

Investigating Labeler Bias in Face Annotation for Machine Learning

Luke Haliburton
LMU Munich
Munich, Germany
luke.haliburton@ifi.lmu.de

Sinksar Ghebremedhin
LMU Munich
Munich, Germany

Robin Welsch
Aalto University
Espoo, Finland

Albrecht Schmidt
LMU Munich
Munich, Germany

Sven Mayer
LMU Munich
Munich, Germany

ABSTRACT

In a world increasingly reliant on artificial intelligence, it is more important than ever to consider the ethical implications of artificial intelligence on humanity. One key under-explored challenge is labeler bias, which can create inherently biased datasets for training and subsequently lead to inaccurate or unfair decisions in health-care, employment, education, and law enforcement. Hence, we conducted a study to investigate and measure the existence of labeler bias using images of people from different ethnicities and sexes in a labeling task. Our results show that participants hold stereotypes that influence their decision-making process and that labeler demographics impact assigned labels. We also discuss how labeler bias influences datasets and, subsequently, the models trained on them. Overall, a high degree of transparency must be maintained throughout the entire artificial intelligence training process to identify and correct biases in the data as early as possible.

CCS CONCEPTS

• **Human-centered computing** → **Human computer interaction (HCI)**.

KEYWORDS

human-computer interaction, machine learning, annotation, bias

ACM Reference Format:

Luke Haliburton, Sinksar Ghebremedhin, Robin Welsch, Albrecht Schmidt, and Sven Mayer. 2023. Investigating Labeler Bias in Face Annotation for Machine Learning. In *arXiv '23, June 23, 2023*. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.48550/arXiv.2301.09902>

1 INTRODUCTION

Artificial intelligence (AI) is rapidly becoming involved in numerous areas of life, making far-reaching decisions such as granting loans and hiring people. Amazon analyzes customers' purchasing

behavior¹, Netflix studies entertainment preferences², and Facebook uses social interactions to tailor content to their users [28]. Data collection, processing, and prediction are key pillars of AI applications. Although AI is a powerful tool, the fundamental reliance on data can be problematic since datasets can be distorted in various ways, creating unintended consequences. One under-investigated contributing factor to biased AI tools is labeler bias, which results from cognitive biases [14] in crowd workers and other dynamics in the labeling process [44]. Many AI applications rely on crowdsourcing platforms to label their data, yet they usually do not consider whether they are utilizing a diverse population of labelers [43]. A biased labeler pool could lead to unfair outcomes for certain groups, such as women, ethnic minorities, or people from disadvantaged neighborhoods. Therefore, it is crucial to examine labeler pools with a critical lens to avoid bias and create a more fair and transparent process.

Investigating labeler bias is essential to understand how labelers influence datasets, but existing studies in this area are limited in scope. Recent work has demonstrated that rater identity plays a significant role in labeling toxicity for online comments [24, 50]. One critical paper by Bender et al. [5] sheds light on how human biases can be unintentionally perpetuated in machine learning, highlighting that biases introduced in the labeling stage can propagate through to end decisions. In response to this issue, researchers in machine learning have attempted to model and correct for bias effects [6, 32, 57]. In general, bias can be partially attributed to stereotypes, which occur when one assigns traits to an individual based on preconceived notions about their group [17]. The Stereotype Content Model (SCM) is an established practical theory explaining stereotypes that has been applied in Human-Computer Interaction (HCI) (e.g., [40, 53]). However, there is a lack of work applying the SCM to characterize biases introduced into datasets by crowdsourced labelers.

In this paper, we address the gap in the existing literature by investigating stereotypes and bias in labeling tasks. We conducted a survey (n=98) asking participants to label a series of human faces from the FairFace dataset [33]. We selected faces with equal representation from seven ethnicities and two sexes and recruited labelers with the same balanced demographic distribution. We asked

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
arXiv '23, June 23, 2023.

© 2023 Association for Computing Machinery.
ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00
<https://doi.org/10.48550/arXiv.2301.09902>

¹<https://www.gigaspace.com/blog/amazon-found-every-100ms-of-latency-cost-them-1-in-sales>

²<https://about.netflix.com/en/news/four-years-after-house-of-cards-netflix-members-elect-their-owntv-schedule>

labelers to rate the portraits based on income and perceived warmth, competence, status, and competition. In this way, we investigate relationships between stereotype perceptions and income within and between demographic groups. We developed this study design to answer the following research questions:

RQ1 What effect do stereotypes have on the way people estimate others in labeling tasks?

RQ2 What is the impact of the ethnicity or sex of labelers on their judgment of others?

Our results show that labelers possess stereotypes independent of their own demographics and that these stereotypes impact the labels they assign. Further, we show that the ethnicities of both labelers and portraits impact the predictions assigned by the labelers³. Our results indicate that labeler bias is a function of both labeler demographics and characteristics of the labeled subject, suggesting that recruiting a diverse labeler pool may not be enough to counteract the bias.

Overall, this paper contributes evidence for the existence of labeler bias and discusses the consequences. In particular, we contribute an investigation of bias in the context of a labeling task involving human faces using a publicly available dataset. Our findings raised awareness for labeler bias; in return, we hope that researchers and practitioners will ensure a diverse labeler sample to counteract the raised concerns. Thus, revisiting current practices is needed to develop AI systems with fewer biases.

2 RELATED WORK

In this section, we first present prior work on data bias in machine learning, examining historical bias and labeler bias. Next, we introduce the stereotype content model (SCM) and how it relates to computer science and our work.

2.1 Bias in Machine Learning

Machine learning models and algorithms can have biases introduced at multiple stages. Algorithms can contain systematic biases embedded by the moral concepts of developers [27, 55]. However, data bias is a much more prevalent concern in intelligent systems. As Mueller [45] describes, the quality of a system is coupled “with the quality of the data provided, following the old slogan *garbage in, garbage out*”. It follows that an AI system will make biased decisions if it is trained on biased data. Notable sources of bias, which we detail below, include historical bias, non-representative sampling, and labeler bias.

2.1.1 Historical and Sampling Bias. Historical bias occurs when a system is trained on data resulting from real, biased scenarios. In an early example, St. George’s University Medical School introduced a new computer system that systematically denied admission to women and people with ‘foreign-sounding names’ based on historical data [36, 52]. Similar issues persist in many modern systems. A CV screening tool at Amazon preferentially hired men even after explicit references to gender, race, and sexual orientation were removed from the dataset [34]. Algorithms for crime prediction

typically rely on historical crime data in which ethnic minorities are over-represented⁴. Even using online proxies to remove sensitive characteristics continues to lead to biased decisions [36, 59]. Such systems can have a drastic impact on the lives of real people. For example, minority groups have an increased likelihood of being stopped and searched by the police, based only on immutable characteristics which they cannot control [2].

Generating datasets based on emerging data can also lead to biased results. Bender et al. [5] describe the experience of collecting data through the internet: “white supremacist and misogynistic, ageist, etc. views are over-represented in the training data, not only exceeding their prevalence in the general population but also setting up models trained on these datasets to amplify biases and harms further.” Women are also underrepresented on platforms used for data collection [4], which results in decisions based on male-skewed data sources. Although issues of historical bias and non-representative sampling continue to be an issue, they are not the focus of this paper. Rather, we aim to investigate the under-explored potential impact of bias within the people labeling data to generate biased datasets.

Sampling bias has also been identified in face annotation tasks. Da Silva and Pedrini [54] found that an emotion classifier trained on one cultural group was inaccurate when used on a different cultural group. Scheuerman et al. [51] recommend embedding race and gender information into databases and classifications to increase transparency.

2.1.2 Labeler Bias. Labeler bias occurs when a group of people annotates a dataset and embeds bias into the resulting data. This often occurs unintentionally, as Wall et al. [56] have argued that unconscious biases can influence judgments and lead to inaccurate conclusions in visual analytic feedback tasks. In one poignant example of labeler bias, prior work has shown that people from Western cultures tend to rate people from other cultures as being less attractive than themselves [30]. Consequently, an attractiveness dataset labeled primarily by people from Western cultures risks having this bias embedded⁵. Any system that uses such a dataset as a basis for decisions would subsequently perpetuate the bias of the labelers. Past work in machine learning has shown that fairness and accuracy can be aligned, motivating a push towards more fairness in training as a simultaneous push towards increased accuracy [58]. Prior research has found that even highly experienced labelers fail to produce unbiased labels [29]. Perhaps most relevant to our study is recent work in CSCW by Goyal et al. [24] on rater identity. They found that rater identity (i.e., African American, LGBTQ, or neither) significantly influenced how raters annotated toxicity in online comments. They suggest that raters who self-identify with the identities targeted in online comments provide additional nuance and more inclusivity in trained models. Similar work found that other social variables (e.g., conservatism) also impact toxicity labels [50]. Past work has also identified that labeler bias also stems from socio-economic contexts and the application of power structures within annotation companies [44]. In response, researchers have proposed more fair and human-centric

³Please note from the outset that we do not analyze which groups are biased in which directions. Our aim is not to apply judgments to specific groups but rather to investigate whether ratings are impacted by demographics in an effort to promote fairness and ethics in machine learning.

⁴<https://blogs.scientificamerican.com/voices/i-know-some-algorithms-are-biased-because-i-created-one/>

⁵Attractiveness Test: <https://attractiveness-test.com/>

crowdsourcing frameworks considering demographics and appropriate compensation [3]. In the presented examples, we can see that labeler bias can stem from a number of characteristics of the labeling group. In this paper, we extend the body of work on labeler bias by exploring whether demographic clusters of labelers can have an impact on labeling task estimates.

Past work has attempted to find solutions to account for labeler bias. Several strategies have been applied in an attempt to model labeler bias, including using knowledge about the ground truth [32], bayesian methods [57], or multi-task Gaussian Processes [6]. Geva et al. [22] recommend that labelers for testing and training datasets be distinct groups since they found that subjective NLP labels produced by a group do not generalize well. Instructions for annotators have also been found to embed bias [42, 47]. In the context of face annotation, Engelmann et al. [15] argue that ‘secondary’ (i.e., subjective) characteristics may not be appropriate attributes for facial recognition systems to predict.

2.2 Stereotype Content Model

The Stereotype Content Model (SCM) is a theory that explains how people develop stereotypes about other people. The SCM proposes that people primarily use two dimensions, warmth and competence, to assess other people. These dimensions prognosticate emotional prejudices, which in turn lead to discrimination [17–19]. People group themselves based on what they perceive to be similarities between themselves and others. Across-group prejudice is a judgment on an emotional, cognitive, or behavioral level towards another group [12].

The warmth-competence model is a well-researched [9, 13, 26, 38, 53] fundamental theory in social psychology with broad implications for both social perception and social interaction [7]. The model helps to explain how we form stereotypes about different groups of people, positing that we judge groups based on how warm and competent we perceive them to be [8, 16]. Warmth is the evaluation criterion by which a person is perceived as friendly and trustworthy, while competence shows how capable and compelling a person is perceived to be. Generally, our reaction to individuals from certain groups is governed by how warm and competent we perceive them to be. For example, people generally see elderly individuals as warm but not competent and therefore react with pity. Groups seen as competent but not warm are met with envy, while those perceived as both competent and warm tend to be admired [19]. The SCM has also been extended to include competition and status, which are particularly influenced by how an individual views a specific group relative to their own group [19].

Nicolas et al. [46] made a significant contribution by developing dictionaries for stereotypical content. These dictionaries simplify the study of stereotypes and speed up the identification of social biases in AI, social media, and other textual sources. Subsequent work has used these dictionaries to automatically identify the stereotypical language in news coverage [35] or to mitigate stereotypical language through anti-stereotypes [21]. The SCM has also been used to understand human-machine interactions. For example, McKee et al. [41] used the SCM to explore how individuals react to different digital avatars. They found that users increasingly perceive the system as being warm if it appears in the role of an assistant

and cold if it appears in the role of a competitor. The SCM has also been used in HCI to investigate stereotypes in personas [40] and determine the social acceptability of mobile devices [53]. In our work, we are investigating the way that labeler stereotypes influence the labels they assign during annotation tasks. As such, we use the SCM, a well-established model, to understand how stereotypes vary across our labeler population.

3 METHOD

We conducted an online survey where participants from various demographic backgrounds completed labeling tasks. We asked participants to label portraits of varying ethnicities and sexes based on income and the SCM. For this, we balanced the participants’ self-reported ethnicity.

Since ethnicity labels are not clearly defined [48], we aligned our ethnicity categorization with the FairFace dataset. The following seven groups listed in FairFace will be referred to as ethnicities⁶ in this work: (1) Black, (2) East Asian, (3) Indian, (4) Latino, (5) Middle Eastern, (6) Southeast Asian, and (7) White.

3.1 Dataset Preprocessing and Portrait Selection

For our study, we selected portraits from the FairFace dataset (License CC BY 4.0) [33]. This dataset consists of images of people and was specifically developed to be balanced in ethnicity, sex, and age. The authors generated this dataset to foster the development of fair and inclusive machine learning models. We selected the FairFace dataset because it provides us with a high probability of finding high-quality images across a wide range of demographic categories. The ethnicity, gender, and age tags for the images were labeled by Amazon Mechanical Turk (MTurk) users based on a two-thirds majority vote [33]. The majority vote labeling process is common practice for labeling [60], but consequently, we do not have a ground truth for the labels. However, to our knowledge, this is the most extensive dataset available with diverse ethnicities and sexes and was, therefore, the most suitable choice for this study.

We defined several criteria to filter the images in the dataset and create a subset for our study. Our research questions concern ethnicity and sex, so we designated age as a control variable and only selected images in the dataset within the range of 27 to 40 years old (the narrowest age filter provided by the dataset). To avoid any confounding factors, we also defined the following criteria: 1) Only one visible, camera-facing person with a neutral facial expression, 2) A neutral background with no cropped edges, 3) In color (i.e., no black & white images), and 4) No glasses/sunglasses/headgear.

As the FairFace dataset contains more than 100,000 images, we could not manually filter all images by our predefined criteria. Therefore, we created a preprocessing script with the following functions: (1) We removed any images with age labels lower than 27 or higher than 40 (2) We detected faces within each image using the python library “face-recognition”⁷ and removed any images where the number of faces was equal to zero or greater than one. (3) We detected face landmarks using the python library *mediapipe* [39]

⁶We use the term “ethnicity,” as it encompasses more social aspects and is a broader term than “race,” although they are often used interchangeably in practice. We also note that ethnic distinctions are non-specific, but such labels are commonly used in Machine Learning applications such as the FairFace dataset.

⁷https://github.com/ageitgey/face_recognition




Figure 1: We implemented a script to exclude non-frontal faces and images containing more than one person.

and used these landmarks (see Figure 1) to remove images where the subject is not facing the camera. A face is considered to be non-frontal if it deviates on the x-axis (see Figure 1b) or the y-axis (see Figure 1c) beyond a threshold of 0.09, or if both thresholds are crossed, indicating that the face is completely non-frontal (see Figure 1d). If none of the thresholds are crossed, the face is considered to be frontally aligned (see Figure 1a). (4) We detected facial expressions using a machine learning model based on⁸. We selected only images with neutral facial expressions. (5) Finally, we were left with 1,834 images which we manually filtered. Three authors evaluated each of the remaining images and only selected those that fit all of the above criteria. The three authors triple-coded and only included images where all agreed. This resulted in 56 portraits, four for each sex and ethnicity combination.

3.2 Participants

As our study investigates the existence and impact of labeler bias, it was crucial for us to have an accurate picture of our participants' demographic information. We recruited 98 participants (49 female and 49 male) from Prolific⁹. Participants were between 18 and 52 years old ($M = 26.1$, $SD = 6.9$). The participants were equally distributed among the seven ethnicity categories from the FairFace dataset, and age. They live in 21 countries¹⁰. Participants were compensated at a rate of 10€ per hour for a total of 3€ (approximately 2.54£). The study was approved by the ethics committee within the University Faculty¹¹. Participants' income, as indicated in £ on Prolific, varied across ethnicities, shown in Figure 2. We analyzed the correlation between participant income and their income estimates for the portraits to prevent estimation bias. A Pearson correlation showed no significant correlation ($r = -.025$, $p = .403$), so we assume that participants' own income does not influence the results.

3.3 Study Procedure

We used Prolific, a crowdsourcing marketplace, to gather data as it can provide demographic information about the participants. We created a separate posting for each sex and ethnicity category

and used built-in demographic filtering features to ensure that we recruited an equal number of participants from each category.

The participants were first given a brief overview before providing informed consent and completing a demographic questionnaire. The demographic information was also provided by Prolific, but we collected it in the survey to ensure that the ethnicity and sex labels matched the terminology used in the dataset. We then asked participants to respond to questions associated with portraits of people. Each participant was presented with one randomly chosen portrait from each ethnicity and sex category, resulting in 14 portraits per participant.

For each portrait, the participants responded to SCM questions about perceptions of warmth, competence, status, and competition based on [19]¹². Several sub-scale questions are averaged to score ratings of warmth, competence, status, and competition on a scale. The SCM is an established method of measuring stereotype attitudes [19]. We also asked participants to estimate the income for each portrait. To prevent country- and currency-specific biases, participants were asked to assign a value between "high" or "low" with a slider rather than a dollar amount. This task mimics several realistic scenarios. For example, it is customary in some European to include a photo on a CV when applying for a job. Hiring personnel, therefore, make judgments associating income and suitability for a job based on a photo of a face. Finally, we included attention checks (correctly answering a multiple choice question about information written in a short text) in the survey to prevent spam responses, which is common practice in crowdsourced tasks [1].

4 RESULTS

To investigate the relationship between demographics and labels, we performed two-way ANOVA models (Type III, $\alpha = .05$) using Greenhouse-Geisser correction [25] where the sphericity assumption is violated. Note that although it would be possible to compare all levels on all factors and their interaction with post hoc tests, we refrain from doing so. First, test-corrections will be very conservative for pairwise comparisons due to the high number of levels. Second, we were only interested in showing that our factors can

⁸https://github.com/priya-dwivedi/face_and_emotion_detection

⁹Prolific: <https://www.prolific.co>

¹⁰Country information is included in the supplementary material.

¹¹Details removed for anonymization purposes

¹²The full list of questions is provided in the supplementary material.

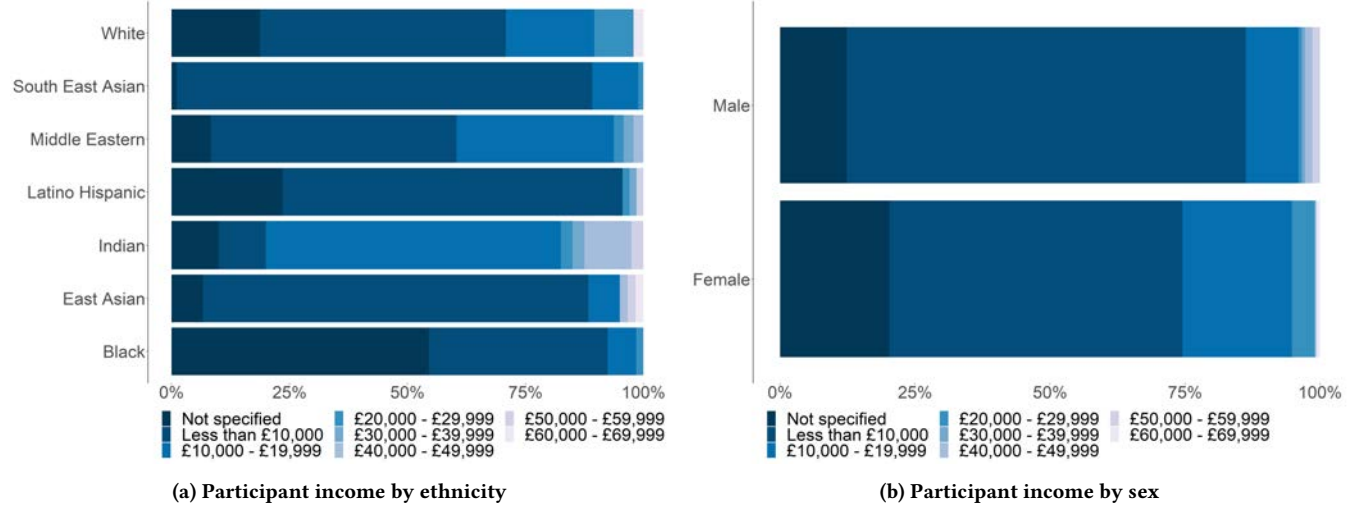


Figure 2: The average participant income by ethnicity and age. A Pearson correlation showed that income variation had no significant impact on the results

explain variation on the dependent variable. As such, we do not analyze which specific biases are present in specific groups.

We present the results in line with our research questions. First, we investigate whether stereotypes impact labeling (RQ1) and then analyze how ethnicity and sex influence labels (RQ2).

4.1 The Impact of Stereotypes on Estimations (RQ1)

In line with the SCM [19], Figure 3 shows the Warmth-Competence ratings assigned to the portraits by the labelers. Figure 3a shows the stereotypes assigned to each $PORTRAIT_{Ethnicity}$ while Figure 3b presents the stereotypes assigned by each $LABELER_{Ethnicity}$. All of the ratings are clustered near neutral warmth and neutral competence.

We conducted a Pearson correlation analysis for warmth, competence, status, and competition against the estimated income. Table 1 shows that three out of four stereotype variables have a significant positive correlation with income. Competence, status, and competition all covary significantly with income, with status being the most positively correlated. Only warmth did not significantly covary with the income estimates.

In general, the biases were in the predicted directions [19]. For example, participants generally estimated a low income for a person they rated as low status, and vice versa. Figure 4 illustrates this phenomenon, showing a positive correlation between Status and Income. Additional plots for all stereotype variables are in Appendix A.

4.2 The Impact of Demographics on Estimations (RQ2)

Our second research question asks whether participant demographics impact their perceptions of stereotypes and their income estimations. Figure 5 shows the income labels assigned to each $PORTRAIT_{Ethnicity}$ demonstrates that $LABELER_{Ethnicity}$ and $PORTRAIT_{Ethnicity}$ have

by each $LABELER_{Ethnicity}$. The corresponding plots for each of the stereotype variables are included in Appendix B.

We conducted an ANOVA using the interaction effect of $LABELER_{Ethnicity}$ and $PORTRAIT_{Ethnicity}$. The results, shown in Table 2, reveal that the $LABELER_{Ethnicity}$ significantly impacts status, the $PORTRAIT_{Ethnicity}$ significantly impacts all variables, and the interaction effect is significant for all variables except warmth. For an exemplary chart that shows income as a function of $LABELER_{Ethnicity}$ and $PORTRAIT_{Ethnicity}$, see Figure 5.

We also performed an ANOVA to investigate the impact of Sex on the estimates. We found no significant main- or interaction-effects on income (all $p > .05$), competence (all $p > .05$), status (all $p > .05$) or competition (all $p > .05$) estimates. Note however, that in the ANOVA on warmth estimates, we found a main effect of $PORTRAIT_{Sex}$, $F(1, 96) = 10.06$, $p = .002$, $\eta_p^2 = .09$. None of the other ANOVA-terms were significant (all $p > .05$).

5 DISCUSSION

In this work, we set out to answer two research questions: (RQ1) What effect do stereotypes have on the way people estimate others in labeling tasks? (RQ2) What is the impact of the ethnicity or sex of labelers on their judgment of others? In the following, we discuss how our results address the research questions, the consequences for machine learning practice, and the limitations of our study.

5.1 Labelers Exhibit Bias

Our results provide evidence that labeler bias exists in two ways. First, our findings indicate that **labelers hold stereotypes about the people depicted in the portraits** they are tasked with annotating (see Figure 3a) and that these stereotypes have an impact on their subsequent income labels (see Table 1). Second, Table 2 demonstrates that $LABELER_{Ethnicity}$ and $PORTRAIT_{Ethnicity}$ have

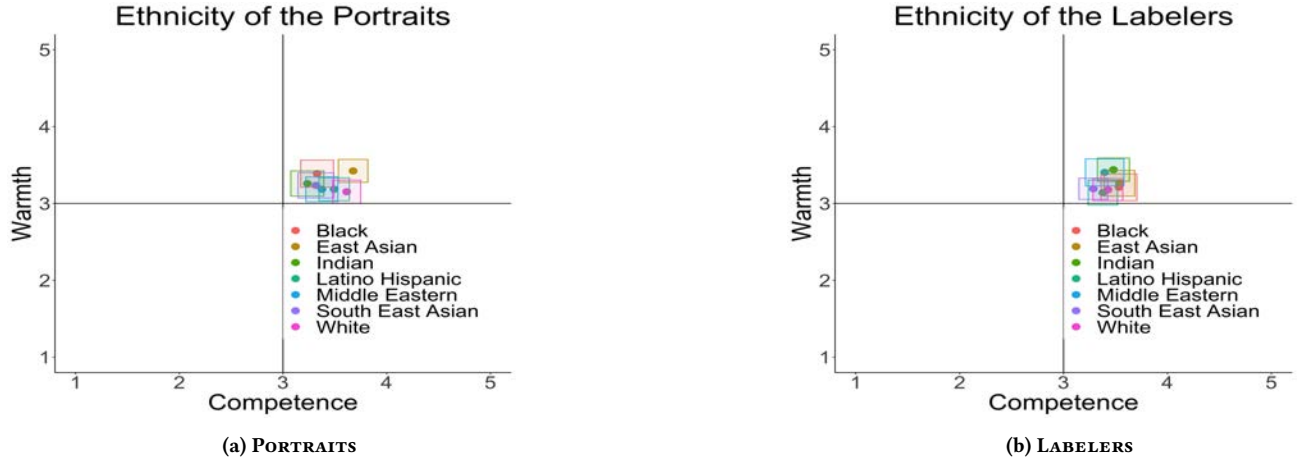


Figure 3: Warmth-Competence ratings displayed from the perspective of the portraits and the labelers, including a 95% confidence interval.

Table 1: The Pearson correlations for each of the stereotype variables and the estimated income. ** indicates $p < .01$; * indicates $p < .05$. We have also computed Linear mixed models that can take into account the nested structure in the data, however, they did not show any noteworthy difference from the simple correlations. For the sake of brevity, we only report the Pearson correlations.

	p	r
Warmth	.093	.045
Competence	<.001**	.541
Status	<.001**	.773
Competition	<.001**	.431

Table 2: The two-way ANOVA results for the Income estimates and the four stereotype variables for $\text{LABELER}_{\text{Ethnicity}}$ and $\text{PORTRAIT}_{\text{Ethnicity}}$.

	$\text{LABELER}_{\text{ethnicity}}$					$\text{PORTRAIT}_{\text{ethnicity}}$					$L \times P$				
	df	dfs	F	p	η_p^2	df	dfs	F	p	η_p^2	df	dfs	F	p	η_p^2
Income	6	91	0.92	.486	.06	5.28	481	49.2	<.001	.35	31.7	481	2.39	<.001	.14
Warmth	6	91	1.2	.313	.07	4.92	448	3.9	.002	.04	36.	546	1.35	.089	.08
Competence	6	91	0.71	.642	.04	6.	546	11.2	<.001	.11	36.	546	1.64	.012	.1
Status	6	91	4.33	<.001	.22	6.	546	40.2	<.001	.31	36.	546	2.15	<.001	.12
Competition	6	91	2.66	.02	.15	4.81	438	31.6	<.001	.26	28.9	438	1.84	.006	.11

a significant impact on perceptions of stereotypes and income labels¹³. Our results show that $\text{LABELER}_{\text{Ethnicity}}$ significantly impacts status, $\text{PORTRAIT}_{\text{Ethnicity}}$ significantly impacts all variables, and the interaction effect is significant for all variables except warmth.

We found that income estimates were correlated with stereotype perceptions (RQ1) and that both stereotype perceptions and income

¹³Note that we did not test all comparisons post hoc as this procedure would not be informative to our study for two main reasons: First, one would need to apply very conservative α -level corrections in order to avoid Type-II error inflation. Second, we were not interested in characterizing the specific bias of a certain group but rather intend to show that variation within the estimates can be explained by considering the interaction of labeler demographics and portrait demographics.

estimates were impacted by $\text{LABELER}_{\text{Ethnicity}}$ and the $\text{PORTRAIT}_{\text{Ethnicity}}$ (RQ2). Interestingly, SEX only had a significant effect on warmth for this task. This is in line with psychological research showing how stereotype judgments are made across cultures [8, 38] and, in particular, how gender stereotypes influence perception [7]. In this domain, where labelers are annotating portraits of humans, we can conclude that labeler bias exists, depends on labeler demographics, and can be explained using stereotype content [20]. These results are in line with recent findings in CSCW by Goyal et al. [24] demonstrating that toxicity labels for online content are influenced by labeler self-identification.

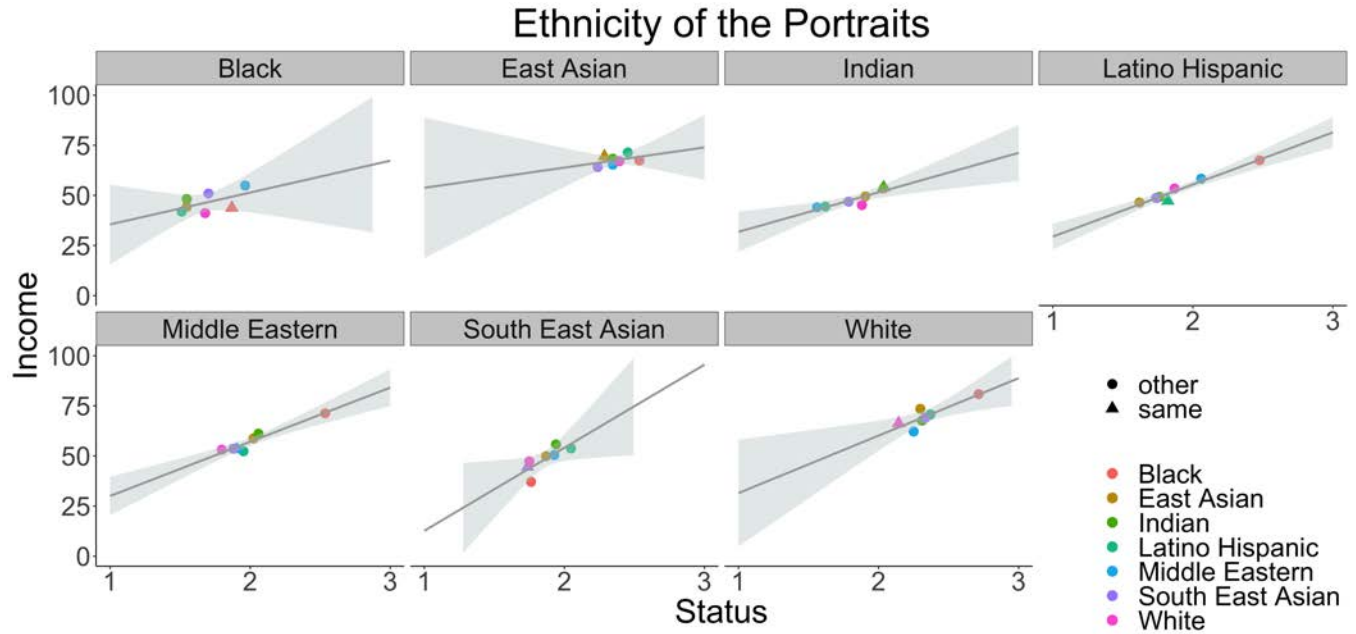


Figure 4: Correlation between mean status and income. Each subplot represents a portrait ethnicity and the points in each plot show how labelers of each ethnicity rated the portraits.

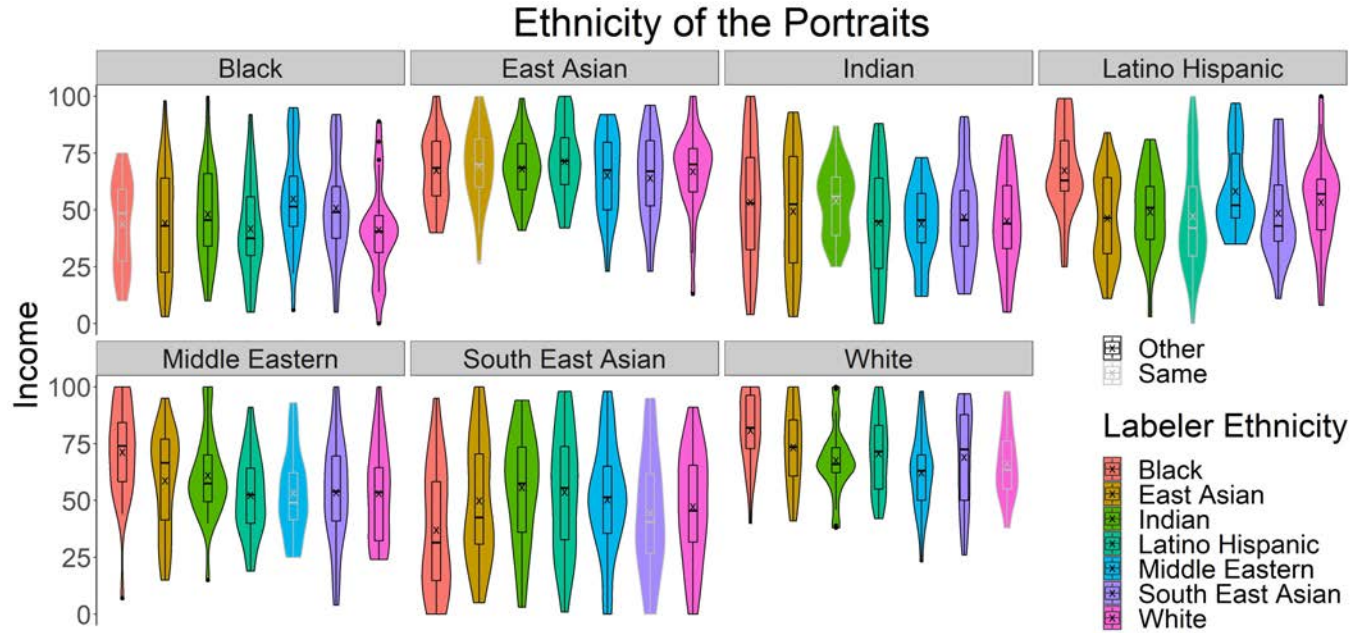


Figure 5: Estimated income as a function of $\text{LABELER}_{\text{Ethnicity}}$ and $\text{PORTRAIT}_{\text{Ethnicity}}$. Grey borders indicate the cases where $\text{LABELER}_{\text{Ethnicity}}$ and $\text{PORTRAIT}_{\text{Ethnicity}}$ match.

5.2 Implications for Machine Learning

Our results shed some light on the impact of the human-aspect of machine learning. We have shown that labels vary with labeler

demographics for annotation tasks involving portraits. This is important because prior work has shown that the majority of workers on MTurk are from the United States and India [10, 31, 49], and Levay et al. [37] found that over 70% of MTurkers identify as white. Consequently, the status quo of gathering labels from crowdsourcing

platforms without considering *who* is doing the labeling should be reconsidered, as it will almost certainly lead to a non-balanced labeler pool and subsequently biased labels. However, as shown in Figure 3a and Table 2, stereotypes occur in labeling tasks involving images of people independent of labeler demographics. Recruiting labelers from a wide variety of populations should result in labels that are biased in a way that is consistent with societal biases, but the bias will still exist. Therefore, it remains an open research question as to whether it is possible to combat labeler bias through recruiting or to correct it post hoc. Post hoc methods have been proposed for bias in machine learning labels (e.g., [32]), but typically require knowledge of a ground truth. Ground truths for social issues are complicated by the fact that representation in data is connected to the unequal distribution of power in society [11]. On such social issues, where ground truths may be fuzzy or non-existent, and society may be biased against particular groups, there is space for crucial future research to develop solutions. Past work by Miceli et al. [43] suggests expanding data documentation and incorporating social contexts, which could be an important step toward ensuring fair, transparent data pipelines.

5.3 Limitations & Future Work

The most notable limitation of our study is the fact that we have only explored one specific labeling task, namely annotating secondary characteristics of faces. This labeling scenario was chosen because it was likely to highlight the impact of stereotypes while still being rooted in a realistic scenario, such as making hiring decisions, which have been shown to be impacted by stereotypes [23]. However, further work is required to understand how stereotypes and labeler demographics impact more abstract labeling tasks, such as image segmentation or product labeling. There are many high-stakes scenarios relevant to specific tasks across the field of machine learning that may be impacted by this phenomenon. Since we have now shown that there is a bias in this scenario, we call on future work to explore other scenarios which are relevant to specific tasks in machine learning. Quantifying how these results generalize to other tasks is crucial to understanding when labeler demographics and their interactions with content must be accounted for.

Another limitation derives from the methodology used to create the FairFace dataset. The ethnicity, sex, and age labels in the dataset were created using a majority vote method on Amazon Mechanical Turk. The labels, therefore, are not necessarily a true ground truth. In a similar vein, there is no ground truth for the income estimates since this information was not included in the FairFace dataset. Despite these limitations, this is the best available dataset we could locate with a balanced representation and labeled demographics and was, therefore, the best option available for this work. As we only show that there is a potential for biases, future studies should look at the magnitude of bias as a difference from the ground truth. Resolving this issue is not trivial, as it likely requires that a new database of images be generated with associated income levels provided by the image subjects, but it would be a worthy endeavor to further investigate and characterize this information.

6 CONCLUSION

In this paper, we investigated the existence and impact of labeler bias. We recruited 98 participants to engage in an online survey where we asked them to estimate the income and rate the perceived warmth, competence, status, and competition of people from multiple ethnicities and sexes portrayed in a series of images. We selected the portraits from the FairFace dataset using predefined exclusion criteria to create a balanced set of portraits. We found a significant relationship between income estimates and perceptions of competence, status, and competition. Additionally, the results indicate that the main- and interaction-effects of $\text{LABELER}_{\text{Ethnicity}}$ and $\text{PORTRAIT}_{\text{Ethnicity}}$ significantly affect estimations, while $\text{LABELER}_{\text{Sex}}$ and $\text{PORTRAIT}_{\text{Sex}}$ did not significantly impact the results. This insight poses a major challenge for AI applications, as it implies that datasets annotated by a non-diverse set of labelers are likely to carry stereotypes. Thus, we recommend that further research develops guidelines for responsible dataset generation.

REFERENCES

- [1] James D. Abbey and Margaret G. Meloy. 2017. Attention by design: Using attention checks to detect inattentive respondents and improve data quality. *Journal of Operations Management* 53-56 (2017). <https://doi.org/10.1016/j.jom.2017.06.001>
- [2] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2016. Machine bias. (2016), 254–264. <https://doi.org/10.1201/9781003278290-37>
- [3] Natã M. Barbosa and Monchu Chen. 2019. Rehumanized Crowdsourcing: A Labeling Framework Addressing Bias and Ethics in Machine Learning. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3290605.3300773>
- [4] Michael Barera. 2020. Mind the Gap: Addressing Structural Equity and Inclusion on Wikipedia. <http://hdl.handle.net/10106/29572>
- [5] Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (Virtual Event, Canada) (FAccT '21). ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3442188.3445922>
- [6] Trevor Cohn and Lucia Specia. 2013. Modelling Annotator Bias with Multi-task Gaussian Processes: An Application to Machine Translation Quality Estimation. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Sofia, Bulgaria, 32–42. <https://aclanthology.org/P13-1004>
- [7] Amy J. C. Cuddy, Susan T. Fiske, and Peter Glick. 2008. Warmth and Competence as Universal Dimensions of Social Perception: The Stereotype Content Model and the BIAS Map. *Advances in experimental social psychology* 40 (2008), 61–149. [https://doi.org/10.1016/S0065-2601\(07\)00002-0](https://doi.org/10.1016/S0065-2601(07)00002-0)
- [8] Amy J. C. Cuddy, Susan T. Fiske, Virginia S. Y. Kwan, Peter Glick, Stephanie Demoulin, Jacques-Philippe Leyens, Michael Harris Bond, Jean-Claude Croizet, Naomi Ellemers, Ed Sleebos, Tin Tin Htun, Hyun-Jeong Kim, Greg Maio, Judi Perry, Kristina Petkova, Valery Todorov, Rosa Rodríguez-Bailón, Elena Morales, Miguel Moya, Marisol Palacios, Vanessa Smith, Rolando Perez, Jorge Vala, and Rene Ziegler. 2009. Stereotype content model across cultures: Towards universal similarities and some differences. *British Journal of Social Psychology* 48, 1 (2009). <https://doi.org/10.1348/014466608X314935>
- [9] Adamantios Diamantopoulos, Arnd Florack, Georgios Halkias, and Johanna Palcu. 2017. Explicit versus implicit country stereotypes as predictors of product preferences: Insights from the stereotype content model. *J Int Bus Stud* 48, 8 (2017). <https://doi.org/10.1057/s41267-017-0085-9>
- [10] Djellel Difallah, Elena Filatova, and Panos Ipeirotis. 2018. Demographics and Dynamics of Mechanical Turk Workers. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining (WSDM '18)*. ACM, New York, NY, USA. <https://doi.org/10.1145/3159652.3159661>
- [11] Catherine D'Ignazio and Lauren F. Klein. 2020. *Data Feminism*. MIT Press.
- [12] John F. Dovidio and Samuel L. Gaertner. 2010. Intergroup bias. In *Handbook of social psychology, Vol. 2, 5th ed.* John Wiley & Sons, Inc., Hoboken, NJ, US. <https://doi.org/10.1002/9780470561119.socpsy002029>
- [13] Federica Durante, Courtney Bearns Tablante, and Susan T. Fiske. 2017. Poor but Warm, Rich but Cold (and Competent): Social Classes in the Stereotype Content Model. *Journal of Social Issues* 73, 1 (2017). <https://doi.org/10.1111/josi.12208>

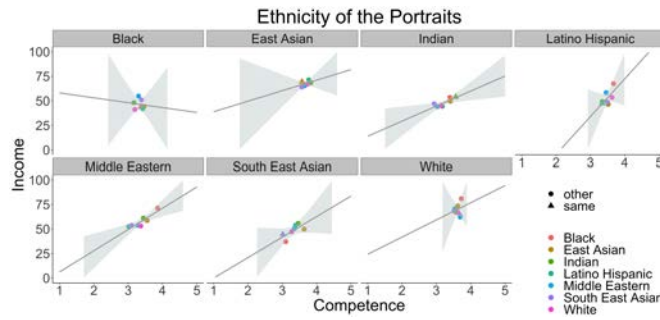
- [14] Carsten Eickhoff. 2018. Cognitive Biases in Crowdsourcing. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining* (Marina Del Rey, CA, USA) (WSDM '18). Association for Computing Machinery, New York, NY, USA, 162–170. <https://doi.org/10.1145/3159652.3159654>
- [15] Severin Engelmann, Chiara Ullstein, Orestis Papakyriakopoulos, and Jens Grossklags. 2022. What People Think AI Should Infer From Faces. In *2022 ACM Conference on Fairness, Accountability, and Transparency (FACCT '22)*. Association for Computing Machinery, New York, NY, USA, 128–141. <https://doi.org/10.1145/3531146.3533080>
- [16] Susan T. Fiske. 2015. Intergroup Biases: A Focus on Stereotype Content. *Curr Opin Behav Sci* 3 (2015). <https://doi.org/10.1016/j.cobeha.2015.01.010>
- [17] Susan T. Fiske. 2018. Stereotype Content: Warmth and Competence Endure. *Current Directions in Psychological Science* 27, 2 (2018). <https://doi.org/10.1177/0963721417738825>
- [18] Susan T. Fiske, Amy J. C. Cuddy, and Peter Glick. 2007. Universal dimensions of social cognition: Warmth and competence. *Trends in cognitive sciences* 11, 2 (2007). <https://doi.org/10.1016/j.tics.2006.11.005>
- [19] 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, 6 (2002). <https://doi.org/10.1037/0022-3514.82.6.878>
- [20] Susan T. Fiske and Shelley E. Taylor. 1991. *Social cognition, 2nd ed.* McGraw-Hill Book Company, New York, NY, England.
- [21] Kathleen C. Fraser, Isar Nejadgholi, and Svetlana Kiritchenko. 2021. Understanding and Countering Stereotypes: A Computational Approach to the Stereotype Content Model. *arXiv:2106.02596 [cs]* (2021). <http://arxiv.org/abs/2106.02596>
- [22] Mor Geva, Yoav Goldberg, and Jonathan Berant. 2019. Are We Modeling the Task or the Annotator? An Investigation of Annotator Bias in Natural Language Understanding Datasets. <https://doi.org/10.48550/arXiv.1908.07898>
- [23] M. José González, Clara Cortina, and Jorge Rodríguez. 2019. The Role of Gender Stereotypes in Hiring: A Field Experiment. *European Sociological Review* 35, 2 (2019). <https://doi.org/10.1093/esr/jcy055>
- [24] Nitesh Goyal, Ian D. Kivlichan, Rachel Rosen, and Lucy Vasserman. 2022. Is Your Toxicity My Toxicity? Exploring the Impact of Rater Identity on Toxicity Annotation. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 363 (nov 2022), 28 pages. <https://doi.org/10.1145/3555088>
- [25] Samuel W. Greenhouse and Seymour Geisser. 1959. On methods in the analysis of profile data. *Psychometrika* 24, 2 (1959).
- [26] Dmitry Grigoryev, Susan T. Fiske, and Anastasia Batkhina. 2019. Mapping Ethnic Stereotypes and Their Antecedents in Russia: The Stereotype Content Model. *Frontiers in Psychology* 10 (2019). <https://doi.org/10.3389/fpsyg.2019.01643>
- [27] Thilo Hagendorff. 2020. The Ethics of AI Ethics: An Evaluation of Guidelines. *Minds & Machines* 30, 1 (2020). <https://doi.org/10.1007/s11023-020-09517-8>
- [28] Keach Hagey and Jeff Horwitz. 2021. Facebook Tried to Make Its Platform a Healthier Place. It Got Angrier Instead. *Wall Street Journal* (2021). <https://www.wsj.com/articles/facebook-algorithm-change-zuckerberg-11631654215>
- [29] Christoph Hube, Besnik Fetahu, and Ujwal Gadiraju. 2019. Understanding and Mitigating Worker Biases in the Crowdsourced Collection of Subjective Judgments. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3290605.3300637>
- [30] Jevan Hutson, Jessie G. Taft, Solon Barocas, and Karen Levy. 2018. Debiasing Desire: Addressing Bias & Discrimination on Intimate Platforms. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW (2018). <https://doi.org/10.1145/3274342>
- [31] Panagiotis G. Ipeirotis. 2010. *Demographics of Mechanical Turk*. Technical Report 1585030. Rochester, NY. <https://papers.ssrn.com/abstract=1585030>
- [32] Heinrich Jiang and Ofir Nachum. 2020. Identifying and Correcting Label Bias in Machine Learning. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*. PMLR. <https://proceedings.mlr.press/v108/jiang20a.html>
- [33] Kimmo Karkkainen and Jungseock Joo. 2021. FairFace: Face Attribute Dataset for Balanced Race, Gender, and Age for Bias Measurement and Mitigation. In *Winter Conference on Applications of Computer Vision (WACV)*. IEEE, Waikoloa, HI, USA. <https://doi.org/10.1109/WACV48630.2021.00159>
- [34] Akhil Alfons Kodyan. 2019. An overview of ethical issues in using AI systems in hiring with a case study of Amazon's AI based hiring tool. *Researchgate Preprint* (2019).
- [35] Anne C. Kroon, Damian Trilling, and Tamara Raats. 2021. Guilty by Association: Using Word Embeddings to Measure Ethnic Stereotypes in News Coverage. *Journalism & Mass Communication Quarterly* 98, 2 (2021). <https://doi.org/10.1177/1077699020932304>
- [36] Nicol Turner Lee, Paul Resnick, and Genie Barton. 2019. Algorithmic bias detection and mitigation: Best practices and policies to reduce consumer harms. *Brookings Institute: Washington, DC, USA* (2019).
- [37] Kevin E. Levay, Jeremy Freese, and James N. Druckman. 2016. The Demographic and Political Composition of Mechanical Turk Samples. *SAGE Open* 6, 1 (2016). <https://doi.org/10.1177/2158244016636433>
- [38] Monica H. Lin, Virginia S. Y. Kwan, Anna Cheung, and Susan T. Fiske. 2005. Stereotype Content Model Explains Prejudice for an Envid Outgroup: Scale of Anti-Asian American Stereotypes. *Pers Soc Psychol Bull* 31, 1 (2005). <https://doi.org/10.1177/0146167204271320>
- [39] Camillo Lugaes, Jiuqiang Tang, Hadon Nash, Chris McClanahan, Esha Ubweja, Michael Hays, Fan Zhang, Chuo-Ling Chang, Ming Guang Yong, Juhyun Lee, Wan-Teh Chang, Wei Hua, Manfred Georg, and Matthias Grundmann. 2019. MediaPipe: A Framework for Building Perception Pipelines. (2019). <http://arxiv.org/abs/1906.08172>
- [40] Nicola Marsden and Maren Haag. 2016. Stereotypes and Politics: Reflections on Personas. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. Association for Computing Machinery, New York, NY, USA, 4017–4031. <https://doi.org/10.1145/2858036.2858151>
- [41] Kevin McKee, Xuechunzi Bai, and Susan Fiske. 2021. *Understanding Human Impressions of Artificial Intelligence*. Technical Report. PsyArXiv. <https://doi.org/10.31234/osf.io/5ursp>
- [42] Milagros Miceli and Julian Posada. 2022. The Data-Production Dispositif. <https://doi.org/10.48550/arXiv.2205.11963>
- [43] Milagros Miceli, Julian Posada, and Tianling Yang. 2022. Studying Up Machine Learning Data: Why Talk About Bias When We Mean Power? *Proceedings of the ACM on Human-Computer Interaction* 6, GROUP (Jan. 2022), 34:1–34:14. <https://doi.org/10.1145/3492853>
- [44] Milagros Miceli, Martin Schuessler, and Tianling Yang. 2020. Between Subjectivity and Imposition: Power Dynamics in Data Annotation for Computer Vision. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW2, Article 115 (oct 2020), 25 pages. <https://doi.org/10.1145/3415186>
- [45] Vincent C. Mueller. 2021. Ethics of Artificial Intelligence and Robotics. (2021). <https://plato.stanford.edu/archives/sum2021/entries/ethics-ai/>
- [46] Gandalf Nicolas, Xuechunzi Bai, and Susan T. Fiske. 2021. Comprehensive stereotype content dictionaries using a semi-automated method. *European Journal of Social Psychology* 51, 1 (2021). <https://doi.org/10.1002/ejsp.2724>
- [47] Mihir Parmar, Swaroop Mishra, Mor Geva, and Chitta Baral. 2022. Don't Blame the Annotator: Bias Already Starts in the Annotation Instructions. <https://doi.org/10.48550/arXiv.2205.00415>
- [48] Jean S. Phinney. 1996. When we talk about American ethnic groups, what do we mean? *American Psychologist* 51, 9 (1996). <https://doi.org/10.1037/0003-066X.51.9.918>
- [49] Joel Ross, Lilly Irani, M. Six Silberman, Andrew Zaldivar, and Bill Tomlinson. 2010. Who are the crowdworkers? shifting demographics in mechanical turk. In *CHI '10 Extended Abstracts on Human Factors in Computing Systems (CHI EA '10)*. ACM, New York, NY, USA. <https://doi.org/10.1145/1753846.1753873>
- [50] Maarten Sap, Swabha Swayamdipta, Laura Vianna, Xuhui Zhou, Yejin Choi, and Noah A. Smith. 2022. Annotators with Attitudes: How Annotator Beliefs And Identities Bias Toxic Language Detection. <https://doi.org/10.48550/arXiv.2111.07997>
- [51] Morgan Klaus Scheuerman, Kandrea Wade, Caitlin Lustig, and Jed R. Brubaker. 2020. How We've Taught Algorithms to See Identity: Constructing Race and Gender in Image Databases for Facial Analysis. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW1 (May 2020), 58:1–58:35. <https://doi.org/10.1145/3392866>
- [52] Oscar Schwartz. 2019. Untold history of AI: Algorithmic bias was born in the 1980s. *IEEE Spectrum* (2019).
- [53] Valentin Schwind, Niklas Deierlein, Romina Poguntke, and Niels Henze. 2019. Understanding the Social Acceptability of Mobile Devices using the Stereotype Content Model. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA. <https://doi.org/10.1145/3290605.3300591>
- [54] Flávio Altinier Maximiano da Silva and Helio Pedrini. 2015. Effects of cultural characteristics on building an emotion classifier through facial expression analysis. *Journal of Electronic Imaging* 24, 2 (March 2015), 023015. <https://doi.org/10.1117/1.JEL.24.2.023015>
- [55] Andreas Tsamados, Nikita Aggarwal, Josh Cows, Jessica Morley, Huw Roberts, Mariarosaria Taddeo, and Luciano Floridi. 2021. The ethics of algorithms: key problems and solutions. *AI & Soc* (2021). <https://doi.org/10.1007/s00146-021-01154-8>
- [56] Emily Wall, Leslie M. Blaha, Celeste Lyn Paul, Kristin Cook, and Alex Endert. 2018. Four Perspectives on Human Bias in Visual Analytics. In *Cognitive Biases in Visualizations*. Springer International Publishing, Cham. https://doi.org/10.1007/978-3-319-95831-6_3
- [57] Fabian L. Wauthier and Michael Jordan. 2011. Bayesian Bias Mitigation for Crowdsourcing. In *Advances in Neural Information Processing Systems*, Vol. 24. Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2011/hash/0768281a05da9f27df178b5c39a51263-Abstract.html>
- [58] Michael Wick, Swetasudha Panda, and Jean-Baptiste Tristan. 2019. Unlocking Fairness: a Trade-off Revisited. *Advances in neural information processing systems* 32 (2019). <https://proceedings.neurips.cc/paper/2019/file/373e4c5d8edfa8b74fd4b6791d0c6dc-Paper.pdf>
- [59] Tal Zarsky. 2015. *Understanding Discrimination in the Scored Society*. Technical Report ID 2550248. Social Science Research Network, Rochester, NY. <https://doi.org/10.2139/ssrn.2550248>

[//papers.ssrn.com/abstract=2550248](https://papers.ssrn.com/abstract=2550248)

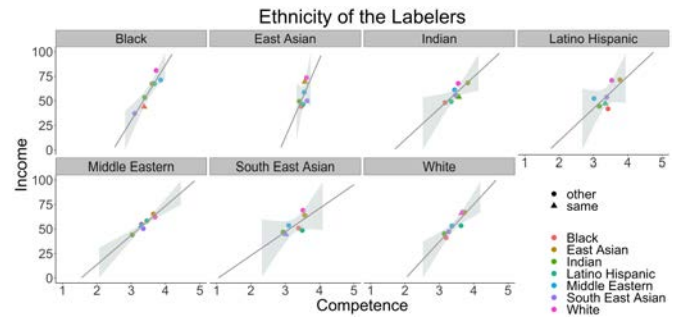
- [60] Jing Zhang, Xindong Wu, and Victor Sheng. 2016. Learning from crowdsourced labeled data: a survey. *Artificial Intelligence Review* 46 (2016). [https://doi.org/10.](https://doi.org/10.1007/s10462-016-9491-9)

1007/s10462-016-9491-9

A INCOME CORRELATION

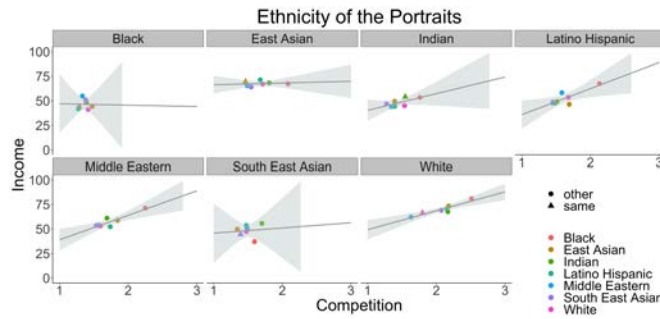


(a) Correlation between mean status and income. Each subplot represents a portrait ethnicity and the points in each plot show how labelers of each ethnicity rated the portraits.

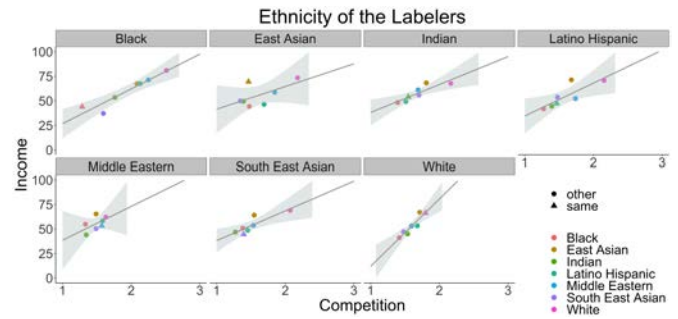


(b) Correlation between mean status and income. Each subplot represents a portrait ethnicity and the points in each plot show how labelers of each ethnicity rated the portraits.

Figure 6: Correlations between income and competence

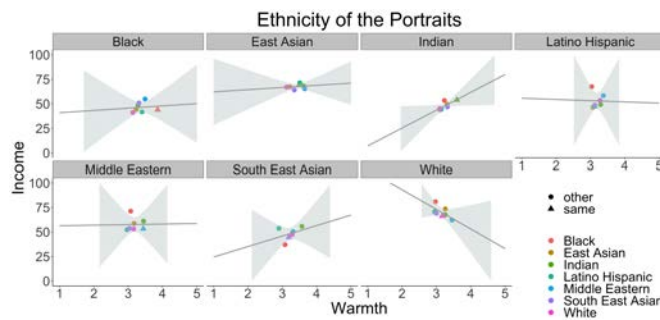


(a) Correlation between mean status and income. Each subplot represents a portrait ethnicity and the points in each plot show how labelers of each ethnicity rated the portraits.

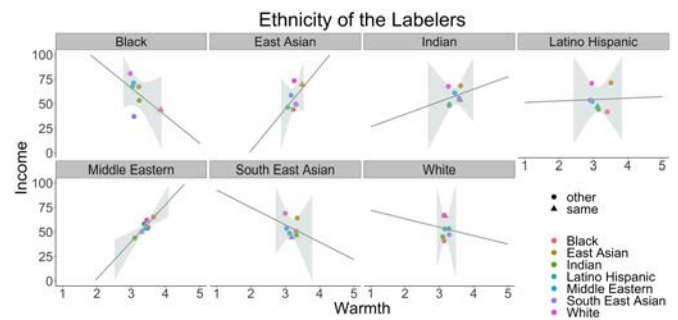


(b) Correlation between mean status and income. Each subplot represents a portrait ethnicity and the points in each plot show how labelers of each ethnicity rated the portraits.

Figure 7: Correlations between income and competition

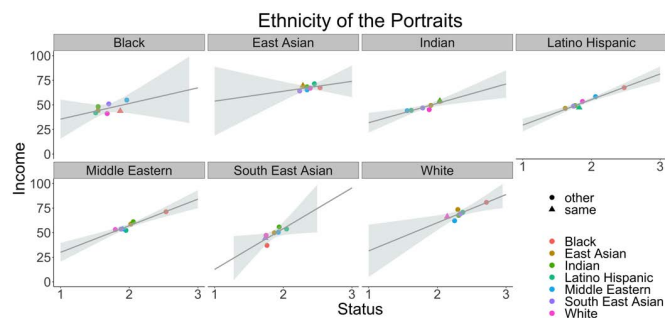


(a) Correlation between mean status and income. Each subplot represents a portrait ethnicity and the points in each plot show how labelers of each ethnicity rated the portraits.

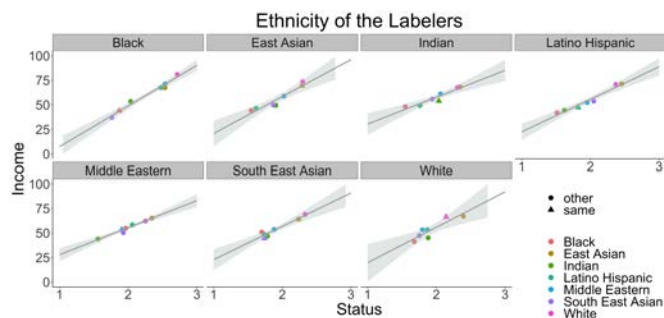


(b) Correlation between mean status and income. Each subplot represents a portrait ethnicity and the points in each plot show how labelers of each ethnicity rated the portraits.

Figure 8: Correlations between income and warmth




(a) Correlation between mean status and income. Each subplot represents a portrait ethnicity and the points in each plot show how labelers of each ethnicity rated the portraits.



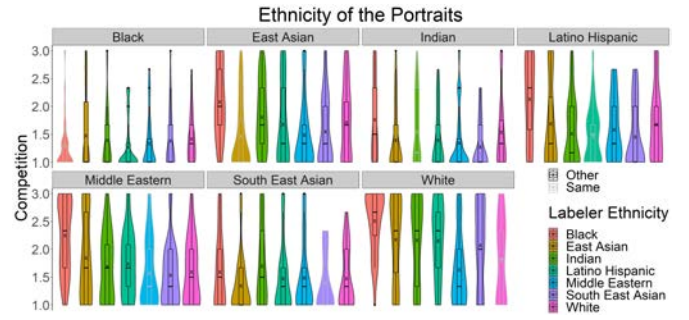
(b) Correlation between mean status and income. Each subplot represents a portrait ethnicity and the points in each plot show how labelers of each ethnicity rated the portraits.

Figure 9: Correlations between income and status

B VARIATION WITH ETHNICITY




(a) Estimated competence as a function of $\text{LABELER}_{\text{Ethnicity}}$ and $\text{PORTRAIT}_{\text{Ethnicity}}$. Grey borders indicate the cases where $\text{LABELER}_{\text{Ethnicity}}$ and $\text{PORTRAIT}_{\text{Ethnicity}}$ match.

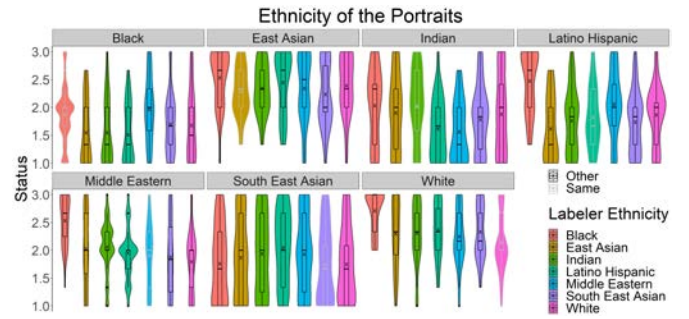


(b) Estimated Competition as a function of $\text{LABELER}_{\text{Ethnicity}}$ and $\text{PORTRAIT}_{\text{Ethnicity}}$. Grey borders indicate the cases where $\text{LABELER}_{\text{Ethnicity}}$ and $\text{PORTRAIT}_{\text{Ethnicity}}$ match.

Figure 10: Competence and competition ratings for each $\text{LABELER}_{\text{Ethnicity}}$ and $\text{PORTRAIT}_{\text{Ethnicity}}$ intersection.



(a) Estimated warmth as a function of $\text{LABELER}_{\text{Ethnicity}}$ and $\text{PORTRAIT}_{\text{Ethnicity}}$. Grey borders indicate the cases where $\text{LABELER}_{\text{Ethnicity}}$ and $\text{PORTRAIT}_{\text{Ethnicity}}$ match.



(b) Estimated status as a function of $\text{LABELER}_{\text{Ethnicity}}$ and $\text{PORTRAIT}_{\text{Ethnicity}}$. Grey borders indicate the cases where $\text{LABELER}_{\text{Ethnicity}}$ and $\text{PORTRAIT}_{\text{Ethnicity}}$ match.

Figure 11: Status and warmth ratings for each $\text{LABELER}_{\text{Ethnicity}}$ and $\text{PORTRAIT}_{\text{Ethnicity}}$ intersection.