

Investigating Labeler Bias in Face Annotation for Machine Learning

Luke HALIBURTON^{a,b,1}, Sinksar GHEBREMEDHIN^a, Robin WELSCH^c,
Albrecht SCHMIDT^a, and Sven MAYER^a

^aLMU Munich, Germany

^bMunich Center for Machine Learning (MCML), Germany

^cAalto University, Finland

ORCID ID: Luke Haliburton <https://orcid.org/0000-0002-5654-2453>, Sinksar Ghebremedhin <https://orcid.org/0000-0002-2874-2909>, Robin Welsch <https://orcid.org/0000-0002-7255-7890>, Albrecht Schmidt <https://orcid.org/0000-0003-3890-1990>, Sven Mayer <https://orcid.org/0000-0001-5462-8782>

Abstract. In a world increasingly reliant on artificial intelligence, it is more important than ever to consider the ethical implications of artificial intelligence. One key under-explored challenge is labeler bias — bias introduced by individuals who label datasets — which can create inherently biased datasets for training and subsequently lead to inaccurate or unfair decisions in healthcare, employment, education, and law enforcement. Hence, we conducted a study ($N=98$) 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.

Keywords. bias, machine learning, crowdworkers, annotation, labeler bias

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 behavior², Netflix studies entertainment preferences³, and Facebook uses social interactions to tailor content to their users [1]. 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 due to the potential for bias to be em-

¹Corresponding Author: Luke Haliburton, luke.haliburton@ifi.lmu.de.

²<https://www.gigaspace.com/blog/amazon-found-every-100ms-of-latency-cost-them-1-in-sales>

³<https://about.netflix.com/en/news/four-years-after-house-of-cards-netflix-members-elect-their-owntv-schedule>

bedded in datasets, creating unintended consequences. One under-investigated contributing factor to biased AI tools is labeler bias, which results from cognitive biases [2] in crowd workers and other dynamics in the labeling process [3]. 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 [4]. 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 [5, 6]. One critical paper by Bender et al. [7] 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 [8, 9, 10]. 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 [11]. The Stereotype Content Model (SCM) is an established practical theory explaining stereotypes, such as perceived warmth and competence, that has been applied in Human-Computer Interaction (HCI) (e.g., [12, 13]). 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 crowd workers to label a series of human faces from the FairFace dataset [14]. We selected faces with equal representation from seven ethnicities and two sexes and recruited labelers with the same balanced demographic distribution. We asked 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 is the impact of *stereotypes held by labelers* on their estimates of others in labeling tasks?

RQ2 What is the impact of the *ethnicity or sex of labelers* on their estimates of others in labeling tasks?

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 labelers⁴. 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 its consequences. In particular, we contribute an investigation of bias in the context of a face-labeling task using a publicly available

⁴Please 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.

dataset. Our findings raise awareness for labeler bias, which we hope leads to researchers and practitioners critically examining and revisiting current practices in data labeling.

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

Biases can be introduced in machine learning models and algorithms at multiple stages. Algorithms can contain systematic biases embedded by the moral concepts of developers [15, 16]. However, data bias is a more prevalent concern in intelligent systems. As Mueller [17] 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 [18, 19]. 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 [20]. Algorithms for crime prediction typically rely on historical crime data in which ethnic minorities are over-represented⁵. Even using online proxies to remove sensitive characteristics continues to lead to biased decisions [19, 21]. 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 [22].

Generating datasets based on emerging data can also lead to biased results. Bender et al. [7] 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 [23], 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 [24] found that an emotion classifier trained on one cultural group was inaccurate when

⁵<https://blogs.scientificamerican.com/voices/i-know-some-algorithms-are-biased-because-i-created-one/>

used on a different cultural group. Scheuerman et al. [25] recommend embedding race and gender information into databases and classifications to increase transparency.

2.1.2. *Labeler Bias*

Labeler bias occurs when individuals annotate a dataset and embed bias into the resulting data. This often occurs unintentionally, as Wall et al. [26] 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 [27]. Consequently, an attractiveness dataset labeled primarily by people from Western cultures risks having this bias embedded⁶. 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 [28]. 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. [5] 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 [6]. Past work has identified that labeler bias also stems from socio-economic contexts and the application of power structures within annotation companies [3]. In response, researchers have proposed more fair and human-centric crowdsourcing frameworks considering demographics and appropriate compensation [30]. 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 to model labeler bias include using knowledge about the ground truth [8], bayesian methods [9], or multi-task Gaussian Processes [10]. Geva et al. [31] 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 [32, 33]. In the context of face annotation, Engelmann et al. [34] 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 [11, 35, 36]. People group themselves based on what they perceive to be similarities between themselves and oth-

⁶Attractiveness Test: <https://attractivenessstest.com/>

ers. Cross-group prejudice is a judgment on an emotional, cognitive, or behavioral level towards another group [37].

The warmth-competence model is a well-researched [38, 39, 40, 12, 41] fundamental theory in social psychology with broad implications for both social perception and social interaction [42]. 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 [43, 44]. 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 [36]. 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 [36].

Nicolas et al. [45] 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 [46] or to mitigate stereotypical language through anti-stereotypes [47]. The SCM has also been used to understand human-machine interactions. For example, McKee et al. [48] 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 [13] and determine the social acceptability of mobile devices [12]. 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 labeled 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 [49], we aligned our ethnicity categorization with the FairFace dataset. The following seven groups listed in FairFace will be referred to as ethnicities⁷ 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) [14]. This dataset consists of images of people and was specifically developed to be balanced

⁷We 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.



(a) Face is frontal, included. (b) Eyes are not centered, excluded. (c) Face is not centered on y-Axis, excluded. (d) Distance of edges is not centered, excluded.

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

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 [14]. The majority vote labeling process is common practice for labeling [50], 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”⁸ 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* [51] 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 on⁹. 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.

⁸https://github.com/ageitgey/face_recognition

⁹https://github.com/priya-dwivedi/face_and_emotion_detection

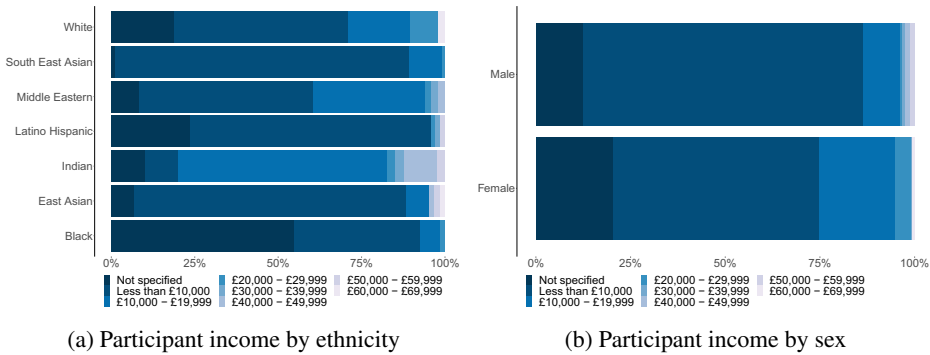


Figure 2. The average participant income by ethnicity and age. A Pearson correlation showed that income variation had no significant impact on the results. Amounts are shown in GBP (£) as this is the currency used by Prolific.

3.2. Participants

We recruited 98 participants (49 female and 49 male) from Prolific¹⁰. 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 countries¹¹. Participants were compensated at a rate of 10€ per hour for a total of 3€. The study was approved by the ethics committee within the LMU Munich University Faculty. 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 [36]¹¹. Several sub-scale questions are averaged to score ratings of warmth, competence, status, and competition on

¹⁰Prolific: <https://www.prolific.co>

¹¹Full study information is provided at <https://github.com/mimuc/labeler-bias>.

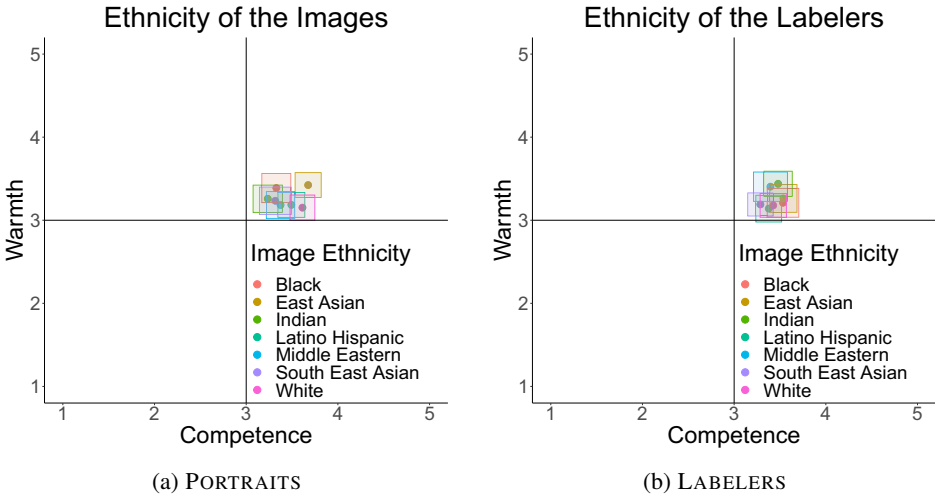


Figure 3. Warmth-Competence ratings displayed from the perspective of the portraits and the labelers, including a 95% confidence interval. All ratings are clustered near neutral (3) for both warmth and competence.

a scale. The SCM is an established method of measuring stereotype attitudes [36]. 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 countries 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 [52].

4. Results

To investigate the relationship between demographics and labels, we performed two-way ANOVA models (Type III, $\alpha = .05$) using Greenhouse-Geisser correction [53] 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 explain variation on the dependent variable. As such, we do not analyze which specific biases are present in specific groups.

4.1. The Impact of Stereotypes on Estimations (RQ1)

In line with the SCM [36], Figure 3 shows the Warmth-Competence ratings assigned to the portraits by the labelers. Figure 3a shows the stereotypes assigned to each $\text{PORTRAIT}_{\text{Ethnicity}}$ while Figure 3b presents the stereotypes assigned by each $\text{LABELER}_{\text{Ethnicity}}$. All ratings are clustered near neutral warmth and neutral competence.

Table 1. The Pearson correlations for each of the stereotype variables and the estimated income. We also computed Linear mixed models that take into account the nested structure in the data. However, showed no noteworthy difference from the simple correlations, so we only report the Pearson correlations for brevity.

	<i>p</i>	<i>r</i>
Warmth	.093	.045
Competence	<.001	.541
Status	<.001	.773
Competition	<.001	.431



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.

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 [36]. 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 Income¹²

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 $PORTRAIT_{Ethnicity}$ by each $LABELER_{Ethnicity}$ ¹².

We conducted an ANOVA using the interaction effect of $LABELER_{Ethnicity}$ and $PORTRAIT_{Ethnicity}$. The results, shown in Table 2, reveal that the $LABELER_{Ethnicity}$ sig-

¹²Additional plots for all variables are available at <https://github.com/mimuc/labeler-bias>

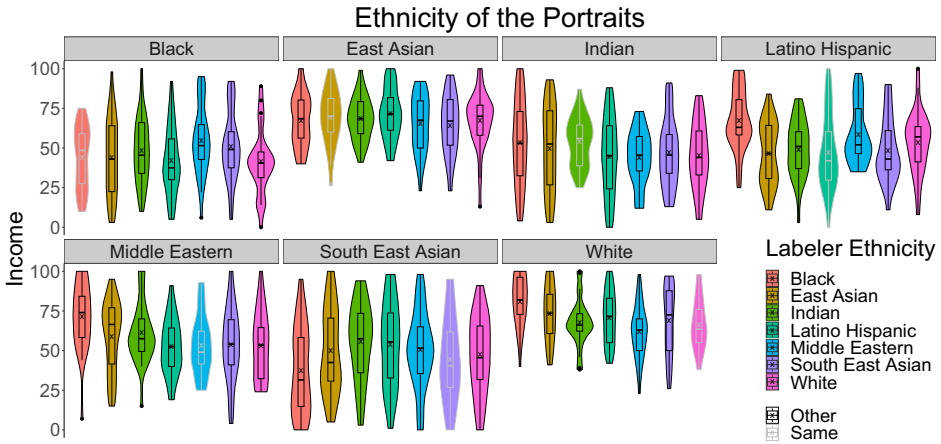


Figure 5. Estimated income as a function of LABELER_{Ethnicity} and PORTRAIT_{Ethnicity}. Grey borders indicate the cases where LABELER_{Ethnicity} and PORTRAIT_{Ethnicity} match.

Table 2. The two-way ANOVA results for the Income estimates and the four stereotype variables for LABELER_{Ethnicity} and PORTRAIT_{Ethnicity}.

	LABELER _{ethnicity}					PORTRAIT _{ethnicity}					L × 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

nificantly 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. 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 our two research questions. 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 a significant impact on perceptions of stereotypes and income labels¹³. Our results show that LABELER_{Ethnicity} significantly impacts status, PORTRAIT_{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 estimates were impacted by LABELER_{Ethnicity} and the PORTRAIT_{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 [44, 41] and, in particular, how gender stereotypes influence perception [42]. 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 [54]. These results are in line with recent findings in CSCW by Goyal et al. [5] demonstrating that toxicity labels for online content are influenced by labeler self-identification.

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 [55, 56, 57], and Levay et al. [58] found that over 70% of MTurkers identify as white. Consequently, the status quo of gathering labels on 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., [8]), but typically require knowledge of 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 [59]. 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. [4] suggests expanding data documentation and incorporating social contexts, which could be an important step toward ensuring fair, transparent data pipelines.

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.

5.3. Limitations & Future Work

The most notable limitation of our study is 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 [60]. 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 additional scenarios that are relevant to specific tasks in machine learning. For example, labeler bias may be relevant in detecting and classifying objects in autonomous driving tasks. 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 ground truth. Similarly, 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. Future studies should investigate 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 and that researchers and practitioners reconsider the status quo for data labeling.

References

- [1] Hagey K, Horwitz J. Facebook Tried to Make Its Platform a Healthier Place. It Got Angrier Instead. *Wall Street Journal*. 2021. Available from: <https://www.wsj.com/articles/facebook-algorithm-change-zuckerberg-11631654215>.
- [2] Eickhoff C. Cognitive Biases in Crowdsourcing. In: *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining. WSDM '18*. New York, NY, USA: Association for Computing Machinery; 2018. p. 162–170. doi:10.1145/3159652.3159654.
- [3] Miceli M, Schuessler M, Yang T. Between Subjectivity and Imposition: Power Dynamics in Data Annotation for Computer Vision. *Proc ACM Hum-Comput Interact*. 2020 oct;4(CSCW2). doi:10.1145/3415186.
- [4] Miceli M, Posada J, Yang T. Studying Up Machine Learning Data: Why Talk About Bias When We Mean Power? *Proceedings of the ACM on Human-Computer Interaction*. 2022 Jan;6(GROUP):34:1-34:14. doi:10.1145/3492853.
- [5] Goyal N, Kivlichan ID, Rosen R, Vasserman L. Is Your Toxicity My Toxicity? Exploring the Impact of Rater Identity on Toxicity Annotation. *Proc ACM Hum-Comput Interact*. 2022 nov;6(CSCW2). doi:10.1145/3555088.
- [6] Sap M, Swayamdipta S, Vianna L, Zhou X, Choi Y, Smith NA. Annotators with Attitudes: How Annotator Beliefs And Identities Bias Toxic Language Detection; 2022. doi:10.48550/arXiv.2111.07997.
- [7] Bender EM, Gebru T, McMillan-Major A, Shmitchell S. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? In: *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency. FAccT '21*. New York, NY, USA: Association for Computing Machinery; 2021. p. 610–623. doi:10.1145/3442188.3445922.
- [8] Jiang H, Nachum O. Identifying and Correcting Label Bias in Machine Learning. In: *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics. PMLR*; 2020. p. 702-12. Available from: <https://proceedings.mlr.press/v108/jiang20a.html>.
- [9] Wauthier FL, Jordan M. Bayesian Bias Mitigation for Crowdsourcing. In: *Advances in Neural Information Processing Systems*. vol. 24. Curran Associates, Inc.; 2011. p. 1-9. Available from: <https://proceedings.neurips.cc/paper/2011/hash/0768281a05da9f27df178b5c39a51263-Abstract.html>.
- [10] Cohn T, Specia L. 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)*. Sofia, Bulgaria: Association for Computational Linguistics; 2013. p. 32-42. Available from: <https://aclanthology.org/P13-1004>.
- [11] Fiske ST. Stereotype Content: Warmth and Competence Endure. *Current Directions in Psychological Science*. 2018;27(2). doi:10.1177/0963721417738825.
- [12] Schwind V, Deierlein N, Poguntke R, Henze N. 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*. New York, NY, USA: Association for Computing Machinery; 2019. p. 1–12. doi:10.1145/3290605.3300591.

- [13] Marsden N, Haag M. Stereotypes and Politics: Reflections on Personas. In: Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems. CHI '16. New York, NY, USA: Association for Computing Machinery; 2016. p. 4017-31. doi:10.1145/2858036.2858151.
- [14] Karkkainen K, Joo J. FairFace: Face Attribute Dataset for Balanced Race, Gender, and Age for Bias Measurement and Mitigation. In: Winter Conference on Applications of Computer Vision (WACV). Waikoloa, HI, USA: IEEE; 2021. p. 1547-57. doi:10.1109/WACV48630.2021.00159.
- [15] Hagendorff T. The Ethics of AI Ethics: An Evaluation of Guidelines. *Minds & Machines*. 2020;30(1). doi:10.1007/s11023-020-09517-8.
- [16] Tsamados A, Aggarwal N, Cowls J, Morley J, Roberts H, Taddeo M, et al. The ethics of algorithms: key problems and solutions. *AI & Soc*. 2021. doi:10.1007/s00146-021-01154-8.
- [17] Mueller VC. Ethics of Artificial Intelligence and Robotics. *The Stanford Encyclopedia of Philosophy*. 2021. Available from: <https://plato.stanford.edu/archives/sum2021/entries/ethics-ai/>.
- [18] Schwartz O. Untold history of AI: Algorithmic bias was born in the 1980s. *IEEE Spectrum*. 2019. Available from: <https://spectrum.ieee.org/untold-history-of-ai-the-birth-of-machine-bias>.
- [19] Lee NT, Resnick P, Barton G. Algorithmic bias detection and mitigation: Best practices and policies to reduce consumer harms. Brookings Institute: Washington, DC, USA. 2019. Available from: <https://policycommons.net/artifacts/4141276/algorithmic-bias-detection-and-mitigation/4949849/>.
- [20] Kodiyan AA. An overview of ethical issues in using AI systems in hiring with a case study of Amazon's AI based hiring tool; 2019.
- [21] Zarsky T. Understanding Discrimination in the Scored Society. Rochester, NY: Social Science Research Network; 2015. ID 2550248. Available from: <https://papers.ssrn.com/abstract=2550248>.
- [22] Angwin J, Larson J, Mattu S, Kirchner L. Machine bias. *Ethics of Data and Analytics*. 2016:254-64. doi:10.1201/9781003278290-37.
- [23] Barera M. Mind the Gap: Addressing Structural Equity and Inclusion on Wikipedia. University of Texas at Arlington: University of Texas at Arlington; 2020. Available from: <http://hdl.handle.net/10106/29572>.
- [24] Silva FAMd, Pedrini H. Effects of cultural characteristics on building an emotion classifier through facial expression analysis. *Journal of Electronic Imaging*. 2015 Mar;24(2):023015. doi:10.1117/1.JEI.24.2.023015.
- [25] Scheuerman MK, Wade K, Lustig C, Brubaker JR. How We've Taught Algorithms to See Identity: Constructing Race and Gender in Image Databases for Facial Analysis. *Proceedings of the ACM on Human-Computer Interaction*. 2020 May;4(CSCW1):58:1-58:35. doi:10.1145/3392866.
- [26] Wall E, Blaha LM, Paul CL, Cook K, Endert A. Four Perspectives on Human Bias in Visual Analytics. In: *Cognitive Biases in Visualizations*. Cham: Springer International Publishing; 2018. p. 29-42. doi:10.1007/978-3-319-95831-6_3.
- [27] Hutson J, Taft JG, Barocas S, Levy K. Debiasing Desire: Addressing Bias & Discrimination on Intimate Platforms. *Proc ACM Hum-Comput Interact*. 2018;2(CSCW). doi:10.1145/3274342.

- [28] Wick M, Panda S, Tristan JB. Unlocking Fairness: a Trade-off Revisited. *Advances in neural information processing systems*. 2019;32. Available from: <https://proceedings.neurips.cc/paper/2019/file/373e4c5d8edfa8b74fd4b6791d0cf6dc-Paper.pdf>.
- [29] Hube C, Fetahu B, Gadiraju U. 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. New York, NY, USA: Association for Computing Machinery; 2019. p. 1-12. doi:10.1145/3290605.3300637.
- [30] Barbosa NaM, Chen M. 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*. CHI '19. New York, NY, USA: Association for Computing Machinery; 2019. p. 1-12. doi:10.1145/3290605.3300773.
- [31] Geva M, Goldberg Y, Berant J. Are We Modeling the Task or the Annotator? An Investigation of Annotator Bias in Natural Language Understanding Datasets; 2019. doi:10.48550/arXiv.1908.07898.
- [32] Parmar M, Mishra S, Geva M, Baral C. Don't Blame the Annotator: Bias Already Starts in the Annotation Instructions; 2022. doi:10.48550/arXiv.2205.00415.
- [33] Miceli M, Posada J. The Data-Production Dispositif; 2022. doi:10.48550/arXiv.2205.11963.
- [34] Engelmann S, Ullstein C, Papakyriakopoulos O, Grossklags J. What People Think AI Should Infer From Faces. In: *2022 ACM Conference on Fairness, Accountability, and Transparency*. FAccT '22. New York, NY, USA: Association for Computing Machinery; 2022. p. 128-41. doi:10.1145/3531146.3533080.
- [35] Fiske ST, Cuddy AJC, Glick P. Universal dimensions of social cognition: Warmth and competence. *Trends in cognitive sciences*. 2007;11(2). doi:10.1016/j.tics.2006.11.005.
- [36] Fiske ST, Cuddy AJC, Glick P, Xu J. A model of (often mixed) stereotype content: Competence and warmth respectively follow from perceived status and competition. *Journal of Personality and Social Psychology*. 2002;82(6). doi:10.1037/0022-3514.82.6.878.
- [37] Dovidio JF, Gaertner SL. Intergroup bias. In: *Handbook of social psychology*, Vol. 2, 5th ed. Hoboken, NJ, US: John Wiley & Sons, Inc.; 2010. p. 1084-121. doi:10.1002/9780470561119.socpsy002029.
- [38] Durante F, Tablante CB, Fiske ST. Poor but Warm, Rich but Cold (and Competent): Social Classes in the Stereotype Content Model. *Journal of Social Issues*. 2017;73(1). doi:10.1111/josi.12208.
- [39] Grigoryev D, Fiske ST, Batkhina A. Mapping Ethnic Stereotypes and Their Antecedents in Russia: The Stereotype Content Model. *Frontiers in Psychology*. 2019;10. doi:10.3389/fpsyg.2019.01643.
- [40] Diamantopoulos A, Florack A, Halkias G, Palcu J. Explicit versus implicit country stereotypes as predictors of product preferences: Insights from the stereotype content model. *J Int Bus Stud*. 2017;48(8). doi:10.1057/s41267-017-0085-9.
- [41] Lin MH, Kwan VSY, Cheung A, Fiske ST. Stereotype Content Model Explains Prejudice for an Envied Outgroup: Scale of Anti-Asian American Stereotypes. *Pers Soc Psychol Bull*. 2005;31(1). doi:10.1177/0146167204271320.

- [42] Cuddy AJC, Fiske ST, Glick P. Warmth and Competence as Universal Dimensions of Social Perception: The Stereotype Content Model and the BIAS Map. *Advances in experimental social psychology*. 2008;40:61-149. doi:10.1016/S0065-2601(07)00002-0.
- [43] Fiske ST. Intergroup Biases: A Focus on Stereotype Content. *Curr Opin Behav Sci*. 2015;3. doi:10.1016/j.cobeha.2015.01.010.
- [44] Cuddy AJC, Fiske ST, Kwan VSY, Glick P, Demoulin S, Leyens JP, et al. Stereotype content model across cultures: Towards universal similarities and some differences. *British Journal of Social Psychology*. 2009;48(1). doi:10.1348/014466608X314935.
- [45] Nicolas G, Bai X, Fiske ST. Comprehensive stereotype content dictionaries using a semi-automated method. *European Journal of Social Psychology*. 2021;51(1). doi:10.1002/ejsp.2724.
- [46] Kroon AC, Trilling D, Raats T. Guilty by Association: Using Word Embeddings to Measure Ethnic Stereotypes in News Coverage. *Journalism & Mass Communication Quarterly*. 2021;98(2). doi:10.1177/1077699020932304.
- [47] Fraser KC, Nejadgholi I, Kiritchenko S. Understanding and Countering Stereotypes: A Computational Approach to the Stereotype Content Model; 2021. doi:10.48550/arXiv.2106.02596.
- [48] McKee K, Bai X, Fiske S. Understanding Human Impressions of Artificial Intelligence; 2021. doi:10.31234/osf.io/5ursp.
- [49] Phinney JS. When we talk about American ethnic groups, what do we mean? *American Psychologist*. 1996;51(9). doi:10.1037/0003-066X.51.9.918.
- [50] Zhang J, Wu X, Sheng V. Learning from crowdsourced labeled data: a survey. *Artificial Intelligence Review*. 2016;46. doi:10.1007/s10462-016-9491-9.
- [51] Lugaresi C, Tang J, Nash H, McClanahan C, Uboweja E, Hays M, et al. MediaPipe: A Framework for Building Perception Pipelines; 2019. doi:10.48550/arXiv.1906.08172.
- [52] Abbey JD, Meloy MG. Attention by design: Using attention checks to detect inattentive respondents and improve data quality. *Journal of Operations Management*. 2017;53-56. doi:10.1016/j.jom.2017.06.001.
- [53] Greenhouse SW, Geisser S. On methods in the analysis of profile data. *Psychometrika*. 1959;24(2). doi:10.1007/BF02289823.
- [54] Fiske ST, Taylor SE. *Social Cognition* (2nd ed). New York, NY, England: McGraw-Hill Book Company; 1991. Available from: <https://psycnet.apa.org/record/1991-97723-000>.
- [55] Difallah D, Filatova E, Ipeirotis P. Demographics and Dynamics of Mechanical Turk Workers. In: *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining. WSDM '18*. New York, NY, USA: Association for Computing Machinery; 2018. p. 135–143. doi:10.1145/3159652.3159661.
- [56] Ipeirotis PG. *Demographics of Mechanical Turk*. Rochester, NY; 2010. 1585030. Available from: <https://papers.ssrn.com/abstract=1585030>.
- [57] Ross J, Irani L, Silberman MS, Zaldivar A, Tomlinson B. Who are the crowdworkers? shifting demographics in mechanical turk. In: *CHI '10 Extended Abstracts on Human Factors in Computing Systems. CHI EA '10*. New York, NY, USA: Association for Computing Machinery; 2010. p. 2863–2872. doi:10.1145/1753846.1753873.

- [58] Levay KE, Freese J, Druckman JN. The Demographic and Political Composition of Mechanical Turk Samples. *SAGE Open*. 2016;6(1). doi:10.1177/2158244016636433.
- [59] D'Ignazio C, Klein LF. *Data Feminism*. MIT Press; 2020. Available from: <https://mitpress.mit.edu/9780262547185/data-feminism/>.
- [60] González MJ, Cortina C, Rodríguez J. The Role of Gender Stereotypes in Hiring: A Field Experiment. *European Sociological Review*. 2019;35(2). doi:10.1093/esr/jcy055.