] Matjaž Kukar and Igor Kononenko. 1998. Cost-Sensitive Learning with Neural Networks. *European Conference on Artificial Intelligence* (1998).
 - [67] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Shahab Kamali, Jordi Pont-Tuset, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, Tom Duerig, and Vittorio Ferrari. 2020. The Open Images Dataset V4: Unified image classification, object detection, and visual relationship detection at scale. *International Journal of Computer Vision (IJCV)* (2020).
 - [68] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft COCO: Common objects in context. *European Conference on Computer Vision (ECCV)* (2014).
 - [69] Walter Lippmann. 1922. Public opinion. *MacMillan Co* (1922).
 - [70] Mercedes López-Sáez and Ana Lisbona. 2014. Descriptive and prescriptive features of gender stereotyping. Relationships among its components. *International Journal of Social Psychology* 24 (2014). Issue 3.
 - [71] Arti P. Makwana, Kristof Dhont, Jonas De keersmaecker, Parisa Akhlaghi-Ghaffarokh, Marine Masure, and Arne Roets. 2018. The Motivated Cognitive Basis of Transphobia: The Roles of Right-Wing Ideologies and Gender Role Beliefs. *Sex Roles* 79 (2018).
 - [72] Thomas Manzini, Yao Chong Lim, Yulia Tsvetkov, and Alan W. Black. 2019. Black is to Criminal as Caucasian is to Police: Detecting and Removing Multiclass Bias in Word Embeddings. *Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)* (2019).
 - [73] Danaë Metaxa, Michelle A. Gan, Su Goh, Jeff Hancock, and James A. Landay. 2021. An Image of Society: Gender and Racial Representation and Impact in Image Search Results for Occupations. *ACM Conference on Human-Computer Interaction (CSCW)* (2021).
 - [74] Thekla Morgenroth and Michelle K. Ryan. 2018. Gender Trouble in Social Psychology: How Can Butler's Work Inform Experimental Social Psychologists' Conceptualization of Gender? *Frontiers in Psychology* (2018).
 - [75] Thekla Morgenroth and Michelle K. Ryan. 2020. The Effects of Gender Trouble: An Integrative Theoretical Framework of the Perpetuation and Disruption of the Gender/Sex Binary. *Perspectives on Psychological Science* 16 (2020).
 - [76] Craig T. Nagoshi, J. Raven Cloud, Louis M. Lindley, Julie L. Nagoshi, and Lucas J. Lothamer. 2019. A Test of the Three-Component Model of Gender-Based Prejudices: Homophobia and Transphobia Are Affected by Raters' and Targets' Assigned Sex at Birth. *Sex Roles* 80 (2019).
 - [77] Nikita Nangia, Clara Vania, Rasika Bhalerao, and Samuel R. Bowman. 2020. CrowS-Pairs: A Challenge Dataset for Measuring Social Biases in Masked Language Models. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu (Eds.). Association for Computational Linguistics, Online, 1953–1967. <https://doi.org/10.18653/v1/2020.emnlp-main.154>
 - [78] Gandalf Nicolas, Xuechunzi Bai, and Susan T Fiske. 2022. A spontaneous stereotype content model: Taxonomy, properties, and prediction. *Journal of personality and social psychology* (2022).
 - [79] Jennifer A. Noble. 2012. Minority voices of crowdsourcing: why we should pay attention to every member of the crowd. *ACM Conference On Computer-Supported Cooperative Work And Social Computing (CSCW)* (2012).
 - [80] Safiya Umoja Noble. 2018. Algorithms of Oppression: How Search Engines Reinforce Racism. *NYU Press* (2018).
 - [81] Omaiith O'Dowd. 2018. Microaggressions: A Kantian Account. *Ethical Theory and Moral Practice* 21 (2018).
 - [82] Jahna Otterbacher, Jo Bates, and Paul Clough. 2017. Competent Men and Warm Women: Gender Stereotypes and Backlash in Image Search Results. *Conference on Human Factors in Computing Systems (CHI)* (2017).
 - [83] Jahna Otterbacher, Alessandro Checco, Gianluca Demartini, and Paul Clough. 2018. Investigating User Perception of Gender Bias in Image Search: The Role of Sexism. *ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)* (2018).
 - [84] Alicia Parrish, Angelica Chen, Nikita Nangia, Vishakh Padmakumar, Jason Phang, Jana Thompson, Phu Mon Htut, and Samuel Bowman. 2022. BBQ: A hand-built bias benchmark for question answering. In *Findings of the Association for Computational Linguistics: ACL 2022*, Smaranda Muresan, Preslav Nakov, and Aline Villavicencio (Eds.). Association for Computational Linguistics, Dublin, Ireland, 2086–2105. <https://doi.org/10.18653/v1/2022.findings-acl.165>
 - [85] Przemysław Palka. 2023. AI, Consumers & Psychological Harm. *Cambridge University Press* (2023).
 - [86] Shauli Ravfogel, Yanai Elazar, Hila Gonen, Michael Twiton, and Yoav Goldberg. 2020. Null It Out: Guarding Protected Attributes by Iterative Nullspace Projection. *Annual Conference of the Association for Computational Linguistics (ACL)* (2020).
 - [87] Regina Rini. 2020. The Ethics of Microaggression. *Routledge Taylor & Francis Group* (2020).
 - [88] Chiara Rollero, Peter Glick, and Stefano Tartaglia. 2014. Psychometric properties of short versions of the Ambivalent Sexism Inventory and Ambivalence Toward

- Men Inventory. *TPM-Testing, Psychometrics, Methodology in Applied Psychology* 21 (2014). Issue 2.
- [89] Laurie A. Rudman, Corinne A. Moss-Racusin, Peter Glick, and Julie E. Phelan. 2012. Reactions to Vanguard: Advances in Backlash Theory. *Advances in experimental social psychology* 45 (2012).
- [90] Laurie A. Rudman, Corinne A. Moss-Racusin, Julie E. Phelan, and Sanne Nauts. 2012. Status incongruity and backlash effects: Defending the gender hierarchy motivates prejudice against female leaders. *Journal of Experimental Social Psychology* 48 (2012).
- [91] Beatrice Savoldi, Marco Gaido, Luisa Bentivogli, Matteo Negri, and Marco Turchi. 2021. Gender Bias in Machine Translation. *Transactions of the Association for Computational Linguistics* 9 (2021), 845–874. https://doi.org/10.1162/tacl_a_00401
- [92] Candice Schumann, Susanna Ricco, Utsav Prabhu, Vittorio Ferrari, and Caroline Pantofaru. 2021. A Step Toward More Inclusive People Annotations for Fairness. *ACM Conference on Artificial Intelligence, Ethics, and Society (AIES)* (2021).
- [93] Charles R. Seger, Ishani Banerji, Sang Hee Park, Eliot R. Smith, and Diane M. Mackie. 2017. Specific emotions as mediators of the effect of intergroup contact on prejudice: findings across multiple participant and target groups. *Cognition and Emotion* 31 (2017). Issue 5.
- [94] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*. 618–626.
- [95] Seungjae Shin, Kyungwoo Song, JoonHo Jang, Hyemi Kim, Weonyoung Joo, and Il-Chul Moon. 2020. Neutralizing Gender Bias in Word Embedding with Latent Disentanglement and Counterfactual Generation. *Findings of EMNLP* (2020).
- [96] Nurit Shnabel, Yoav Bar-Anan, Anna Kende, Orly Bareket, and Yael Lazar. 2016. Help to perpetuate traditional gender roles: Benevolent sexism increases engagement in dependency-oriented cross-gender helping. *Journal of personality and social psychology* 110, 1 (2016), 55.
- [97] Tom Simonite. 2018. When It Comes to Gorillas, Google Photos Remains Blind. *Wired, January* (2018).
- [98] Anna Sotnikova, Yang Trista Cao, Hal Daumé III, and Rachel Rudinger. 2021. Analyzing Stereotypes in Generative Text Inference Tasks. *Findings of the Association for Computational Linguistics: ACL-IJCNLP* (2021).
- [99] Steven J. Spencer, Christine Logel, and Paul G. Davies. 2015. Stereotype Threat. *Annual Review of Psychology* 67 (2015).
- [100] Megha Srivastava, Hoda Heidari, and Andreas Krause. 2019. Mathematical Notions vs. Human Perception of Fairness: A Descriptive Approach to Fairness for Machine Learning. *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)* (2019).
- [101] Claude M. Steele and Joshua Aronson. 1995. Stereotype threat and the intellectual test performance of African Americans. *Journal of Personality and Social Psychology* 69 (1995). Issue 5.
- [102] Chadly Stern and Nicholas O. Rule. 2017. Physical Androgyny and Categorization Difficulty Shape Political Conservatives' Attitudes Toward Transgender People. *Social Psychological and Personality Science* (2017).
- [103] Emiel van Miltenburg. 2016. Stereotyping and Bias in the Flickr30K Dataset. *Proceedings of the Workshop on Multimodal Corpora* (2016).
- [104] Joseph A. Vandello, Jennifer K. Bosson, Dov Cohen, Burnaford Rochelle M, and Jonathan R. Weaver. 2008. Precarious manhood. *Journal of Personality and Social Psychology* (2008).
- [105] Madalinea Vlasceanu and David M. Amodio. 2022. Propagation of societal gender inequality by internet search algorithms. *Proceedings of the National Academy of Sciences of the United States of America (PNAS)* 119 (2022).
- [106] Yixin Wan, George Pu, Jiao Sun, Aparna Garimella, Kai-Wei Chang, and Nanyun Peng. 2023. "Kelly is a Warm Person, Joseph is a Role Model": Gender Biases in LLM-Generated Reference Letters. *Conference on Empirical Methods in Natural Language Processing (EMNLP Findings)* (2023).
- [107] Angelina Wang, Alexander Liu, Ryan Zhang, Anat Kleiman, Leslie Kim, Dora Zhao, Iroha Shirai, Arvind Narayanan, and Olga Russakovsky. 2022. REVISE: A Tool for Measuring and Mitigating Bias in Visual Datasets. *International Journal of Computer Vision (IJCV)* (2022).
- [108] Angelina Wang and Olga Russakovsky. 2021. Directional Bias Amplification. *International Conference on Machine Learning (ICML)* (2021).
- [109] Clarice Wang, Kathryn Wang, Andrew Bian, Rashidul Islam, Kamrun Naher Keya, James Foulds, and Shimei Pan. 2021. User Acceptance of Gender Stereotypes in Automated Career Recommendations. *arXiv:2106.07112* (2021).
- [110] Tianlu Wang, Jieyu Zhao, Mark Yatskar, Kai-Wei Chang, and Vicente Ordonez. 2019. Balanced Datasets Are Not Enough: Estimating and Mitigating Gender Bias in Deep Image Representations. *International Conference on Computer Vision (ICCV)* (2019).
- [111] Zeerak Waseem. 2016. Are You a Racist or Am I Seeing Things? Annotator Influence on Hate Speech Detection on Twitter. *Proceedings of the First Workshop on NLP and Computational Social Science* (2016).
- [112] David Watson, Lee Anna Clark, and Auke Tellegen. 1988. Development and validation of brief measures of positive and negative affect: the PANAS scales. *Journal of Personality and Social Psychology* 54 (1988).
- [113] John E. Williams and Deborah L. Best. 1977. Sex Stereotypes and Trait Favorability on the Adjective Check List. *Educational and Psychological Measurement* 37 (1977). Issue 1.
- [114] Allison Woodruff, Sarah E. Fox, Steven Rousso-Schindler, and Jeffrey Warshaw. 2018. A Qualitative Exploration of Perceptions of Algorithmic Fairness. *Conference on Human Factors in Computing Systems (CHI)* (2018).
- [115] Alison Wylie. 2003. Why Standpoint Matters. *Science and Other Cultures: Issues in Philosophies of Science and Technology* (2003).
- [116] Mohammad Yaghini, Andreas Krause, and Hoda Heidari. 2021. A Human-in-the-loop Framework to Construct Context-aware Mathematical Notions of Outcome Fairness. *AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society* (2021).
- [117] Dora Zhao, Angelina Wang, and Olga Russakovsky. 2021. Understanding and Evaluating Racial Biases in Image Captioning. *International Conference on Computer Vision (ICCV)* (2021).
- [118] Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. 2017. Men also like shopping: Reducing gender bias amplification using corpus-level constraints. *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)* (2017).
- [119] Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. 2018. Gender Bias in Coreference Resolution: Evaluation and Debiasing Methods. *North American Chapter of the Association for Computational Linguistics* (2018).
- [120] Jieyu Zhao, Yichao Zhou, Zeyu Li, Wei Wang, and Kai-Wei Chang. 2018. Learning Gender-Neutral Word Embeddings. *Empirical Methods in Natural Language Processing (EMNLP)* (2018).

A Methods Details

Here we present additional details on our methods.

A.1 Participants

While men and women generally tend to hold the same gender stereotypes [34, 53, 70, 113], we still collect equal numbers of participants who identify as men and women, and use this variable as a covariate throughout. Due to limitations in the survey platform which only allow us to specify gender as “male” or “female,” this formulation excludes people who identify as non-binary, which is a harmful limitation. Because we do not control for race in the recruitment of participants, our sample diverges from a nationally representative sample. For the gender stereotype scope of our current work, we find this to be an acceptable limitation, especially given that one defining feature of stereotypes is they are largely shared through a cultural consensus [63].

We did not use quality check questions in any of our surveys, because our pilot studies showed high quality responses. Instead, we used filters on Cloud Research to only recruit participants who have had at least 50 HITs approved, and have a HIT approval rate of 98%.

A.2 Studies 1, 4: Distinguishing Errors by Stereotype

When asking about which machine learning errors are stereotypes, we make sure to ask participants about their perception of stereotypes held by Americans, rather than for their personal beliefs [27]. The Study 1 text is “In this survey, you will answer questions about 20 objects and their associated stereotypes. A stereotype is an over-generalized belief about a group of people. Please note that we are not interested in your personal beliefs, but in how you think these groups are viewed by American Society. For each object, please indicate which demographic of men or women are stereotypically associated with the object. If no group is, only select the answer

choice for “neither.” If both groups are, then select both groups. As an example, ping pong balls are sometimes stereotypically associated with Asian people.”

A.3 Study 3a: Measuring Pragmatic Harm

We conduct a between-subjects survey experiment on participants who are exposed to an image search result page that contain one of three types of errors: stereotype-reinforcing, stereotype-violating, or neutral (Fig. 5).³ To have the participants engage with these results we ask them to describe it in 3-4 sentences. Next, we ask them the behavior questions, then re-expose them to the stimulus before asking them the cognitive belief and attitude questions. We analyze changes in cognitive beliefs, attitudes, and behaviors as pragmatic harms resulting from stereotype-reinforcing errors compared to the two other conditions as controls. In this section when describing our method, we will use as examples oven and women for the stereotype-reinforcing error, oven and men for the stereotype-violating error, and bowl and women for the neutral one. Each question we ask is carefully grounded in the social psychology literature.

The stimuli take the form of an image search result and are pictured in Fig. 5 with teal and orange colored boxes around the component of the image that changes between conditions. The search bar contains the search query, and then eight images that may or may not be correctly retrieved are shown. Each of the eight images is annotated with either “In image” or “Not in image” to make it clear to the participant which images are correct or not. The stereotype-reinforcing condition on the left contains the search query of “oven” with five correctly identified ovens, and three false positive images that all contain women. In other words, this classifier erroneously (and stereotypically) assumes there are ovens in images of women. The stereotype-violating condition contains the same search query, but the mistakes are replaced with false positive images that all contain men. The neutral condition contains all of the exact same images as the stereotype-reinforcing condition, with the only change being that the search query is now “bowl” instead of “oven.” This is because the five correct images were deliberately chosen to contain both bowls and ovens, which allows us to control for the variance between the different search conditions. All false positive images were selected from the actual errors of a Vision Transformer (ViT) model [30] trained on COCO so that they are as realistic as possible to a computer vision model’s errors, and not completely egregious, e.g., a picture of a woman in a sports field as a false positive for “oven” or “bowl.”

A.3.1 Dependent Variables. For most of our measurements, we simply use the measure directly (e.g., the value for competence of women) as the dependent variable to regress on. For the measurements that we do something more complicated, we describe below.

Behavior - Tags. Each participant produces a set of three ordered tags associated with an image of a feminine-presenting person and a set associated with a counterfactual image of a masculine-presenting person. We convert this set of tags by scoring the presence of the object in question, e.g., “hair dryer” (along with common

misspellings such as “hair drier”) based on its position in the ordered list of tags. When the word is present in the first spot it is given 3 points, second spot 2 points, third spot 1 point, otherwise no points. The dependent variable is the score of both the stereotypical and neutral object on the feminine-presenting person. This is intended to capture whether the stereotype-reinforcing condition is able to increase the presence of the stereotype tag more than just the priming effect captured by the neutral object.

Behavior - Captions. We offer some descriptive statistics about the captions in Appendix E. This analysis was mostly exploratory, and we do not find any statistically significant differences. We first ran Study 3a looking at pragmatic harms on the stereotype of women and oven (with bowl as the control). In this iteration, we asked that respondents please describe each person in the image in separate sentences. However, there was too much noise in how respondents interpreted this set of instructions, such that the data became hard to interpret. Thus, in our second iteration of this study using the stereotype of women and hair dryer (with toothbrush as the control), we have two separate text entry boxes to caption each person in the image. We only present the results of this iteration in the table, as we were unable to parse anything differentiating in the first iteration.

Cognitive - Object Use. In this measurement, we have a value from -10 (mostly men) to 10 (mostly women) for both the stereotypical and neutral object. The dependent variable is the summation of both values. Again, this is intended to capture whether the stereotype-reinforcing condition is able to change the value of its associated object more than the control condition is able to.

A.4 Study 3b, 4: Measuring Experimental Harm

In Study 3b, in addition to personal discomfort, we also ask about societal harm. This way, even if the participant does not personally feel harmed, they may feel it on behalf of the stereotyped group. However, we find that participants’ responses to both personal and societal harm are extremely correlated, and leave the results for the latter in Appendix F.

B Sample Size Justification

Our sample size selection method is recorded on Open Science Framework and is done as follows: Study 1 we selected the number such that each COCO object was labeled by 10 participants from each gender; once we saw there was sufficient consensus from Study 1, for the first part of Study 4 where we labeled OpenImages objects, we selected the number such that each OpenImages object was labeled by 5 participants from each gender; Study 3a we had three stimulus conditions across two objects, so for this between-subjects study selected the number to have 50 participants from each gender for each object-condition setting; Study 3b we had three stimulus conditions across four objects, but this is a within-subjects study so each participant sees all possible scenarios, and thus we again selected the number to have 50 participants from each gender; Study 4 we had 40 objects and as our last study ended up having the budget to have around 37.5 participants per object.

³The people pictured in our search results pages are predominantly White, which is the majority group in the dataset we employ.

Table 1: The time, pay, and reported races of the participants for each of our five studies. The full column names of races from left to right are: American Indian or Alaska Native, Asian, Black or African American, Hispanic or Latinx, Native Hawaiian or Other Pacific Islander, White, Multi-Racial / Other, and Prefer not to say.

Study	Time (min)	Pay (\$)	Gender	AI/AN	Asian	Black	H/L	NHOPI	White	MR/O	PNTS	Total
1 and 2	7	1.75	Women	0	3	5	0	0	25	6	1	40
			Men	1	4	2	2	0	30	1	0	40
3a	10	2.50	Women	1	11	32	8	0	229	19	0	300
			Men	0	19	35	10	1	211	22	2	300
3b	5	1.25	Women	0	4	7	3	1	35	5	0	50
			Men	0	4	2	3	1	35	5	0	50
4 (Labeling)	4	1	Women	0	5	8	0	0	42	4	1	60
			Men	0	2	6	5	1	44	2	0	60
4 (Harms)	5	1.25	Women	0	5	15	1	0	120	7	2	150
			Men	1	9	17	6	1	107	9	0	150

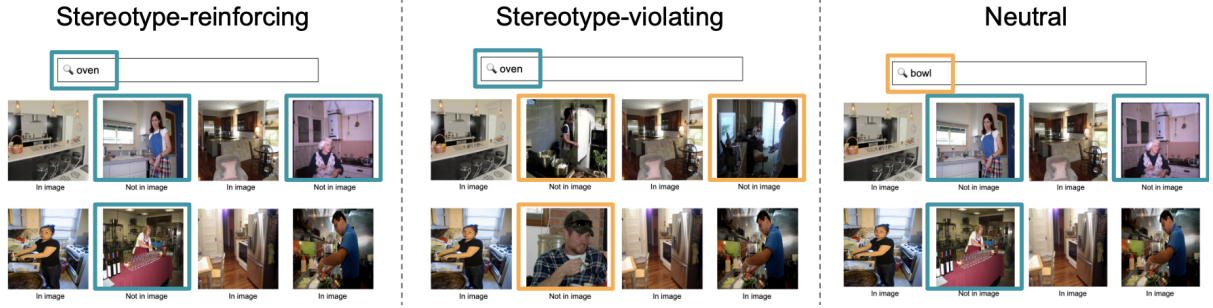


Figure 5: Study 3 Stimuli. Our three different stimuli are shown for the conditions: stereotype-reinforcing, stereotype-violating, and neutral. They are all image search results containing minimal changes from each other, each of which indicates whether the search query is pictured in the image, i.e., if the image search retrieval was correct or not. The teal and orange squares indicate that the only difference between the stimuli, as all images which contain an oven also contain a bowl, and all which do not contain an oven also do not contain a bowl. This was a deliberate choice to control for all potential confounding factors from the images in the study.

C Additional Results from Study 1

We show the results of harm annotation in the abstract on the y-axis of Fig. 7 for the 13 objects marked as stereotypes. We see large variations within stereotypical objects for whether the association is perceived to be harmful when it is disconnected from a particular impact.

D Additional Results from Study 2

Here we present the full analyses we perform on the open-ended responses we received in Study 2 regarding why participants believe an object is a stereotype, and if so, why they find it harmful or not.

Our categorization for why an object is a stereotype or not are as follows (some responses did not fall into any of the categories):

- Descriptive (45%), e.g., for handbag and women: “women are often seen wearing handbags and buying them”
- Occupation/role (22%), e.g., for oven and women: “Women are stereotyped to always be in the kitchen cooking while the men go out and work”

- Trait (11%), e.g., for chair and men: “sometimes men would be seen as coming home and just being lazy and lounging in their chair”
- Pop culture (8%), e.g., for cow and women: “Most people who describe a women as a cow are being harmful and hurtful”
- Connection to another object (5%), e.g., for vase and women: “I think women are seen as liking flowers, which are often put into a vase”
- Prescriptive (3%), e.g., for handbag and women: “society generally believes that only women should carry handbags”

Our categorization for why a stereotypical object is harmful is as follows (some responses did not fall into any of the categories):

- Proscriptive, i.e., excluding (40%), e.g., for dining table and women: “it makes it looked down upon if a man cooks dinner.”
- Prescriptive, i.e., restricting (26%), e.g., for dining table and women: “I think it puts women in a box that says they must prepare dinner”




Figure 6: To measure behavioral tendencies, we ask participants to complete a realistic data annotation task on images which are created and manipulated by DALL-E2. The left pair is for the annotation of image tags, and the right pair is for image captions. Each participant is shown one image from each pair, and then we perform a between-subjects analysis to understand whether perceived gender expression affects the tags, and whether object shown influences how people of different perceived genders are described.

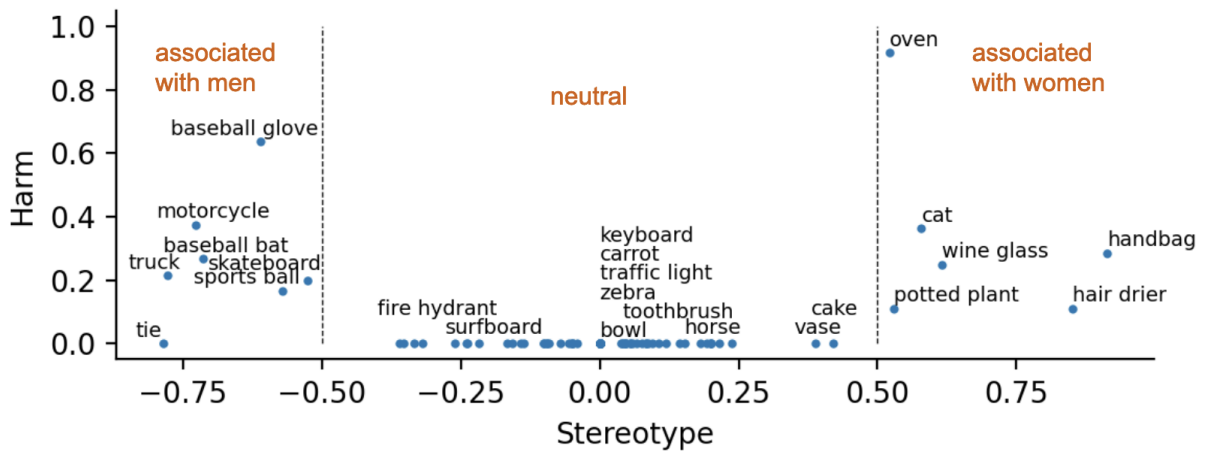


Figure 7: Study 1 Results. Participant responses for 80 objects in COCO dataset. The x-axis indicates the percentage of participants who indicated an object is a stereotype, where negative numbers indicate it is a stereotype about men, and positive numbers about women. For objects where more than half of the respondents indicate it is a stereotype, the y-axis indicates the percentage who marked it to be harmful.

- Negative Trait (13%), e.g., for handbag and women: “It is harmful because it implies that women cares more about looks and their appearance.”
- Demeans (10%), e.g., for cow and women: “cow is a typical insult for a women a man doesn’t like (‘stupid cow’)”
- Objectifies (5%), e.g., for cup and women: “It is harmful because a cup is an object and it’s comparing it to a woman”
- Sexism (3%), e.g., for sandwich and women: “make me a sandwich meme sexism”
- Incorrect (3%), e.g., for sandwich and women: “It’s an old, tired stereotype that holds no merit.”

The following are the two reasons respondents listed a stereotype to not be harmful: not negative (96%), e.g., for tie and men: “I don’t think it’s harmful because it’s just a fashion choice”; positive stereotype (4%), e.g., for cake and women: “cake can be used to

describe a woman as sweet and nice looking. For that reason I don’t find it harmful.”

E Behavior Caption Analysis from Study 3a

Here we present an exploratory analysis of the captions produced by participants for the behavior task in Study 3a. In Tbl. 2 we compare the captions generated by participants across conditions, and find no statistically significant results.

F Additional Results from Study 3b

For Study 3b not only did we ask for personal experiential harm from an error, but also societal harm, so that even if the participant does not personally feel harmed, they may feel it on behalf of the stereotyped group. There is a high correlation between the responses to these two versions of the question for each error,

Table 2: Descriptive statistics about the captions annotated as a part of Study 3a’s behavior measure for the stereotype of women and hair dryer, where toothbrush serves as the control neutral object.

Condition	Mention of Hair Dryer / Mention of Toothbrush		Warmth		Competence	
	Women	Men	Women	Men	Women	Men
Stereotype-reinforcing	1.095 (0.679-1.800)	1.125 (0.679-1.800)	6.750 ± 0.386	1.107 ± 0.327	1.222 ± 0.263	0.933 ± 0.225
Stereotype-violating	.810 (0.458-1.286)	.857 (0.222-2.500)	1.182 ± 0.417	0.905 ± 0.384	0.632 ± 0.265	0.750 ± 0.248
Neutral	0.913 (0.562-1.450)	0.571 (0.100-1.750)	1.300 ± 0.425	0.769 ± 0.300	0.783 ± 0.293	1.182 ± 0.229

which is why we only reported results from personal harm in the main text.

We regress on harm value with the independent variables of stereotype condition, personal harm compared to societal harm, and their interaction effect. We find that for stereotype-reinforcing vs neutral errors, the coefficient for personal or societal harm is $b = .22$ (95% *CI* [-.07, .51], $p = .143$), and that for the interaction effect to be $b = -.18$ (95% *CI* [-.59, .23], $p = .398$). For stereotype-reinforcing vs stereotype-violating errors, the coefficient for personal or societal harm is $b = .26$ (95% *CI* [-.03, .55], $p = .076$), and that for the interaction effect to be $b = -.22$ (95% *CI* [-.63, .19], $p = .288$). In other words, there is no statistically significant difference in the results for societal harm compared to personal harm.

G Object Selection from OpenImages

In OpenImages, of the 600 objects, we select the 20 marked with the most agreement to be stereotypically associated with men, and the 20 marked with the most agreement to be stereotypically associated with women. We then randomly select amongst 20 objects that are marked to have no gender stereotypes associated with them. Left out of this are all human-related categories: *boy, girl, human eye, human face, human body, human ear, human arm, human board, human hand*; as well as *invertebrate* because there was confusion amongst pilot testers about what this word meant.

The 20 objects stereotyped about men are: *football helmet, football, cowboy hat, hammer, sports equipment, jet ski, truck, tie, golf ball, beer, skateboard, briefcase, plumbing fixture, tire, wrench, suit, missile, tool, rifle, shotgun*. The 4 that we consider “clothing,” i.e., able to be worn, are *football helmet, cowboy hat, tie, suit*.

The 20 objects stereotyped about women are: *ladybug, doll, hair spray, lily, hair dryer, perfume, kitchenware, cat, wine glass, fashion accessory, necklace, flower, handbag, lipstick, bathtub, face powder, cosmetics, rose, oven, brassiere*. The 9 that we consider “clothing,” i.e., able to be worn, are *necklace, face powder, fashion accessory, lipstick, brassiere, cosmetics, hair spray, handbag, perfume*.

The 20 neutral objects are: *pillow, owl, giraffe, balloon, jellyfish, stop sign, french fries, eraser, shower, orange, chopsticks, window, personal flotation device, bed, goldfish, zebra, raccoon, sea lion, microphone, popcorn*.

H Bias Amplification

Bias amplification is a statistical notion that rests on the idea that any amplification of an existing bias is undesirable, and often used to implicitly capture stereotypes [51, 108, 110, 118]. In this line of work, a “bias” is measured in the dataset, e.g., that women are correlated with object A, and so any amplification of this in the model’s test-time predictions is considered undesirable, and likely the application of something like a stereotype. This “bias” is determined statistically, and two possible formulations come from Zhao et al. [118] (Bias Amp) and Wang and Russakovsky [108] (Directional Bias Amp). As an example, Zhao et al. [118] measures oven, wine glass, and potted plant, to all be biased towards men. From our human annotations, we find all these of these objects to be biased towards women. Thus, mitigation algorithms directed at reducing either of these formulations of bias amplification would actually likely *increase* certain types of harmful errors in an attempt to reduce overall bias amplification. This formulation also assumes that every label is biased in a way such that one direction of error is worse than another, missing that many labels can be neutral in certain respects, e.g., bowl and table.

We quantify two aspects of each bias amplification metric, which are its abilities to identify either objects as stereotypes (measured by calculating the percentage overlap between the top- n “biased” objects and n stereotypes) or the gender direction of the stereotype’s alignment (measured by calculating the gender direction on the n stereotyped objects). In Tbl. 3 we can see that while both bias amplification metrics are able to approximate the gender that a stereotyped object is correlated with in the COCO dataset reasonably well, this is not true for identifying which objects are stereotypes, nor the gender alignment in the larger OpenImages dataset. Thus, attempts to reduce either metric of bias amplification are likely to inadvertently increase the number of stereotypical errors in an attempt to reduce a “bias amplification” error that may not actually be stereotypically harmful.

I Automatic Discovery of Stereotypes

Evaluation is sometimes considered secondary to algorithm development, and thus rapid and fully-automated evaluations are often prioritized over those requiring human input. Thus, one might imagine trying to automate the determination of which labels are

Table 3: Comparison of how correlated to human perceptions prior measures of bias amplification, which approximate the directions of bias that are more harmful, are.

Dataset	Metric	Pearson R	Identification of Stereotypes	Alignment of Stereotypes
COCO	Bias Amp [118]	.5722	7/13 (54%)	10/13 (77%)
	Directional Bias Amp [108]	.6507	6/13 (46%)	13/13 (100%)
OpenImages	Bias Amp	.3912	124/249 (50%)	141/249 (57%)
	Directional Bias Amp	.1502	120/249 (48%)	153/249 (61%)

stereotypes, rather than soliciting judgments from human annotators. To test the limits of this approach, we train a variety of models (Support Vector Machine, Random Forest, and Multi-Layer Perceptron) with hyperparameter search over the number of features and find the highest ROC AUC for predicting whether an OpenImages object is a stereotype given an input of BERT word embeddings [28] to be 74%. Erroneous predictions include carnivore (is stereotype), tennis ball (is not stereotype), infant bed (is stereotype), soap dispenser (is not stereotype), handbag (is stereotype). Given that an object is a stereotype, the highest ROC AUC at predicting which gender is being stereotyped is 85%. Erroneous predictions include mixer (stereotyped with women), doughnut (stereotyped with men), houseplant (stereotyped with women), wheel (stereotyped with men). These inadequate performance rates indicate that stereotypes are highly contextual, and even with the use of powerful word embeddings which capture bias and social context [43], they are insufficient without human input. As we note in the main text, even if the growing power of large language models enables us to predict with higher accuracy which objects are stereotypes, we likely still may want to ensure these annotations come from people themselves [3, 55], thus allowing room for positionality, explanation, and critical reflection.

space [24, 25, 61, 62, 72, 86, 95, 120], all eight of these works would, as far as we can tell, attempt just as much to debias words like “table” and “apple” as they would “homemaker” and “doll.” While it is not clear what exactly is the desired state of debiasing (e.g., describing the world as it is, prescribing the world as it ought to be, aligning with people’s existing stereotypic expectations [109]) it surely seems that more thinking should be done on the different implications of debiasing stereotypes as opposed to debiasing more neutral concepts.

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009

J Connection to Open-Ended ML Tasks without Correctness

We scoped our work to machine learning tasks which have a clear notion of error, i.e., ground-truth labels. Here, we consider the implications of our findings for other machine learning tasks which do not have such a clear notion of an error, for example in text generation. We can also consider the implications of our findings for other machine learning tasks which do not have such a clear notion of an error, for example in text generation. Prior work brought to light that word embeddings mirror stereotypes in our society [11, 15], such as about occupations and attributes from the Implicit Association Test [48]. Since then, most follow-up work in this space seeks to remove nearly all gendered associations in text, conflating each such association with harmful “bias.” Again we see a similar pattern to the logical fallacy of the harm of one type of error, e.g., a correlation of some set of stereotypical occupations to gender, extending to *all* errors. The nuance is lost when gendered associations of all kinds in word embeddings are equated to stereotypes, and most notions of gender are targeted to be removed from the embeddings. To put this into perspective, in the large body of literature that has followed the discovery of gender biases in the embedding