Inclusive Portraits: Race-Aware Human-in-the-Loop Technology

CLAUDIA FLORES-SAVIAGA, Civic AI Lab, Northeastern University, USA

CHRISTOPHER CURTIS, Civic AI Lab, Northeastern University, USA

SAIPH SAVAGE, Civic AI Lab, Northeastern University; Universidad Nacional Autonoma de Mexico (UNAM), MX

Fig. (1)   Overview of Inclusive Portait's functionality.

AI has revolutionized the processing of various services, including the automatic facial verification of people. Automated approaches have demonstrated their speed and efficiency in verifying a large volume of faces, but they can face challenges when processing content from certain communities, including communities of people of color. This challenge has prompted the adoption of "human-in-the-loop" (HITL) approaches, where human workers collaborate with the AI to minimize errors. However, most HITL approaches do not consider workers' individual characteristics and backgrounds. This paper proposes a new approach, called Inclusive Portraits (IP), that connects with social theories around race to design a racially-aware human-in-the-loop system. Our experiments have provided evidence that incorporating race into human-in-the-loop (HITL) systems for facial verification can significantly enhance performance, especially for services delivered to people of color. Our findings also highlight the importance of considering individual worker characteristics in the design of HITL systems, rather than treating workers as a homogenous group. Our research has significant design implications for developing AI-enhanced services that are more inclusive and equitable.

1   INTRODUCTION

Facial verification is the process by which a system confirms a person's identity via a facial biometric scan [53]. Facial verification is a one-to-one mapping that responds to the question: *"Is the person who they say they are?"* Given its success in verifying a large number of people, facial verification has been integrated into a number of services. For example, several universities use facial verification for exam identity validation [5]; banks and hospitals use facial

verification to give people access to online banking or medical care [9]; governments use facial verification to give citizens access to protected buildings, or to classified government information [44].

Current facial verification systems usually use AI (machine learning models) to verify people based on their facial features [10, 92, 94]. AI-based facial recognition methods have gained popularity due to their ability to be easily deployed at scale, and their high accuracy in verifying faces [40, 66]. Automated facial recognition methods are particularly effective when used under optimal lighting conditions, with frontal faces, and with high-quality images [65]. However, previous research demonstrated that past automated facial verification systems exhibited biases and errors [16, 74]. For instance, they might falsely identify people of color as criminal suspects at higher rates than white people [17, 41, 103]. In general, at least historically, automated approaches for verifying people can struggle to accurately recognize people of color [28, 69]. One major factor is the lack of diversity in training data used to develop facial verification algorithms [84]. When datasets do not include a representative sample of people with darker skin tones, the algorithms may struggle to correctly identify and verify these individuals [78]. Additionally, some algorithms may be biased because they are trained on data that over-represents certain demographics or rely on features that are more prominent in certain groups, leading to inaccurate results for people of color [23, 58]. Finally, poor lighting and low-quality images can also impact the system's ability to detect and recognize facial features accurately [43, 58], which is usually more prevalent in images of people with darker skin tones [17].

One approach for mitigating the above concerns is technical: the creation of more accurate and more fair algorithms. Along these lines, the U.S. National Institute of Standards and Technology created a program for evaluating face recognition technologies [2], and indeed today's leading proprietary systems are significantly better than the past [54]. In parallel, the limitations above have also led researchers to integrate human workers into the machine learning pipeline, to assist the automated methods in the facial verification process [71]. Humans may be better at recognizing other humans under a wide range of settings and can also recognize people as they evolve and change through time [61, 62]. Facial verification systems may therefore improve further through the integration of human workers [63]. Such approaches are typically known as "human-in-the-loop" (HITL) [100]. Examples of HITL approaches within this context involve crowdworkers helping to verify unknown soldier portraits [63], or detecting fraudulent passports [96]. While HITL systems can potentially mitigate some of the biases found in human-only or automated-only systems, they are still susceptible to biases and errors [46]. Errors in HITL systems often arise because human workers may struggle to accurately identify people from different races [93], particularly if they have limited exposure to facial expressions and features that are unique to those races [67, 101]. The combination of inaccuracies and biases in facial verification systems has resulted in people of color being disproportionately affected [17]. Hence, it is crucial to explore effective ways to integrate human and machine decision-making to address biases present in facial verification systems.

We believe that part of the problem is that most HITL systems either lack sufficient training of workers and/or usually overlook the diverse backgrounds of the human workers involved [4, 31, 35, 72]. To address the latter problem, "jury learning" techniques have emerged [35]. These techniques involve modeling workers and their backgrounds to better understand their perspectives and improve the quality of the labeling process [35]. However, these techniques lack consideration for the racial backgrounds of workers, which can be a significant limitation [45]. Racial nuances can be essential to ensure that workers' unique experiences are accurately represented in their labor.

This paper introduces a novel racially-aware human-in-the-loop (HITL) system called "Inclusive Portraits" (IP). One of the key features of IP is its "Same Race" mechanism, which is based on social theory research related to race [50, 97]. This research suggests that, at least without training, people are more effective when working with data and information from their own race. The novelty of IP lies in its ability to leverage this insight to improve the accuracy

and reliability of facial verification algorithms for people of color. By designing HITL systems that take into account individual worker characteristics, including race, IP can help to ensure that AI systems are more inclusive. We also discuss other approaches, such as more proactive training of workers, including inclusivity and diversity training, that are not considered here but could be opportunities for future study.

Context. Before proceeding with our paper outline, we now step back and provide a broad, industry-level context for our research. Commercial face recognition systems have advanced significantly over recent years, with NIST's Face Recognition Vendor Test (FRVT) providing significant motivation for commercial advancement along with rigorous evaluations of the capabilities of commercial systems [2]. Reflecting upon NIST's evaluations, leading proprietary systems have significantly fewer biases than past systems [54]. That said, (1) not all entities using facial recognition technologies will have access to proprietary systems—some entities may build upon less accurate, openly available facial recognition technologies. Further, (2) biases do manifest in other types of technologies, beyond facial recognition systems. And, (3) even if current commercial face recognition technologies are better than past technologies, there is always value in further reducing any remaining biases in the overall systems that use those technologies.

As we explore in this paper, human-in-the-loop systems (if designed well) have the potential to reduce remaining biases; the resulting system may have less bias than either machine-alone systems or human-alone systems. Our experiments are not with leading proprietary face recognition technologies, per the NIST FRVT results [2], but face recognition systems openly available to the research community. As such, we hope that our work is valuable to any entity building a composite system with openly available facial recognition technologies (1), that our results are valuable to future researchers and systems designers seeking to build systems leveraging other types of automated classification technologies (2), and that our results are valuable to entities with access to proprietary facial recognition technologies who wish to reduce any remaining biases even further.

Paper Outline. In the following sections, we provide an overview of research on bias in facial verification and discuss current work being done on human-in-the-loop (HITL) systems. We then introduce our novel racially-aware HITL system, Inclusive Portraits (IP), and describe its design. We then present the results of our evaluation of IP, which involved a study where crowdworkers used IP to verify the facial images of individuals from different racial groups. We compared the performance of IP with a state-of-the-art HITL approach and found that our racially-aware HITL system resulted in significantly higher accuracy in facial verification, especially for people of color. Finally, we discuss the implications and lessons that can be drawn from our research, including the importance of incorporating social theories of race into the design of AI systems as one path to ensure inclusivity.

Positionality Statement As HCI researchers, we are committed to designing and evaluating technology that considers diverse perspectives and experiences. We recognize that facial verification technology can pose difficulties for people of color, especially when it is not designed with their needs and concerns in mind. Despite these challenges, we believe that this technology has the potential to be valuable in certain contexts and to benefit people of color. Thus, our research aims to create more inclusive technologies that address and combat biases against people of color. Our team consists of authors from diverse backgrounds, including two Latinas, one of whom holds a leadership role in their institution related to diversity, equity, and inclusion. Additionally, our academic work includes incorporating technology ethics and inequities into curricula. We acknowledge our own privilege and biases and strive to address these issues through a critical lens and by actively seeking diverse perspectives and experiences. Our goal is to create technology that is inclusive, accessible, and equitable for everyone. By contributing to the development of technology that truly serves and supports all communities, we hope to make a positive impact. We believe that by designing technology with a focus on inclusivity, we can help reduce the harm caused by biases and create a more just and equitable world.

2 RELATED WORK

Literature related to our research includes three main pieces: (1) Facial Verification and Biases; (2) Human-in-the-loop systems; and (3) Race and Computing.

2.1 Facial Verification and Biases.

In a recent survey, Wang et al [94] describe the problem of facial verification as the prominent biometric technique for identity authentication. Facial verification has been widely used in many areas, such as military, finance, public security, and daily life [7, 15, 51, 60]. The field itself started to rise in popularity in the 1990s after the introduction of several breakthrough approaches, and the field continued to evolve, with deep learning methods taking the stage with aplomb in the early and mid-2010's [3]. Within the last decade the field has seen advances and the technology has become employed in law enforcement, social media, marketing, and more [83].

Face recognition is not a single type of technology, but rather a class of technologies. Face recognition can be divided into two types: face verification and face identification [34]. Face identification systems focus on identifying a single person out of a large set of people (1:N facial matching), while facial verification systems take two images and ask whether they are of the same person (1:1 facial matching). For both classes, it is important to consider the potential for bias. Indeed, there is a real and prevalent concern regarding the bias in the datasets and models that are a part of learning systems that involve human actors and the consequences for the populations these AI-powered systems are applied to [14, 26, 37]. One well-known bias is that some verification systems often have a higher identification accuracy for male faces than for female faces [85]. Scheuerman et al. also identify that in the case of human categories, race is an area of particular concern regarding bias in machine learning fairness literature [83]. Recent research showed that current machine-learning models have higher error rates when verifying the faces of US underrepresented populations, e.g., Asians or African-Americans [49, 73]. According to Phillips et al., [70], like humans, even state-of-the-art face recognition algorithms — or at least open, non-proprietary systems available for academic study — struggle with "other-race face" recognition (see NIST's FRVT results for proprietary systems [2]; the U.S. Department of Homeland Security has also tested proprietary algorithms [1]). Benjamin observed that technologies used in other countries may manifest biases differently, e.g., algorithms "developed in China, Japan, and South Korea recognized East Asian faces far more readily than Caucasians" [14, p. 112]. Cavazos et al. acknowledge that there could be different factors such as lack of algorithm training with representative faces, inaccurately represented sub-populations, quality of the training, or test photographs across subgroups [18]. Currently, there are various research efforts to address forms of racial bias in the machine learning pipeline. Balayan et al. produced a recent survey of practical techniques [11]. Scheuerman et al. explored how computer vision datasets are often at odds between social values and desired metrics as well as provided suggestions on how to reconcile these differences [82]. Wang et al. proposed a reinforcement learning model which was successful in mitigating the racial bias across image datasets [93].

These efforts have highlighted the importance of addressing racial biases in machine learning and computer vision datasets. However, this research has tended to focus on the development of new machine learning models and techniques for quantitatively mitigating bias rather than on leveraging HCI principles and social theories for advancing new tools and dataset creation paradigms. Indeed, much previous research on this topic has failed to consider and take advantage of modern theories around race and computing; as Scheuerman et al. put it *"despite increasing attempts to diversify databases, approaches remain simplistic and lacking in critical and social theories"* [83]. The contribution of this paper is in part to address this exact need. To enhance the accuracy of facial verification, especially for individuals of color,

we integrate insights from HITL (human-in-the-loop) system research and race and computation research to create a race-aware HITL system. While our technical evaluations are focused on openly available, non-proprietary facial verification systems, we believe that all future automated systems can benefit from the thoughtful consideration of how to include humans in the loop.

2.2 Bias in Crowdsourcing and Human-in-the-Loop Systems.

Crowdsourcing and HITL systems face the challenge of human biases, which can affect the accuracy and fairness of AI models [51, 93, 95]. Researchers have proposed various solutions to address biases in crowdsourcing and HITL systems. These include the development of web applications that promote interpersonal accountability to mitigate biases [90], a web application influenced by cognitive psychology to help experts distinguish between labels and reduce biases [61], and modeling crowdworkers to predict their decision-making process, enabling designers of AI systems to choose the right crowdworkers to avoid potential biases [35]. However, despite these efforts, there is still limited work on how to address biases in crowdsourcing and HITL systems, particularly biases that affect people of color. To create more accurate and unbiased outcomes, especially for these communities, this paper suggests incorporating ideas from race and computation into the design of HITL systems. By doing so, we can better understand how workers' different racial backgrounds can help to create solutions that are more inclusive and respectful of diverse perspectives.

2.3 Race and Computing.

Race is studied as a social construct [8, 12] — racial groups are a product of social construct, not nature, and vary across social groups and time [21]. In a computing context, prior work has adapted critical race theory for HCI, emphasizing that race is both real and socially constructed. It highlights how racial systems of oppression can be encoded and perpetuated in computing systems, such as racially biased facial recognition [17]), and that engaging with race meaningfully in a design process is crucial to combating racial bias [64, 87]. In addition to critical race theory as a productive framing for understanding sociocultural systems, we can use it as a tool for design and evaluation of systems (e.g., by tracing lines of power and hierarchy, by drawing out how technical design choices interact with, reproduce, and/or mitigate racial inequities, by incorporating an understanding of how other identities co-mingle with race/ethnicity, etc. — much in the same way that critical disability studies can be used to evaluate assistive technologies [52, 80]). In this paper, we use theories around race and computation to design HITL systems that are better adapted to different races.

3 INCLUSIVE PORTRAITS (IP)

Inclusive Portraits (IP) is a novel Human-in-the-Loop (HITL) system that utilizes social theories related to race to boost the accuracy of facial verification, as illustrated in Figure 1. The system has two main components: (1) an automated component, which receives pairs of facial images for verification and employs state-of-the-art machine learning techniques to complete the process; and (2) a human component, where crowdworkers come into play. For pairs of images where the automated component is uncertain, IP sends them to the human component for verification. This is where the system's uniqueness lies, as it introduces new "Racially Aware Labor Interfaces" that leverage the diversity of crowdworkers to improve the accuracy of facial verification.

3.1 IP's Automated Component.

The automated component of IP begins by receiving pairs of facial images from various datasets or sets of photos that require verification. These pairs are then sent to state-of-the-art facial verification models, as outlined in Table 1. It is worth noting that these AI models are all available for scientific study (excluding proprietary models, as noted in [2]).

Using these models, IP calculates a "confidence score" for each pair of faces. The score represents IP's level of certainty that the faces belong to the same individual. To determine this score, IP takes the median confidence score of each AI model for the given pair of images, creating an overall measure of confidence. Pairs with lower confidence scores are compiled into a list and passed on to the human component for further analysis. This process allows IP to leverage the strengths of both automated and human analysis, ultimately improving the accuracy of the verification process.

Variations. In a production system, designers may instead aim to minimize false negatives or false positives, depending on the specific goals of their system. If minimizing false negatives is the priority, only uncertain negative outputs of the automated system may be sent to the human component. Conversely, if minimizing false positives is the priority, only uncertain positive outputs may be sent for human verification.

Another approach would be to send all negative or all positive outputs to the human component for further analysis. For instance, all negatives (regardless of how certain) could be sent to human verifiers, who could then reclassify them as positives. This would lead to a system that reduces false negatives but increases false positives, resulting in a higher overall count of positives (all automated positives and all automated negatives with human positives). Such an approach may come at the cost of increased resource consumption or processing time. Ultimately, the decision regarding which approach to use will depend on the specific needs and constraints of the system being developed.

3.2 IP's Human Component.

IP's human component incorporates a racially aware interface to efficiently manage the coordination of crowdworkers for the facial verification of images where the automated component is uncertain. Unlike traditional human-in-the-loop interfaces that assign "anonymous" crowdworkers to verification tasks [63], IP's approach takes into account the diversity of workers' backgrounds and experiences. This results in a more robust and reliable verification process, especially when verifying the faces of people of color. By leveraging the diverse perspectives and expertise of the crowdworkers, IP's racially aware interface enables more accurate and inclusive verification of individuals.

3.2.1 Racially Aware Interface. Human-in-the-loop (HITL) interfaces often neglect the diverse backgrounds of the workers involved, which can result in AI systems that fail to capture the perspectives and experiences of certain groups [35, 72]. A promising approach to addressing this challenge is "jury learning," which involves modeling workers and their backgrounds to improve the quality of the labeling process [35]. However, a key limitation of this approach is its failure to account for workers' racial backgrounds. To overcome this limitation, IP introduces a racially aware interface that models workers based on their racial backgrounds. This innovative approach empowers workers to provide more tailored labeling that reflects the distinct needs and perspectives of their respective communities, including those of communities of color. The novelty of IP's racially aware interface lies in its unique annotation interface that incorporates a Same Race Mechanism. This novel mechanism is based on research that has shown that people are less effective when working with data and information from outside their own race [50, 97]. This component thus matches workers with labeling tasks related to their racial background [76]. In specific, IP's racially aware interface matches workers with facial verification tasks that involve faces from their same racial background. This approach allows workers to leverage their unique experiences, resulting in more accurate and racially sensitive facial verifications. By incorporating

workers' racial backgrounds in the task assignment process, IP ensures that the labeling process is more inclusive and representative of diverse communities.

In line with the best practices that recognize race as a social construct [86], IP follows established guidelines and requests first for crowdworkers to self-identify their racial identity. This approach is especially important in large-scale data collection efforts [27]. However, to ensure consistency and accuracy, it is recommended that we then code these identities to pre-existing categories or labels. By adopting these methods, IP effectively accounts for the nuances of race [19]. IP utilizes a matching process that pairs a crowdworker's racial background with a corresponding facial verification task. This ensures that the crowdworker is verifying images that align with their background and experiences, ultimately improving the accuracy and inclusivity of the labeling process. IP leverages the APIs of popular crowdsourcing platforms like Toloka to source workers from geographic regions that have a diverse population from different racial backgrounds.

4 EVALUATION

We study whether IP's racially aware interface can achieve higher accuracy in facial verification. We conducted an IRB-approved study to compare IP's accuracy with a state-of-the-art HITL system, in order to investigate this. The study took place from November 6th to November 15, 2022.

4.1 METHODS

To ensure the validity and reliability of our study, we meticulously recruited human participants and prepared the data, AI models, and HITL approaches used in our evaluation. Note that our participants were the crowdworkers in both the IP and control HITL conditions. Our data included facial photos from diverse racial backgrounds to be verified. The AI models were the first line of verification, sending any data for which the AI models were uncertain or had doubts to the HITL systems. These HITL systems included IP and the control HITL, which we utilized to compare and contrast our approach. Our comprehensive preparation ensured a thorough and accurate evaluation of our system's performance in facial verification.

Data. To study and compare the effectiveness of IP and contrast it with a state-of-the-art HITL system, we utilized the "Racial Faces in the Wild" (RFW) balanced dataset [59]. This dataset contains images that are pre-labeled with the primary race present in each image. The RFW dataset is comprised of a 7K identities per race (Asian, Black or African-American, Caucasian and Indian (2,984 pairs of images). Fig. 2 shows examples of images from the RFW dataset. The RFW dataset includes both negative pairs (two pictures of different individuals) and positive pairs (two pictures of the same individual), making it possible to study the false positives and false negatives of a given approach. This is particularly important for uncovering biases in facial verification systems [95]. Furthermore, the RFW dataset stands out from other datasets because it provides a broad range of racial labels and was purposefully designed to facilitate investigations of biases in AI systems. By working with the RFW dataset and developing racially-aware HITL systems, we hope to contribute towards the development of additional approaches for more equitable and accurate facial verification technologies.

Participants (Crowdworkers). A total of 160 participants were recruited for the study, evenly divided into two groups: 80 in the control HITL condition and 80 in the IP's condition. These participants were crowdworkers recruited through the Toloka crowdsourcing platform[1].
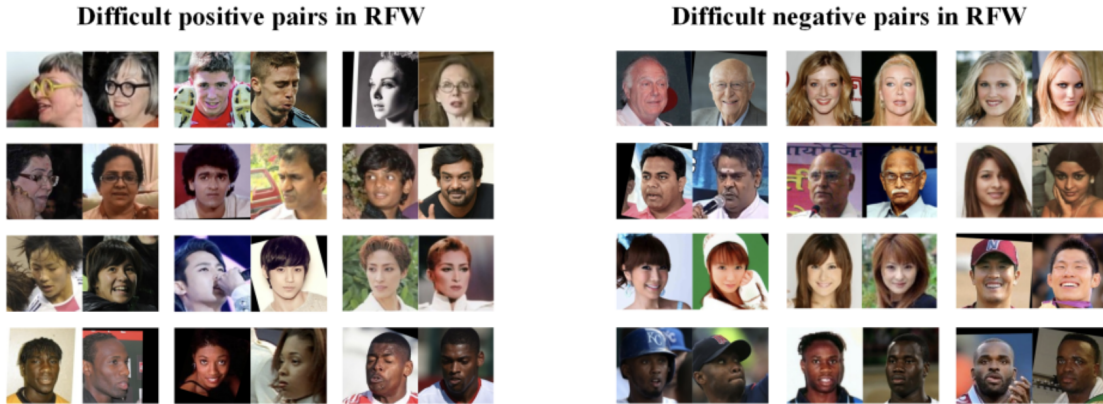
---

[1]https://toloka.ai/tolokers

Fig. (2)   Examples of images in our study, which included pairs of facial photos of people from different races. Note that across races, there were no large differences in pose, age, or gender distribution in the facial images we used for our field experiments.

In our study, it was essential to recruit crowdworkers from diverse racial backgrounds since we were comparing a racially aware HITL interface with a traditional HITL approach. To ensure unbiased and representative results, we aimed to have an equal number of crowdworkers from each racial group in both the control and IP conditions. Because race is a social construct [14, 68], we followed established guidelines and requested that crowdworkers self-identify their racial identity [86]. We then coded these identities into the labels used by the RFW dataset, which included the labels: "Black or African American", "Asian", "Caucasian", and "Indian." This approach is important in large-scale data collection efforts [27], and helps to effectively account for the nuances of race [19]. Additionally, we only selected crowdworkers who had no prior experience with facial verification tasks to control for confounding factors. Workers who were not selected were compensated for answering our recruitment survey.

AI Models. HITL systems in this domain of facial verification typically involve an AI component that performs the initial verification, followed by a human component that verifies and assists the AI with cases that the AI is uncertain about.[2] Given this, we first prepared various AI models to perform automatic verification and then sent the images that all the AI models struggled with to our HITL systems (control and IP) for further verification. We evaluated eight open facial verification models (ArcFace, Facenet, SFace, VGG-Face) and one proprietary model: Amazon Rekognition. We utilized the RFW dataset of facial photos and fed it to the AI models. Images that all the models struggled with were then sent to our HITL systems (control and IP) for verification. To ensure a fair comparison of how the HITL systems performed with different races, we made sure that the sample size of photos associated with each race was the same size across groups. In cases where a certain race had more errors, we randomly selected images until the smallest sample size was achieved, and we also followed best practices from HITL systems in terms of the number of images we gave to crowdworkers to verify [61].

We sent the HITL systems a sample of Asian (2000 pairs of images), Black or African-American (2000 pairs of images), Caucasian (2000 pairs of images), and Indian (2000 pairs of images) photos for verification. Note that both control and IP were given the exact same images for verification.

---

[2]See earlier discussion of possible variations in HITL systems, e.g., systems that seek to maximize positives might send all negatives (and only negatives) to human verifiers.

Additionally, we considered it valuable to study the accuracy of each AI model to get a better sense of the information ecosystem present in this space. To achieve this, we used the RFW testing dataset of facial photos and studied whether a given AI model correctly verified positive image pairs and rejected negative image pairs for each racial group. This allowed us to determine how well each AI model performed for a given race. It is worth highlighting that we did not utilize the RFW dataset to train the AI models that we tested. Our decision was based on the fact that training the models to reduce bias is beyond the scope of this research. Instead, we focused on studying how we could improve HITL systems by being racially aware.

IP & Control HITL. To compare IP with state of the art HITL approaches, we used a between-subjects study design and considered the following conditions:

- **Control Condition**. In this condition, we recruited 80 crowdworkers from different races (African-American, Asian, Caucasian, and Indian), to perform face verification tasks using a traditional HITL interface. The participants were divided into four groups, each consisting of 20 workers from the same race. They verified a total of 32 images (8 images per race), sourced from the photos that the AI models were uncertain about. This condition served as the baseline for comparison with the Inclusive Portraits condition.
  Note that to setup these conditions, we followed best practices from HITL systems in terms of the number of images we gave to crowdworkers. We limited the number of images each worker could verify to 32 to ensure they remained focused and attentive to the task at hand. This also helped to minimize any potential biases or fatigue that could arise from doing repetitive work [42], in this case verifying a large number of images. Additionally, we provided clear instructions and guidelines for the task, and we monitored the workers' performance to ensure the quality of the annotations. Overall, we followed standard best practices in HITL systems to ensure that our study was conducted in a rigorous and reliable manner.
- **Inclusive Portraits Condition**. To make a fair comparison between the Inclusive Portraits (IP) system and the traditional HITL system, we recruited 80 crowdworkers, with 20 from each race, to perform facial verification tasks using IP. We ensured that the workers were assigned a task that involved verifying a similar number of images as in the control group. However, in this scenario, all the images belonged to the same race as the worker's self-identified race. The crowdworkers completed the verification process using the IP system, which was aware of their racial background. This enabled us to compare the performance of the racially aware IP system with the traditional HITL interface.

To prevent bias and ensure reliable data, both conditions used the same interface, see Figure 3. Additionally, we implemented several measures to ensure that participants in our experiment had no prior exposure to the verification task. We meticulously screened participants to verify their lack of prior exposure to this task. Quality control questions were included to identify and exclude participants who were not fully engaged in the task. All workers passed these control questions. To compensate participants for their time and effort, we pay them $10 USD, which exceeds the US minimum wage [38], as the entire experiment took 25 minutes.


Fig. (3)   Screenshots of the interface. The facial images are from the RFW dataset (described below), which allows for images to be included in papers.

4.2 RESULTS.

Our study begins by presenting the accuracy of state-of-the-art AI models in verifying photos of people of color. We then highlight the limitations of these systems and proceed to compare a state-of-the-art HITL system for facial verification with IP's performance. It is important to note that the HITL systems were tested using the same data that the AI models struggled with, i.e., the photos in which they did not have confidence in their verification.

Results: AI Systems. To begin our study, we analyzed the accuracy of cutting-edge AI models in verifying faces from diverse races using the RWF dataset. To evaluate the performance of each AI model, we computed its accuracy for each race by dividing the number of correct predictions by the total number of images provided to the AI model for that specific race. As shown in Table 1, the AI models in general struggled with verifying faces of people of color. Specifically, ArcFace had an accuracy of only 25% for Asian faces, while it had an accuracy of 80% for Caucasian faces. This inconsistency was also observed in other AI models, such as Facenet, which performed poorly on Indian and African-American photos. Notably, Amazon's Recognition AI model had the best overall performance. However, it also exhibited a bias towards better accuracy on photos from Caucasian individuals.

These results are consistent with previous research findings [95], and highlight the challenges that AI models still face when it comes to accurately verifying faces of individuals of color.

Races (Photos)	ArcFace	Facenet	SFace	VGG-Face	Rekognition
African-American	30%	45%	75%	75%	70%
Asian	25%	75%	75%	75%	70%
Caucasian	80%	75%	90%	90%	100%
Indian	45%	40%	80%	80%	70%
Median (AI Model)	30%	45%	75%	75%	80%

Table (1)   Racial bias in open deep facial verification systems. Verification accuracy (65%) evaluated on 500 difficult pairs per race of RFW database are given. See NIST's FRVT [2] for an evaluation of proprietary facial verification systems; commercial deployments, which use facial verification technologies, may also have more advanced methods of curating image quality, which can impact results.

Results: HITL Systems. To understand the performance of our control HITL condition and Inclusive Portraits (IP), we presented the images that were misclassified by (at least one of) the automated systems. In Fig. 4, we present a summary of the median accuracy of these two HITL systems across different races. Table 3, presents an overview of the accuracy and percentage difference in facial verification of different racial groups, using Control HITL and IP conditions.

We observed that IP outperformed the state-of-the-art HITL system across races. To determine whether the differences in accuracy between IP and the control HITL system were statistically significant, we conducted a non-parametric analysis of variance. First, we tested the normality of the accuracy distributions in each condition using the Shapiro-Wilk test. The results showed that both distributions were non-normal, with p-values less than .05. We conducted the Kruskal-Wallis test, a non-parametric alternative to the one-way ANOVA, to study whether the differences in accuracy between these two condictions were significant, see Table 2.

Worker, Race	Statistic	p-value
Caucassian	1.3276	0.2492
Indian	5.6366	0.0176
African	4.3395	0.0372
Asian	4.0060	0.0453

Table (2)   Results of the Kruskal-Wallis Test comparing the accuracy of control and IP for face verification tasks.

Our study demonstrates that there is a significant improvement in accuracy for people of color when using the IP system, as indicated by the positive percentage difference for each race. Specifically, the IP system resulted in a higher
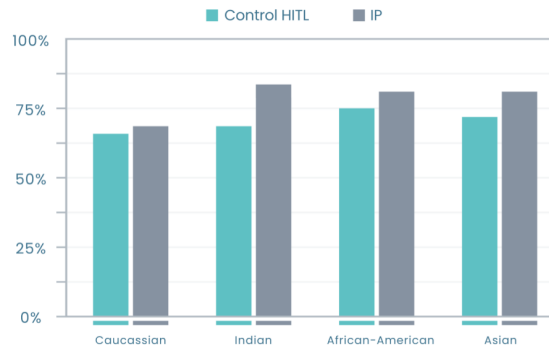
Fig. (4)   Overview of the median accuracy of the HITL control group and IP. The X-axis represents the different races of the crowdworkers performing the verification, while the Y-axis represents their accuracy.

	Control HITL	IP	Percentage
	Accuracy	Accuracy	Difference
African-American	75.00%	81.25%	+ 8.33%
Asian	71.88%	81.25%	+13.04%
Caucasian	65.63%	68.75%	+ 4.75%
Indian	68.75%	84.38%	+ 22.73%

Table (3)   Overview of the accuracy and percentage difference in facial verification of different racial groups, using Control HITL and IP conditions.

accuracy rate for Indian (+22.73%), African American (+8.33%), and Asian (+13.04%) individuals, compared to the control system. The Kruskal-Wallis test statistics demonstrated that this accuracy improvement is statistically significant with values of 4.3395 (p-value 0.0372), 5.6366 (p-value 0.0176), and 4.0060 (p-value 0.0453), respectively. However, for the Caucasian race, there was no significant difference in accuracy between the control HITL condition and IP, as shown by the Kruskal-Wallis test statistics of 1.3276 (p-value 0.2492). This suggests that the IP system can effectively enhance accuracy in facial recognition, especially for people of color.

Our study highlights the potential benefits of using a racially aware HITL system like Inclusive Portraits for improving accuracy in facial verification tasks involving people of color. The fact that we did not observe an improvement in accuracy for the Caucasian race may be attributed to the existing bias towards this group in facial verification systems [17], resulting in a higher quality dataset and more experienced crowdworkers for this group. However, our results also emphasize the importance of having a diverse and inclusive workforce to mitigate bias in AI systems. By utilizing workers with different backgrounds and experiences, Inclusive Portraits was able to improve accuracy for people of color, demonstrating the potential of a more inclusive approach to HITL.

5   DISCUSSION

In previous research, there have been attempts to combine human and machine intelligence to improve facial verification [63, 71]. Our study builds on this work by investigating how human-machine collaborations can be improved by taking into account workers' racial backgrounds. We introduce a racially aware HITL system, Inclusive Portraits, which significantly improves the accuracy of facial verification among untrained workers. By incorporating social theories on race into the design of human-in-the-loop systems [88, 99], we can optimize and improve the outcomes of workers.

To further enhance the effectiveness of Inclusive Portraits, future work could explore the integration of training materials [20, 24, 102], combined with racially aware mechanisms to achieve higher quality labor [50]. Our findings emphasize the value of considering the racial diversity of crowdworkers in facial verification tasks, especially in today's society where diversity and inclusion are increasingly valued. As improved facial verification technology can provide greater access to personalized services, Inclusive Portraits' racially aware interface is a valuable tool for enhancing the accuracy and quality of facial verification tasks.

5.1   Design Implications.

To expand the applicability of our study, we believe that future research should explore how architectures that consider workers' diverse backgrounds could be integrated into other human-in-the-loop systems, not limited to facial verification [35]. The "Circuit of Culture" framework [25, 36, 39, 47] could serve as a useful tool for researchers to identify how workers' diverse backgrounds can influence the design of artifacts, especially in three key areas: (1) data representation, (2) data consumption, and (3) regulation.

1. Data Representation. Representation is the creation of meaning through language and signs[13, 22, 25], which can be shaped by codes and symbols unique to specific communities, and are constantly evolving [13, 91]. We believe that designing human-in-the-loop (HITL) systems to coordinate crowdworkers to provide relevant data representations for end-users with similar backgrounds presents an opportunity for creating more inclusive and equitable AI systems. In natural language processing (NLP) and recommendation systems, HITL interfaces can be designed to coordinate crowdworkers from specific communities to provide translations, define relevant terms, or offer recommendations that are sensitive to their needs and perspectives. Incorporating the backgrounds of crowdworkers into the design of HITL systems has the potential to create more representative and community-driven AI solutions.

2. Data Consumption. Hall's concept of the "decoding" of a message by its public, and the Circuit of Culture's idea of product meaning being shaped by consumption [36, 81], both highlight the importance of involving end-users in the process of defining data representations. We believe there is value in bringing these ideas into HITL systems. Such an approach could help address biases and gaps in the definitions that crowdworkers' and designers' have, and the ones consumers have, and consequently create more representative labelled datasets [13, 25]. However, conflicts may arise when consumers, crowdworkers, and designers have differing perspectives, so it is important to find ways to address these conflicts, as well as consider ways to prevent harmful or exploitative data representations, including through connecting with related cybersecurity research [6, 79].

3. Regulation. To ensure that the development and deployment of racially-aware technologies are done in an ethical and responsible way, it is essential to involve a diverse group of stakeholders, including members of the communities that these technologies are intended to serve, in the regulatory process [13, 25, 36]. As previous researchers have noted [31], there is a need to explore designing crowdsourcing systems that take into account the backgrounds of crowdworkers. Taking into account crowdworkers backgrounds is important to provide input about the type of regulations needed in their community, ensuring that they are effective, equitable, and sensitive to the needs of the communities. Future research could also explore how to govern and design independent oversight bodies to effectively monitor the use of racially-aware technologies to ensure that these technologies are used in a responsible and ethical manner, while promoting social justice, respecting diversity, and protecting individual rights.

5.1.1   Utilizing Own-Race Bias in Design Choices: Implications and Considerations. The own-race bias, a well-documented phenomenon where individuals tend to better remember faces of their own race than those of other races, has important

implications for the design and implementation of facial verification systems [50, 97]. Our research has demonstrated that incorporating this knowledge into crowdsourcing platforms can lead to more effective design choices. However, it is important to carefully consider the ethical and logistical implications of federating and pooling workers by their race. Prior research on cross-race bias suggests that training methods can be used to mitigate this bias and reduce the disadvantage faced by outgroup workers [50, 89]. Employers could use this knowledge to provide training to specific groups of workers who can then act as an auditing or control group or be used when there are not enough task workers to fill desired groups. Nonetheless, it is important to recognize that such strategies may not fully eliminate the own-race bias but rather aim to minimize its impact. Therefore, organizations must continually monitor the performance of their systems and make adjustments as necessary.

5.1.2   Crowdworker Privacy and Racially Aware HITL Systems. It is essential to consider worker privacy when designing and deploying racially-aware human-in-the-loop systems. The collection and processing of sensitive personal information, including race, can raise significant privacy concerns for workers involved in such systems. In many cases, workers may not be aware that they are contributing to the development of racially-aware systems or how their personal information is being used. Moreover, workers may face risks such as stigmatization or discrimination based on their participation in such systems. To mitigate these concerns, it is important to implement clear and transparent policies around the collection, storage, and use of worker data but also respect a user's expectations for how their data will be used [75]. For example, workers should be informed about the nature of the task, how their data will be used, and the steps being taken to protect their privacy. Additionally, organizations should implement strict data security measures and limit access to personal information to authorized personnel only. Further, anonymity and pseudonymity can be employed as ways to reduce the potential for workers' personal information to be associated with their work [48]. Finally, it is important to establish mechanisms for workers to report privacy violations or concerns and provide a means for workers to opt-out of the system if they choose to do so. By considering worker privacy concerns and implementing appropriate measures, organizations can help ensure that racially-aware human-in-the-loop systems are developed and deployed in a responsible and ethical manner.

Limitations and Future work Although our research sheds light on the effects of own-race bias on facial verification, it is important to acknowledge its limitations and the need for future work. In some settings, the race of individuals may not be known and, in such cases, it may not be possible or appropriate to infer race algorithmically; indeed algorithmic approaches for race detection may also exhibit biases [98]. Further, the people being verified may be multiracial and the worker pool may not include people with the exact same demographic backgrounds. Additionally, exploring the impact of various worker backgrounds, such as age [56], gender [57], personality [55] and expertise [29, 32, 33], on the effectiveness of facial verification and other types of tasks would provide valuable insights. Additionally, examining the role of monetary incentives on crowdworkers' performance in facial verification tasks could also be beneficial, as it has been shown to affect task quality in other contexts [30, 77]. A comprehensive understanding of how these factors interact with own-race bias can facilitate the development of more equitable and accurate facial verification systems. Therefore, our research serves as a crucial initial step in creating more inclusive and fair systems, not just for face recognition systems, where industry and government continue to advance technically [2], but for AI systems as a whole. Future work may also study the role of workers in different types of classification systems as surveyed in Section 3.1, e.g., a system designed to minimize false negatives at the expense of a greater number of false positives.

6 CONCLUSION

Humans and technologies each have their own strengths and their own weaknesses, and by combining both via human-in-the-loop systems it is possible to leverage the best of both and produce more equitable results. This paper explores how to push the human-in-the-loop frontier further and, in particular, explores the impact of different methods for involving humans in the loop.

Specifically, we study both a standard and an advanced way of including humans in a composite human-in-the-loop verification system. We experimented with our approach using openly available facial recognition technologies, which are an example of an automated technology with known biases (we did not experiment with the most advanced, proprietary facial recognition technologies according to NIST's FRVT [2]). Although leveraging open systems means that the underlying facial recognition technologies we studied do not have the same performance as leading proprietary systems, leveraging these openly available systems allowed us to explore human-in-the-loop composition methods using resources generally available to the academic community and hence resources that other researchers can build upon. Through our research, we developed Inclusive Portraits, a race-aware human-in-the-loop system that outperforms a state-of-the-art HITL system. Our system is grounded in social theory research on race, and our online study demonstrated its efficacy, particularly in the verification of non-Caucasian faces. Given our findings, we encourage future researchers studying and building advanced human-in-the-loop systems to consider fully the impacts and biases of individual humans in human-in-the-loop systems.

REFERENCES

[1] [n. d.]. 2022 Biometric Technology Rally - The Maryland Test Facility. https://mdtf.org/Rally2022.

[2] [n. d.]. FRVT 1:1 Verification - NIST. https://pages.nist.gov/frvt/html/frvt11.html.

[3] Ghaliya Al Farsi, Jasiya Jabbar, and Ragad M Tawafak. 2019. A Review on models of human face verification techniques. In *2019 International Conference on Fourth Industrial Revolution (ICFIR)*. IEEE, 1–5.

[4] Adriana Alvarado Garcia, Juan F Maestre, Manuhuia Barcham, Marilyn Iriarte, Marisol Wong-Villacres, Oscar A Lemus, Palak Dudani, Pedro Reynolds-Cuéllar, Ruotong Wang, and Teresa Cerratto Pargman. 2021. Decolonial Pathways: Our Manifesto for a Decolonizing Agenda in HCI Research and Design. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–9.

[5] Neetu Anand, Tripti Sharma, and Kumar Gaurav. 2022. Face Recognition for Examinee Verification. *International Journal of Next-Generation Computing* 13, 3 (2022).

[6] Mohd Anwar, Hilda Goins, Tianyang Zhang, and Xiangfeng Dai. 2022. Human Elements in Machine Learning-Based Solutions to Cybersecurity. In *Smart and Intelligent Systems*. CRC Press, 137–148.

[7] Okereke Eze Aru and Ihekweaba Gozie. 2013. Facial verification technology for use in ATM transactions. *American Journal of Engineering Research (AJER)* 2, 5 (2013), 188–193.

[8] Lee David Baker. 1994. *The role of anthropology in the social construction of race, 1896-1954*. Temple University.

[9] TV Bakunova, EA Trofimova, and EV Lapteva. 2019. Biometrics as a method of information security in the banking sector digitalization. In *International Scientific and Practical Conference on Digital Economy (ISCDE 2019)*. Atlantis Press, 929–934.

[10] Stephen Balaban. 2015. Deep learning and face recognition: the state of the art. *Biometric and surveillance technology for human and activity identification XII* 9457 (2015), 68–75.

[11] Agathe Balayn, Christoph Lofi, and Geert-Jan Houben. 2021. Managing bias and unfairness in data for decision support: a survey of machine learning and data engineering approaches to identify and mitigate bias and unfairness within data management and analytics systems. *The VLDB Journal* 30, 5 (2021), 739–768.

[12] Michael J Bamshad and Steve E Olson. 2003. Does race exist. *Scientific American* 289, 6 (2003), 78–85.

[13] Shaowen Bardzell, Jeffrey Bardzell, Jodi Forlizzi, John Zimmerman, and John Antanitis. 2012. Critical design and critical theory: the challenge of designing for provocation. In *Proceedings of the Designing Interactive Systems Conference*. 288–297.

[14] Ruha Benjamin. 2019. Race after technology: Abolitionist tools for the new jim code. *Social forces* (2019).

[15] Cynthia L Bennett, Cole Gleason, Morgan Klaus Scheuerman, Jeffrey P Bigham, Anhong Guo, and Alexandra To. 2021. "It's Complicated": Negotiating Accessibility and (Mis) Representation in Image Descriptions of Race, Gender, and Disability. In *Proceedings of the 2021 CHI Conference*

on Human Factors in Computing Systems. 1–19.

[16] Austin Blanton, Kristen C Allen, Timothy Miller, Nathan D Kalka, and Anil K Jain. 2016. A comparison of human and automated face verification accuracy on unconstrained image sets. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops.* 161–168.

[17] Joy Adowaa Buolamwini. 2017. *Gender shades: intersectional phenotypic and demographic evaluation of face datasets and gender classifiers.* Ph. D. Dissertation. Massachusetts Institute of Technology.

[18] Jacqueline G Cavazos, P Jonathon Phillips, Carlos D Castillo, and Alice J O'Toole. 2020. Accuracy comparison across face recognition algorithms: Where are we on measuring race bias? *IEEE transactions on biometrics, behavior, and identity science* 3, 1 (2020), 101–111.

[19] Yiqun T Chen, Angela DR Smith, Katharina Reinecke, and Alexandra To. 2023. Why, when, and from whom: considerations for collecting and reporting race and ethnicity data in HCI. In *CHI Conference on Human Factors in Computing Systems.*

[20] Chun-Wei Chiang, Anna Kasunic, and Saiph Savage. 2018. Crowd coach: Peer coaching for crowd workers' skill growth. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 1–17.

[21] Rodney D Coates. 2004. *Race and ethnicity: Across time, space and discipline.* Brill.

[22] Patricia A Curtin and T Kenn Gaither. 2005. Privileging identity, difference, and power: The circuit of culture as a basis for public relations theory. *Journal of Public Relations Research* 17, 2 (2005), 91–115.

[23] Roxana Daneshjou, Mary P Smith, Mary D Sun, Veronica Rotemberg, and James Zou. 2021. Lack of transparency and potential bias in artificial intelligence data sets and algorithms: a scoping review. *JAMA dermatology* 157, 11 (2021), 1362–1369.

[24] Steven Dow, Anand Kulkarni, Brie Bunge, Truc Nguyen, Scott Klemmer, and Björn Hartmann. 2011. Shepherding the crowd: managing and providing feedback to crowd workers. In *CHI'11 Extended Abstracts on Human Factors in Computing Systems.* 1669–1674.

[25] Paul Du Gay, Stuart Hall, Linda Janes, Anders Koed Madsen, Hugh Mackay, and Keith Negus. 2013. *Doing cultural studies: The story of the Sony Walkman.* Sage.

[26] Virginia Eubanks. 2018. *Automating inequality: How high-tech tools profile, police, and punish the poor.* St. Martin's Press.

[27] Huaying Fang, Qin Hui, Julie Lynch, Jacqueline Honerlaw, Themistocles L Assimes, Jie Huang, Marijana Vujkovic, Scott M Damrauer, Saiju Pyarajan, J Michael Gaziano, et al. 2019. Harmonizing genetic ancestry and self-identified race/ethnicity in genome-wide association studies. *The American Journal of Human Genetics* 105, 4 (2019), 763–772.

[28] Rachel S Fleischer. 2020. Bias In, Bias Out: Why Legislation Placing Requirements on the Procurement of Commercialized Facial Recognition Technology Must Be Passed to Protect People of Color. *Public Contract Law Journal* 50, 1 (2020), 63–89.

[29] Claudia Flores-Saviaga, Shangbin Feng, and Saiph Savage. 2022. Datavoidant: An AI System for Addressing Political Data Voids on Social Media. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022), 1–29.

[30] Claudia Flores-Saviaga, Ricardo Granados, Liliana Savage, Lizbeth Escobedo, and Saiph Savage. 2020. Understanding the complementary nature of paid and volunteer crowds for content creation. *Avances en Interacción Humano-Computadora* 1 (2020), 37–44.

[31] Claudia Flores-Saviaga, Yuwen Li, Benjamin Hanrahan, Jeffrey Bigham, and Saiph Savage. 2020. The challenges of crowd workers in rural and urban America. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 8. 159–162.

[32] Claudia Flores-Saviaga and Saiph Savage. 2021. Fighting disaster misinformation in Latin America: the# 19S Mexican earthquake case study. *Personal and Ubiquitous Computing* 25 (2021), 353–373.

[33] Claudia Flores-Saviaga, Saiph Savage, and Dario Taraborelli. 2016. Leadwise: using online bots to recruit and guide expert volunteers. In *Proceedings of the 19th ACM Conference on Computer Supported Cooperative Work and Social Computing Companion.* 257–260.

[34] Md Tahmid Hasan Fuad, Awal Ahmed Fime, Delowar Sikder, Md Akil Raihan Iftee, Jakaria Rabbi, Mabrook S Al-Rakhami, Abdu Gumae, Ovishake Sen, Mohtasim Fuad, and Md Nazrul Islam. 2021. Recent advances in deep learning techniques for face recognition. *IEEE Access* (2021).

[35] Mitchell L Gordon, Michelle S Lam, Joon Sung Park, Kayur Patel, Jeff Hancock, Tatsunori Hashimoto, and Michael S Bernstein. 2022. Jury learning: Integrating dissenting voices into machine learning models. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems.*

[36] Stuart Hall. 1997. Representation, meaning and language. *Representation: Cultural representations and signifying practices* (1997), 15–64.

[37] Foad Hamidi, Morgan Klaus Scheuerman, and Stacy M Branham. 2018. Gender recognition or gender reductionism? The social implications of embedded gender recognition systems. In *Proceedings of the 2018 chi conference on human factors in computing systems.* 1–13.

[38] Kotaro Hara, Abigail Adams, Kristy Milland, Saiph Savage, Chris Callison-Burch, and Jeffrey P Bigham. 2018. A data-driven analysis of workers' earnings on Amazon Mechanical Turk. In *Proceedings of the 2018 CHI conference on human factors in computing systems.* 1–14.

[39] John Hartley. 2012. *Communication, cultural and media studies: The key concepts.* Routledge.

[40] Thomas Huang, Ziyou Xiong, and Zhenqiu Zhang. 2005. Face recognition applications. In *Handbook of Face Recognition.* Springer, 371–390.

[41] Yu-Hsuan Huang and Homer H Chen. 2020. Face recognition under low illumination via deep feature reconstruction network. In *2020 IEEE International Conference on Image Processing (ICIP).* IEEE, 2161–2165.

[42] Andreas Ihl, Kim Simon Strunk, and Marina Fiedler. 2018. The influence of utilitarian and hedonic motivation on success in crowd work. (2018).

[43] Niharika Jain, Alberto Olmo, Sailik Sengupta, Lydia Manikonda, and Subbarao Kambhampati. 2022. Imperfect ImaGANation: Implications of GANs exacerbating biases on facial data augmentation and snapchat face lenses. *Artificial Intelligence* 304 (2022), 103652.

[44] Byoungjun Jeon, Boseong Jeong, Seunghoon Jee, Yan Huang, Youngmin Kim, Gee Ho Park, Jungah Kim, Maierdanjiang Wufuer, Xian Jin, Sang Wha Kim et al. 2019. A facial recognition mobile app for patient safety and biometric identification: Design, development, and validation. *JMIR mHealth and uHealth* 7, 4 (2019), e11472.

[45] Shivani Kapania, Ding Wang, and Alex Taylor. 2023. A hunt for the Snark: Annotator Diversity in Data Practices. *ACM Conference on Human Factors in Computing Systems (CHI)* (2023).

[46] Vijay Keswani, Matthew Lease, and Krishnaram Kenthapadi. 2022. Designing closed human-in-the-loop deferral pipelines. *arXiv preprint arXiv:2202.04718* (2022).

[47] Lucy Kimbell. 2011. Rethinking design thinking: Part I. *Design and culture* 3, 3 (2011), 285–306.

[48] Alfred Kobsa and Jörg Schreck. 2003. Privacy through pseudonymity in user-adaptive systems. *ACM Transactions on Internet Technology (TOIT)* 3, 2 (2003), 149–183.

[49] Joshua A Kroll. 2022. ACM TechBrief: Facial Recognition.

[50] Sophie Lebrecht, Lara J Pierce, Michael J Tarr, and James W Tanaka. 2009. Perceptual other-race training reduces implicit racial bias. *PloS one* 4, 1 (2009), e4215.

[51] Christopher Libby and Jesse Ehrenfeld. 2021. Facial recognition technology in 2021: masks, bias, and the future of healthcare. *Journal of Medical Systems* 45, 4 (2021), 1–3.

[52] Jennifer Mankoff, Gillian R Hayes, and Devva Kasnitz. 2010. Disability studies as a source of critical inquiry for the field of assistive technology. In *Proceedings of the 12th international ACM SIGACCESS conference on Computers and accessibility*. 3–10.

[53] Liam M Mayron, Yasser Hausawi, and Gisela Susanne Bahr. 2013. Secure, usable biometric authentication systems. In *International Conference on Universal Access in Human-Computer Interaction*. Springer, 195–204.

[54] Michael McLaughlin and Daniel Castro. 2020. The Critics Were Wrong: NIST Data Shows the Best Facial Recognition Algorithms Are Neither Racist Nor Sexist. https://itif.org/publications/2020/01/27/critics-were-wrong-nist-data-shows-best-facial-recognition-algorithms/.

[55] Ahmed M Megreya and Markus Bindemann. 2013. Individual differences in personality and face identification. *Journal of Cognitive Psychology* 25, 1 (2013), 30–37.

[56] Ahmed M Megreya and Markus Bindemann. 2015. Developmental improvement and age-related decline in unfamiliar face matching. *Perception* 44, 1 (2015), 5–22.

[57] Ahmed M Megreya, Markus Bindemann, and Catriona Havard. 2011. Sex differences in unfamiliar face identification: Evidence from matching tasks. *Acta Psychologica* 137, 1 (2011), 83–89.

[58] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2021. A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)* 54, 6 (2021), 1–35.

[59] Jiani Hu Xunqiang Tao Yaohai Huang Mei Wang, Weihong Deng. [n. d.]. RFW Dataset. http://whdeng.cn/RFW/testing.html.

[60] Christopher S Milligan. 1999. Facial recognition technology, video surveillance, and privacy. *S. Cal. Interdisc. LJ* 9 (1999), 295.

[61] Vikram Mohanty, Kareem Abdol-Hamid, Courtney Ebersohl, and Kurt Luther. 2019. Second opinion: Supporting last-mile person identification with crowdsourcing and face recognition. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 7. 86–96.

[62] Vikram Mohanty, David Thames, Sneha Mehta, and Kurt Luther. 2019. Photo sleuth: Combining human expertise and face recognition to identify historical portraits. In *Proceedings of the 24th International Conference on Intelligent User Interfaces*. 547–557.

[63] Vikram Mohanty, David Thames, Sneha Mehta, and Kurt Luther. 2020. Photo sleuth: Identifying historical portraits with face recognition and crowdsourced human expertise. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 10, 4 (2020), 1–36.

[64] Ihudiya Finda Ogbonnaya-Ogburu, Angela DR Smith, Alexandra To, and Kentaro Toyama. 2020. Critical race theory for HCI. In *Proceedings of the 2020 CHI conference on human factors in computing systems*. 1–16.

[65] Alice J O'Toole, Xaiobo An, Joseph Dunlop, Vaidehi Natu, and P Jonathon Phillips. 2012. Comparing face recognition algorithms to humans on challenging tasks. *ACM Transactions on Applied Perception (TAP)* 9, 4 (2012), 1–13.

[66] Alice J O'Toole, P Jonathon Phillips, Fang Jiang, Janet Ayyad, Nils Penard, and Herve Abdi. 2007. Face recognition algorithms surpass humans matching faces over changes in illumination. *IEEE transactions on pattern analysis and machine intelligence* 29, 9 (2007), 1642–1646.

[67] A O'Toole, P Jonathon Phillips, Fang Jiang, Janet Ayyad, Nils Penard, and Hervé Abdi. 2005. Face recognition algorithms surpass humans. *IEEE Trans. PAMI* (2005).

[68] Robert Ezra Park. 1950. Race and culture. (1950).

[69] Sidney Perkowitz. 2021. The bias in the machine: Facial recognition technology and racial disparities. *MIT Case Studies in Social and Ethical Responsibilities of Computing https://doi. org/10.21428/2c646de5* 62272586, 5 (2021), 15.

[70] P Jonathon Phillips, Fang Jiang, Abhijit Narvekar, Julianne Ayyad, and Alice J O'Toole. 2011. An other-race effect for face recognition algorithms. *ACM Transactions on Applied Perception (TAP)* 8, 2 (2011), 1–11.

[71] P Jonathon Phillips, Amy N Yates, Ying Hu, Carina A Hahn, Eilidh Noyes, Kelsey Jackson, Jacqueline G Cavazos, Géraldine Jeckeln, Rajeev Ranjan, Swami Sankaranarayanan, et al. 2018. Face recognition accuracy of forensic examiners, superrecognizers, and face recognition algorithms. *Proceedings of the National Academy of Sciences* 115, 24 (2018), 6171–6176.

[72] Julian Alberto Posada Gutierrez. 2022. *The Coloniality of Data Work: Power and Inequality in Outsourced Data Production for Machine Learning*. Ph. D. Dissertation.

[73] Inioluwa Deborah Raji and Joy Buolamwini. 2019. Actionable auditing: Investigating the impact of publicly naming biased performance results of commercial ai products. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*. 429–435.

[74] Inioluwa Deborah Raji, Timnit Gebru, Margaret Mitchell, Joy Buolamwini, Joonseok Lee, and Emily Denton. 2020. Saving face: Investigating the ethical concerns of facial recognition auditing. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*. 145–151.

[75] Elissa M Redmiles. 2021. The need for respectful technologies: Going beyond privacy. , 309–313 pages.

[76] Gillian Rhodes, Vance Locke, Louise Ewing, and Emma Evangelista. [n. d.]. Race coding and the other-race effect in face recognition. *Perception* 38, 2 ([n. d.]), 232–241.

[77] Jakob Rogstadius, Vassilis Kostakos, Aniket Kittur, Boris Smus, Jim Laredo, and Maja Vukovic. 2011. An assessment of intrinsic and extrinsic motivation on task performance in crowdsourcing markets. In *Proceedings of the international AAAI conference on web and social media*, Vol. 5.

[78] Drew Roselli, Jeanna Matthews, and Nisha Talagala. 2019. Managing bias in AI. In *Companion Proceedings of The 2019 World Wide Web Conference*.

[79] Sagar Samtani, Murat Kantarcioglu, and Hsinchun Chen. 2020. Trailblazing the artificial intelligence for cybersecurity discipline: a multi-disciplinary research roadmap. , 19 pages.

[80] Saiph Savage, Claudia Flores-Saviaga, Rachel Rodney, Liliana Savage, Jon Schull, and Jennifer Mankoff. 2022. The Global Care Ecosystems of 3D Printed Assistive Devices. *Accessible Computing* 15, 4 (2022).

[81] Jay Scherer and Steven J Jackson. 2008. Cultural studies and the circuit of culture: Advertising, promotional culture and the New Zealand All Blacks. *Cultural Studies? Critical Methodologies* 8, 4 (2008), 507–526.

[82] Morgan Klaus Scheuerman, Alex Hanna, and Emily Denton. 2021. Do datasets have politics? Disciplinary values in computer vision dataset development. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–37.

[83] 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 (2020), 1–35.

[84] Abdallah Hussein Sham, Kadir Aktas, Davit Rizhinashvili, Danila Kuklianov, Fatih Alisinanoglu, Ikechukwu Ofodile, Cagri Ozcinar, and Gholamreza Anbarjafari. 2023. Ethical AI in facial expression analysis: Racial bias. *Signal, Image and Video Processing* 17, 2 (2023), 399–406.

[85] Tomáš Sixta, Julio Jacques Junior, Pau Buch-Cardona, Eduard Vazquez, and Sergio Escalera. 2020. Fairface challenge at eccv 2020: Analyzing bias in face recognition. In *European conference on computer vision*. Springer, 463–481.

[86] Audrey Smedley and Brian D Smedley. 2005. Race as biology is fiction, racism as a social problem is real: Anthropological and historical perspectives on the social construction of race. *American psychologist* 60, 1 (2005), 16.

[87] Angela DR Smith, Adriana Alvarado Garcia, Ian Arawjo, Audrey Bennett, Khalia Braswell, Bryan Dosono, Ron Eglash, Denae Ford, Daniel Gardner, Shamika Goddard, et al. 2021. Keepin'it real about race in HCI. *Interactions* 28, 5 (2021), 28–33.

[88] Siegfried Ludwig Sporer. 2001. Recognizing faces of other ethnic groups: An integration of theories. *Psychology, Public Policy, and Law* (2001).

[89] James W Tanaka and Lara J Pierce. 2009. The neural plasticity of other-race face recognition. *Cognitive, Affective, & Behavioral Neuroscience* 9, 1 (2009), 122–131.

[90] Jacob Thebault-Spieker, Sukrit Venkatagiri, David Mitchell, Chris Hurt, and Kurt Luther. 2019. Pairwise: mitigating political bias in crowdsourced content moderation. *Proceedings of HCOMP* (2019).

[91] Bridget Tombleson and Katharina Wolf. 2017. Rethinking the circuit of culture: How participatory culture has transformed cross-cultural communication. *Public Relations Review* 43, 1 (2017), 14–25.

[92] Mei Wang and Weihong Deng. 1804. Deep face recognition: a survey (2018). *arXiv preprint arXiv:1804.06655* (1804).

[93] Mei Wang and Weihong Deng. 2020. Mitigating bias in face recognition using skewness-aware reinforcement learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 9322–9331.

[94] Mei Wang and Weihong Deng. 2021. Deep face recognition: A survey. *Neurocomputing* 429 (2021), 215–244.

[95] Mei Wang, Weihong Deng, Jiani Hu, Xunqiang Tao, and Yaohai Huang. 2019. Racial faces in the wild: Reducing racial bias by information maximization adaptation network. In *Proceedings of the ieee/cvf international conference on computer vision*. 692–702.

[96] David White, James D Dunn, Alexandra C Schmid, and Richard I Kemp. 2015. Error rates in users of automatic face recognition software. *PloS one* 10, 10 (2015), e0139827.

[97] Hoo Keat Wong, Ian D Stephen, and David RT Keeble. 2020. The own-race bias for face recognition in a multiracial society. *Frontiers in psychology* 11 (2020), 208.

[98] Benjamin P Wood, Man Cao, Michael D Bond, and Dan Grossman. 2017. Instrumentation bias for dynamic data race detection. *Proceedings of the ACM on Programming Languages* 1, OOPSLA (2017), 1–31.

[99] Esther Xiu Wen Wu, Bruno Laeng, and Svein Magnussen. 2012. Through the eyes of the own-race bias: Eye-tracking and pupillometry during face recognition. *Social neuroscience* 7, 2 (2012), 202–216.

[100] Xingjiao Wu, Luwei Xiao, Yixuan Sun, Junhang Zhang, Tianlong Ma, and Liang He. 2022. A survey of human-in-the-loop for machine learning. *Future Generation Computer Systems* (2022).

[101] Jessica L Yaros, Diana A Salama, Derek Delisle, Myra S Larson, Blake A Miranda, and Michael A Yassa. 2019. A memory computational basis for the other-race effect. *Scientific reports* 9, 1 (2019), 1–11.

[102] Lixiu Yu, Paul André, Aniket Kittur, and Robert Kraut. 2014. A comparison of social, learning, and financial strategies on crowd engagement and output quality. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*. 967–978.

[103] Stefanos Zafeiriou, Cha Zhang, and Zhengyou Zhang. 2015. A survey on face detection in the wild: past, present and future. *Computer Vision and Image Understanding* 138 (2015), 1–24.