

WEIRD FAccTs: How Western, Educated, Industrialized, Rich, and Democratic is FAccT?

Original

WEIRD FAccTs: How Western, Educated, Industrialized, Rich, and Democratic is FAccT? / Septiandri, Ali Akbar; Constantinides, Marios; Tahaei, Mohammad; Quercia, Daniele. - (2023), pp. 160-171. (Intervento presentato al convegno 6th ACM Conference on Fairness, Accountability, and Transparency, FAccT 2023 tenutosi a Chicago IL (USA) nel June 12 - 15, 2023) [10.1145/3593013.3593985].

Availability:

This version is available at: 11583/2996109 since: 2025-01-02T14:12:07Z

Publisher:

Association for Computing Machinery

Published

DOI:10.1145/3593013.3593985

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

WEIRD FAccTs: How Western, Educated, Industrialized, Rich, and Democratic is FAccT?

Ali Akbar Septiandri, Marios Constantinides, Mohammad Tahaei, Daniele Quercia
 {ali.septiandri,marios.constantinides,mohammad.tahaei,daniele.quercia}@nokia-bell-labs.com

Nokia Bell Labs
 Cambridge, UK

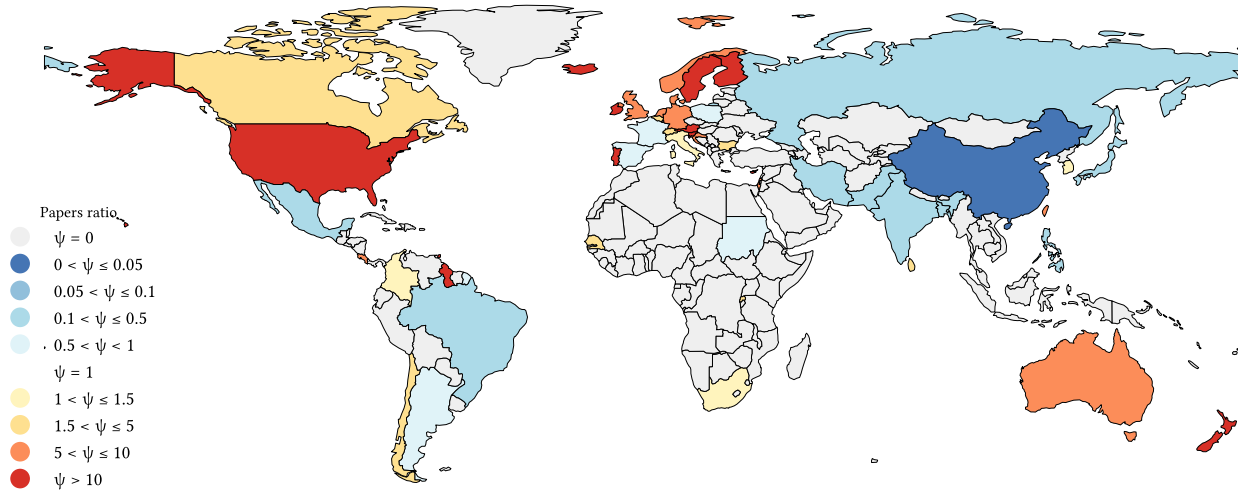


Figure 1: Worldwide distribution of FAccT papers ratio (ψ_{sc}) between 2018–2022, showing over-represented ($\psi_{sc} > 1$) and under-represented countries ($\psi_{sc} < 1$). The papers ratio ψ_{sc} is the number of papers from a given country c (where the study participants are from) over the number of all papers in FAccT, divided by c 's population over the world's population. Countries in light gray ($\psi_{sc} = 0$) did not have participants in FAccT between 2018–2022, while those in darker shades of blue and red indicate under- and over-represented countries, respectively.

ABSTRACT

Studies conducted on Western, Educated, Industrialized, Rich, and Democratic (WEIRD) samples are considered atypical of the world's population and may not accurately represent human behavior. In this study, we aim to quantify the extent to which the ACM FAccT conference, the leading venue in exploring Artificial Intelligence (AI) systems' fairness, accountability, and transparency, relies on WEIRD samples. We collected and analyzed 128 papers published between 2018 and 2022, accounting for 30.8% of the overall proceedings published at FAccT in those years (excluding abstracts, tutorials, and papers without human-subject studies or clear country attribution for the participants). We found that 84% of the analyzed papers

were exclusively based on participants from Western countries, particularly exclusively from the U.S. (63%). Only researchers who undertook the effort to collect data about local participants through interviews or surveys added diversity to an otherwise U.S.-centric view of science. Therefore, we suggest that researchers collect data from under-represented populations to obtain an inclusive world-view. To achieve this goal, scientific communities should champion data collection from such populations and enforce transparent reporting of data biases.

CCS CONCEPTS

• Social and professional topics; • Computing methodologies; • Software and its engineering; • Human-centered computing → Human computer interaction (HCI);

ACM Reference Format:

Ali Akbar Septiandri, Marios Constantinides, Mohammad Tahaei, Daniele Quercia. 2023. WEIRD FAccTs: How Western, Educated, Industrialized, Rich, and Democratic is FAccT?. In *2023 ACM Conference on Fairness, Accountability, and Transparency (FAccT '23)*, June 12–15, 2023, Chicago, IL, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3593013.3593985>

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

FAccT '23, June 12–15, 2023, Chicago, IL, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0192-4/23/06...\$15.00

<https://doi.org/10.1145/3593013.3593985>

1 INTRODUCTION

The ACM Conference on Fairness, Accountability, and Transparency (FAccT)¹ is an annual conference that aims to bring together a diverse community of scholars interested in exploring the fairness, accountability, and transparency of Artificial Intelligence (AI) and Machine Learning (ML) systems. One of the FAccT's challenges concerns biases, particularly as AI is increasingly used in decision-making contexts [26]. One way biases can manifest is through the overuse of specific datasets when evaluating ML models, leading to unfair or inaccurate results. For example, ImageNet, a popular dataset for computer vision tasks, was found to be biased [42, 54].

In recent years, there has also been increasing attention on the concept of **Western, Educated, Industrialized, Rich, and Democratic (WEIRD)** research, which refers to studies conducted on participants drawn from WEIRD populations [19]. These samples are often considered atypical of the global population and may not represent human behavior more globally. For example, an analysis of the Conference on Human Factors in Computing Systems (CHI), the prominent conference for empirical studies focused on human-centered computing, showed that 73% of CHI papers come from Western samples, representing less than 12% of the world's population [30]. Such a focus on WEIRD samples results in the design of technologies used by the rest of the population. However, it does not consider those populations' cultural norms, characteristics, and expectations [33, 37].

Given the potential biases introduced by WEIRD samples, it is therefore vital to consider the extent to which a conference's published papers rely on such samples. In this study, we are interested in quantifying the *WEIRD-ness* of the FAccT conference by examining the datasets used in its proceedings, following Linxen et al. [30]'s method for analyzing CHI's WEIRD-ness. FAccT is a noteworthy conference to study WEIRD-ness as one of its main objectives is to bring fairness, accountability, and transparency to AI. In so doing, we made three main contributions:

- We collected the FAccT proceedings between 2018 and 2022, resulting in an initial set of 416 papers. Out of these papers, we analyzed a total of **128 papers** that met the inclusion criteria for our analysis (§3).²
- We found that **84%** of the analyzed papers were exclusively based on participants from Western countries (§4).
- Compared to CHI, FAccT has a more significant proportion of Western study participants (§4). By analyzing the intersection of the study participants from the two conferences, we found that the difference in "EIRD" (i.e., Educated, Industrialized, Rich, and Democratic) between the two conferences is still significant. To explain these differences, we looked at the sample size, types of datasets used (off-the-shelf vs. author-collected), and authors' affiliation country. However, we did not find evidence to attribute these differences to any of these factors.

In light of these results, we discuss the theoretical and practical implications of our findings (§5), which speak to the need to take steps such as diversifying authors, encouraging collaborations

through workshops, and using cross-cultural online research platforms to make FAccT, and more broadly computing conferences, less WEIRD.

2 RELATED WORK

In 2010, Henrich et al. did a comprehensive study about the generalizability of research findings. They state that most research findings are based on a small world population, often called the WEIRD population. In psychology, 96% of research samples come from the WEIRD population, which only accounts for 12% of the world population. One reason for such bias toward WEIRD samples is that many authors (73%) come from American universities. The situation exacerbates when authors and papers often assume and claim that their findings are universally valid and generalizable, asserting applicability to "humans." Nonetheless, the study population sometimes barely goes beyond undergraduate students from the authors' home institute [19]. Additionally, studies have demonstrated an unequal distribution of participation in the AI ethics debate [25]. There is also a lack of awareness regarding the potential consequences of AI adoption across nations, such as the adoption of face recognition technology in public administration in India despite known biases associated with women and minorities [40].

Since then, in the past decade, there has been an emerging interest in the Computer Science community about the potential biases introduced by WEIRD samples. A few examples of such research that informs the biases introduced by WEIRD samples include the consideration of: cultural differences when designing products that are used in many countries across the globe [6, 31, 33, 51], minorities in the design of computing education materials [35], and countries in the Global South for the design of digital accessibility [34]. To shed light on these biases, we provide details for two studies: (1) In the case of personal trackers, the goals of different user types can vary. Arab users generally see a fitness tracker more as a monitoring device. In contrast, WEIRD populations have been found to see the same device as an assistance to achieve their goal. In these two scenarios, a monitoring device would give users advice and recommendations, while a tool with more authority acts as a coach and gives rules and guidelines to the user [33]; (2) Smartphone privacy is influenced by the users' cultural norms. The current design may assume that device sharing does not happen, and privacy is preserved as long as a phone lock feature is provided. However, in some countries (e.g., India, Pakistan, and Bangladesh), users, particularly female users, may not be able to exercise their autonomy. It may be considered socially unacceptable not to let other people use their phones. In such scenarios, giving deletion options and private mode browsing could be options to help users maintain their privacy and adhere to their social norms (by letting other people use their phone) [37].

In particular, the two computing conferences that heavily focus on sampling issues are *CHI*, which focuses on human-centered design, and *FAccT*, which aims to democratize AI and advance the development of responsible AI. There has been a recent meta-study on CHI findings from 2016–2020 about the WEIRD-ness of the conference (§2), but not on FAccT. Laufer et al. conducted a meta (reflexive) study on four years of FAccT proceedings to extract research

¹Established in 2018, and initially called FAT.

²We made our dataset publicly available for replication and reproducibility.

topics (e.g., group-level fairness and disinformation) and understand community’s values (e.g., transparency over the peer review process and industry influence over the published research) [29]. Our work contributes to this line of research by exploring FAccT’s WEIRD-ness, and comparing it with another prominent research conference that focuses on empirical research and human-centered computing (CHI).

The Original Study: How WEIRD Is CHI? Linxen et al. [30] collected 3,269 CHI papers published in 2016–2020. Of these papers, 2,768 were on human subjects, and 1,076 papers (38.9%) had information about the participants’ countries, which is the basis of the study. CHI findings span 93 countries, with 7.1% of the samples being from Western countries. The USA, Ireland, and Switzerland were the top three represented countries regarding the number of participants. The correlation between samples and other WEIRD factors was also positive; the higher these factors, the stronger those populations were presented in the dataset. Similar to what Henrich et al. suggested (diversifying authors could be a way to reduce WEIRD research), findings from CHI show that 81.2% of papers studied participants from the authors’ institute’s country. The authors propose ideas such as diversifying authorship, increasing the use of online crowd-sourcing platforms, and appreciating replication studies to make CHI less WEIRD.

3 METHODOLOGY

Based on our related work (§2), we set out to explore the WEIRD-ness of FAccT studies. In so doing, we formulated two Research Questions (RQs): **RQ₁**: How WEIRD is the FAccT conference? **RQ₂**: How FAccT and CHI conferences compare in terms of WEIRD-ness? To answer these RQs, we replicated the Linxen et al. [30]’s methodology. Linxen et al.’s research was the first to operationalize WEIRD-ness in a conference proceeding, and by comparing FAccT with CHI, we demonstrate similarities and differences between the two communities and surface potential areas of oversight. We collected and analyzed 416 papers published at FAccT conferences between 2018–2022 from the ACM Digital Library.

3.1 Positionality Statement

Understanding researcher positionality is essential to demystifying our lens on data collection and analysis [14, 18]. We situate this paper in the United Kingdom in the 21st century, writing as authors who primarily work as academic and industry researchers. We identify as males from Indonesia, Cyprus, Iran, and Italy. Our shared backgrounds include HCI, privacy, security, software engineering, AI, social computing, and urbanism.

3.2 Dataset

The first author manually reviewed a random set of 103 (24.8%) papers relevant to the RQs. From this review, a coding guideline and exclusion criteria were collaboratively developed with three other authors to extract relevant data from the papers. The coding scheme and the exclusion criteria were then reviewed and discussed by all authors to ensure their completeness and accuracy. To augment the information obtained from the papers, we emailed 42 dataset owners requesting any missing information. Out of these, 19 (45.2%) replied, but we still could not get the country attribution for 9

(47.4%) datasets. The collected data were then analyzed using the agreed-upon coding scheme. Overall, this multi-step process aimed at systematically and rigorously preparing the dataset and reducing potential biases. To factor in the true magnitude of participants in the statistical analyses, we multiplied the number of datasets that were used more than once by the number of papers using those datasets. Consider, for example, the German credit dataset with 1,000 participants [20]. This dataset was used in 13 papers, yielding a total of 13,000 participants. This approach emphasized the problem in frequently used datasets, and aimed at motivating the community to consider under-represented countries in future sampling. We opted-in for multiplication because otherwise overused datasets would not show their magnitude in the analysis. For example, the German credit dataset would have been counted only once despite being used 13 times.

3.3 Inclusion & Exclusion Criteria

We manually analyzed 416 papers published at FAccT conferences between 2018 and 2022 and found that 93 (22.36%) of the papers used off-the-shelf datasets instead of manually collecting them through interviews, surveys, or workshops. To differentiate the two types of datasets, we will refer to these manually collected datasets from this point onward as “author-collected datasets.” For our analysis, we define our exclusion criteria as follows:

- *Abstracts and tutorials*: No study or insufficient information about the described studies;
- *Non-human datasets*, e.g., MNIST (handwritten digit images), CIFAR-10 (tiny images of 10 objects), Stanford Dogs (images of 120 breeds of dogs), ImageNet (14 million images of more than 20,000 categories), Singapore Bus Stops [43], Boston Housing [17]: because mixing human and non-human data would make it difficult to compare the results, as the units of measurement would be different;
- *Country attribution is unclear*, e.g., OpenImages MIAP, Last.FM, Amazon reviews: because it was not clear which country the human subjects were from or because the authors were studying issues related to a specific country but not analyzing data from human subjects.

3.4 Defining WEIRD Scores

The goal of this work is to quantify the WEIRD-ness of FAccT. To achieve that, we defined a set of WEIRD scores relying on two strands of research: *a*) work that focused on certain populations of the CHI conference [30]; and, *b*) work that focused on specific products for particular populations without much consideration for a majority of the world’s population [33, 34]. Next, we explain each of the WEIRD scores (Table 1):

Western. To determine whether a country is classified as Western, we used the Huntington classification [22]. This is based on Huntington’s thesis, *the Clash of Civilizations*, where he posits that cultural differences are the main cause of conflicts among humanity. Therefore, the classification is based on cultural and historical factors, including the country’s language and religion. According to Huntington, Western civilization is rooted in the tradition of

ancient Greece and Rome, and is characterized by certain commonalities such as the predominance of Christianity, the use of the Latin alphabet, and the prevalence of democratic political systems. Of particular interest in his thesis is the concept of “torn countries”; countries with no clear-cut classification. For example, Turkey has adopted Western customs, including how its people dress, started using the Latin alphabet, became a member of NATO, and has been trying to become a member of the European Union. However, under Huntington’s classification, Turkey is still classified as non-Western due to its history, culture, and traditions derived from Islamic civilization. All European Union members are classified as Western countries [12]. It is worth mentioning that certain nations, such as Japan, South Korea, Chile, and Argentina, meet the criteria of being Educated, Industrialized, Rich, and Democratic, but do not fall under the category of Western nations.

Educated. We used the mean years of schooling per person as reported in the UNDP Human Development Report [36] to measure the educational level of a country. This measure considers the mean years of schooling for adults aged 25 years. An alternative proxy would be the PISA index from OECD. However, the UNDP index was favored to allow for reproducible results (similar to CHI’s findings [30]).

Industrialized. To determine the industrialization of a country, we used the gross domestic product (GDP) per capita adjusted by purchasing power parity (PPP) [44]. GDP measures the total value of goods and services produced within a country’s borders in a given period, usually a year. It is an indicator of a country’s economic performance and is used to compare the economic output of different countries. The data is reported in current international dollars. The international dollar is an artificial currency used to adjust for differences in purchasing power when comparing economic performance among countries. An alternative to this proxy is to use the Competitive Industrial Performance (CIP) Index by United Nations Industrial Development Organization (UNIDO) [46]. However, to produce comparable results to CHI’s WEIRD-ness, the choice of GDP per capita adjusted by PPP was favored.

Rich. We used the gross national income (GNI) per capita adjusted by purchasing power parity (PPP) [45] to measure the wealth of a country as suggested by Arnett [3] and Linxen et al. [30]. GNI per capita approximates a population’s standard of living as it is calculated by adding up all the income earned by a country’s residents and businesses. The data is reported in current international dollars, allowing for comparison between countries.

Democratic. To determine the democracy of a country, we used the “political rights” scores provided by the Freedom House [13]. Freedom House is a U.S. non-profit organization whose mission is to research democracy, freedom, and human rights. The “political rights” measure considers the level of political freedom and rights enjoyed by the citizens of a country. Alternatively, one could use the Democracy Index [11] as a proxy for political rights, developed by the research division of the Economist Group, the Economist Intelligence Unit (EIU). To allow for reproducible results, we opted in for political rights provided by the Freedom House.

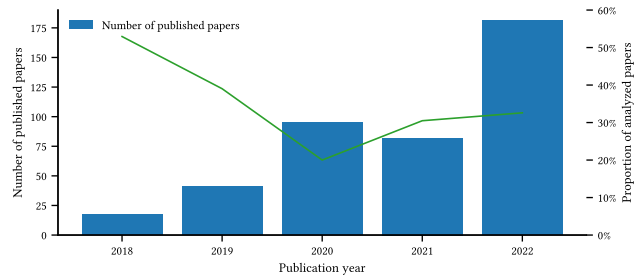


Figure 2: Number of published papers and proportions (green line) of those being analyzed. The proportion of the analyzed papers varied between years due to the number of published papers and the type of those papers. For example, in 2020, 36% of the published papers were abstracts or tutorials. Tutorial papers were not archived in publication years other than 2020.

3.5 Analysis: Terminology and Computing WEIRD

Throughout the paper, we refer to Table 1 for the quantification of the WEIRD scores. The formulae are based on and formalized from methods in [30]. Kendall rank correlation is often used when the data is not normally distributed or when the variables being compared are ordinal rather than interval or ratio variables [2]. The coefficient ranges from -1 to 1, where a value of 1 indicates a perfect positive correlation, a value of -1 indicates a perfect negative correlation, and a value of 0 indicates no correlation. If the coefficient is close to 1, then the researchers at FAcCT conferences fully conducted studies with WEIRD subjects. To get the Kendall rank correlation coefficient closer to zero, researchers should aim to make their study participants proportionate to the country’s population [2].

4 RESULTS

We reviewed 416 papers and discovered that 226 (54.3%) included either synthetic, off-the-shelf, or author-collected datasets. We further narrowed down the selection by excluding papers that were based on non-human participants, had an unspecified number of participants, or did not specify the country of origin, resulting in 128 papers (30.8%) for analysis (§3.3 for details of our (ex)inclusion criteria).

From 2018 to 2022, we observed an increase in the number of published papers in FAcCT (Figure 2). In recent years, there have also been more papers with human participants (32.6% in 2022). We saw this increase after the drop in 2020. The proportion of analyzed papers in 2020 was the lowest because 36% of the published papers that year were either abstracts or tutorials; it was the only year when the tutorials were archived.

Table 2 shows the top 10 countries by the participants ratio (ψ_{pc}), papers ratio (ψ_{sc}), and number of papers (n_{sc}). We found that the U.S. participants are over-represented in the FAcCT community (93.6% of the participants). This is because some datasets (e.g., Adult [10], COMPAS recidivism [24], and ACS [9, 48]) are used multiple times and are based on national surveys. For similar reasons, Germany

Table 1: WEIRD variables and the formula to compute the values. $\mathbb{E}[\cdot]$ means the expected value of a random variable. In Kendall rank correlation (τ), P is the number of concordant pairs, Q is the number of discordant pairs, T is the number of ties in the first variable, and U is the number of ties in the second variable. Concordant pairs are pairs of observations in which the two variables are ranked in the same order, while discordant pairs are pairs of observations in which the two variables are ranked in opposite orders [2]. \vec{X} means a vector of value for variable X from all sampled countries. The formulae in this table are derived from and formalized based on the methods outlined in the paper by Linxen et al. [30].

Symbol	Variable	Formula	Description
c	Country	-	Country where the samples are from
p	Participants	-	Number of participants
s	Papers	-	Number of papers
π	Population	-	Population size of a country based on World Population Prospects 2022 [47]
W_c	Western	1 if $c \in$ Western else 0	Whether country c is Western based on Huntington classification [22]
E_c	Educated	$\mathbb{E}_c[\text{years of schooling}]$	Mean years of schooling for country c based on UNDP Human Development Report (2022) [36]
I_c	Industrialized	GDP per capita $_c$	Level of industrialization for country c based on World Bank GDP per capita, PPP (current Int\$, 2020) [44]
R_c	Rich	GNI per capita $_c$	Wealth of country c based on World Bank GNI per capita, PPP (current Int\$, 2020) [45]
D_c	Democratic	political rights $_c$	Level of democracy for country c based on Freedom House Political Rights (2022) [13]
ψ_{pc}	Participants ratio per country	$\frac{p_c / \sum_c p_c}{\pi_c / \sum_c \pi_c}$	Ratio of the proportion of participants for country c to the proportion of population size for country c
ψ_{sc}	Papers ratio per country	$\frac{s_c / \sum_c s_c}{\pi_c / \sum_c \pi_c}$	Ratio of the proportion of papers for country c to the proportion of population size for country c
$\tau(\cdot, \cdot)$	Kendall rank correlation	$\frac{P-Q}{\sqrt{(P+Q+T) \cdot (P+Q+U)}}$	The similarity of two rankings, e.g. $\vec{\psi}_s$ and \vec{E} ; $\vec{\psi}_p$ and \vec{D}
W -score	Western score	$\frac{1}{N} \sum_c W_c$	Expected value of how Western a conference is from all sampled countries
E -score	Educated score	$\tau(\vec{\psi}_s, \vec{E})$	How correlated papers ratio and mean years of schooling from all sampled countries
I -score	Industrialized score	$\tau(\vec{\psi}_s, \vec{I})$	How correlated papers ratio and level of industrialization from all sampled countries
R -score	Rich score	$\tau(\vec{\psi}_s, \vec{R})$	How correlated papers ratio and level of wealth from all sampled countries
D -score	Democratic score	$\tau(\vec{\psi}_s, \vec{D})$	How correlated papers ratio and level of democracy from all sampled countries

is the second highest country in terms of the number of papers because of the German credit dataset [10] that is used in 13 (10.2%) papers.

4.1 RQ1: How WEIRD is the FAccT Conference?

Table 3 shows that the FAccT community primarily selects participants from Western countries (84% exclusively from Western countries, 63% exclusively from the U.S.). We observed that the proportion of FAccT papers with exclusively Western participants is higher than CHI. The classification of participant countries in a paper is determined by whether the participants are from Western or non-Western countries as defined by Huntington’s framework [22].

Papers are considered: “exclusively Western,” if all participants are from Western countries; “exclusively non-Western,” if all participants are from non-Western countries; and, “mixed,” if participants are from both Western and non-Western countries. The data in the CHI column is not included in the paper by Linxen et al. [30] but for comparison, we reproduced the numbers using their data. In the paper by Linxen et al. [30], the authors reported the number of “participant samples” since a single paper can report participants from multiple countries. As the reported number can be higher than the actual number of papers, our method matches the total number of papers.

Table 2: Top 10 countries of FAccT participants between 2018–2022. ψ_{pc} is the participants ratio, ψ_{sc} is the papers ratio, n_{pc} is the number of participants, n_{sc} is the number of papers, $\%_{pc}$ is the proportion of participants, and $\%_{sc}$ is the proportion of papers. The USA has the most number of papers and participants. One explanation for this is the availability of off-the-shelf datasets such as Adult, COMPAS recidivism, and ACS. These datasets, particularly Adult and ACS, are based on national surveys and, as a result, have many participants. Since the papers ratio ψ_{sc} (middle-column) is calculated by the proportion of papers over the proportion of population size in country c , countries with smaller population size (e.g., Iceland, Guyana, Cyprus) are over-represented.

Top countries by ψ_{pc}				Top countries by ψ_{sc}				Top countries by n_{sc}			
Country	n_{pc}	$\%_{pc}$	ψ_{pc}	Country	n_{sc}	$\%_{sc}$	ψ_{sc}	Country	n_{sc}	$\%_{sc}$	ψ_{sc}
USA	16,621,168	93.62%	21.85	Iceland	1	0.44%	93.79	USA	106	46.49%	10.85
Colombia	472,000	2.66%	4.09	Guyana	1	0.44%	43.14	Germany	20	8.77%	8.25
Portugal	87,963	0.50%	3.77	Ireland	4	1.75%	27.81	United Kingdom	12	5.26%	6.15
Taiwan	120,000	0.68%	2.22	Cyprus	1	0.44%	27.79	India	8	3.51%	0.20
Costa Rica	22,000	0.12%	1.90	Portugal	7	3.07%	23.38	Portugal	7	3.07%	23.38
India	399,528	2.25%	0.13	Trinidad and Tobago	1	0.44%	22.65	Taiwan	4	1.75%	5.77
Germany	15,750	0.09%	0.08	New Zealand	3	1.32%	20.38	Canada	4	1.75%	3.63
Iceland	63	0.00%	0.08	Slovenia	1	0.44%	16.24	Sweden	4	1.75%	13.27
Slovenia	336	0.00%	0.07	Sweden	4	1.75%	13.27	Ireland	4	1.75%	27.81
Sweden	1,471	0.01%	0.06	Finland	2	0.88%	12.44	Australia	4	1.75%	5.36

Table 3: Distribution of papers by participant countries. Most FAccT & CHI papers use participants exclusively from Western countries.

Variable	FAccT		CHI (Linxen et al.)	
	n	%	n	%
Exclusively Western	108	84.38%	817	75.93%
Exclusively non-Western	9	7.03%	197	18.31%
Mixed	11	8.59%	62	5.76%
Total	128	100%	1076	100%

Table 4 shows the Kendall rank correlations (τ) between EIRD values and the ratio of the paper (ψ_{sc}). We provided the 95% Confidence Interval (CI) via bootstrapping.³ We also show the results from Linxen et al. [30] in the same table as a comparison point. As a sanity check, we reproduced their findings using our code and method⁴, and their dataset. According to Table 2, the percentage of participants from the U.S. in FAccT (93.6%) is higher than in CHI (54.8%). Nevertheless, Table 4 suggests that the level of “EIRD-ness” is lower for FAccT. This may come from the EIRD scores—calculated using Kendall tau *rank correlation*—which are not affected by outliers. For example, if half of the participants from the U.S. were removed, the proportion of the U.S. participants compared to all participants and the participant ratio would change. However, the U.S. would still rank first, and the ranks of other countries would remain unchanged. As a result, the EIRD scores would not change.

³Bootstrapping is a statistical method that involves resampling a dataset with replacement to quantify the uncertainty associated with a given estimator. This method is beneficial when the population distribution is unknown or when the sample size is small [23].

⁴Our code repository and dataset are publicly available at <https://github.com/aliakbars/weird-facct>.

Table 4: Kendall rank correlations (τ) of the ratio of the paper ψ_{sc} with measures of Educated, Industrialized, Rich, Democratic. The confidence intervals are generated from 10,000 bootstrap samples. Significance level: * $p < .05$, ** $p < .01$, *** $p < .001$.

Variable	FAccT		CHI (Linxen et al.)	
	τ	95% CI τ	τ	95% CI τ
Educated	0.31**	[0.12, 0.50]	0.46***	[0.34, 0.59]
Industrialized	0.35***	[0.21, 0.50]	0.50***	[0.40, 0.62]
Rich	0.35***	[0.20, 0.50]	0.50***	[0.39, 0.62]
Democratic	0.36***	[0.19, 0.53]	0.50***	[0.38, 0.62]

4.2 RQ2: How FAccT and CHI Conferences Compare in Terms of WEIRD-ness?

In what follows, we sought to find possible explanations for these differences. To explain them, we examined three potential factors for the variations between FAccT and CHI: 1) we aimed to understand whether the sample size of CHI could account for the differences; 2) we evaluated the connection between the datasets used and the EIRD scores, only testing this hypothesis on FAccT as the datasets were only coded for this conference; 3) we evaluated the relationship between cross-country collaboration and the EIRD scores.

Sample Size. Tables 3 and 4 show that FAccT has a larger proportion of papers with exclusively Western participants compared to CHI but is less “EIRD” (i.e., the Educated, Industrialized, Rich, and Democratic scores are lower). However, we expected a positive correlation between Western and EIRD variables. A possible explanation of this counter-intuitive result is that EIRD scores are measured using Kendall tau rank correlation, which is heavily based

Table 5: Kendall rank correlations (τ) of the ratio of the paper ψ_{sc} with measures of Educated, Industrialized, Rich, Democratic by using the same set of countries for both FAccT and CHI. The confidence intervals are generated from 10,000 bootstrap samples. Significance level: * $p < .05$, ** $p < .01$, * $p < .001$.**

Variable	FAccT		CHI (Linxen et al.)	
	τ	95% CI τ	τ	95% CI τ
Educated	0.40***	[0.18, 0.62]	0.54***	[0.37, 0.72]
Industrialized	0.44***	[0.27, 0.60]	0.65***	[0.55, 0.75]
Rich	0.44***	[0.27, 0.60]	0.65***	[0.55, 0.75]
Democratic	0.45***	[0.27, 0.63]	0.58***	[0.40, 0.76]

Table 6: We compared the effect of using off-the-shelf datasets (e.g., those from UCI Machine Learning Repository) and author-collected datasets (e.g., interviews and surveys) to the Educated, Industrialized, Rich, and Democratic scores by randomly shuffling the two types of datasets and calculating the difference in EIRD scores ($\Delta\tau$). We found that the actual differences fall within the 95% confidence interval from the shuffling, indicating that there is insufficient evidence to suggest that these differences are significant.

Variable	$\Delta\tau$	95% CI
Educated	-0.03	[-0.52, 0.15]
Industrialized	0.09	[-0.20, 0.26]
Rich	0.06	[-0.21, 0.28]
Democratic	0.02	[-0.51, 0.18]

on the set of sampled countries. However, we tested whether the difference between the two conferences would impact our results and that was not the case (Table 5). The two conferences have different sets of sampled countries: FAccT from 50 unique countries, whereas CHI from 93 unique countries. Therefore, to test whether EIRD scores are associated with the difference in sampled countries, we computed the Kendall tau rank correlation for each conference using the intersection of the two sets of countries (Table 5). We found that FAccT is still less “EIRD” compared to CHI. We further explored this difference by bootstrapping the number of papers from each conference using a power of 2 (i.e., 16, 32, ..., maximum number of papers for the respective conference). We sampled by replacement from the pool of papers 10,000 times for each sample size and computed the EIRD scores each time. This approach allowed us to quantify the uncertainty associated with the scores. Again, FAccT is less “EIRD” than CHI (Figure 3).

Off-the-Shelf Datasets vs. Author-Collected Datasets. Off-the-shelf datasets are publicly available (e.g., UCI Machine Learning Repository). In contrast, author-collected datasets are those collected through interviews, surveys, or workshops. We examined whether overused datasets may explain the bias in the sample population by comparing studies that used off-the-shelf and author-collected datasets. Over-reliance on certain datasets may lead to

Table 7: The number of unique author affiliation countries is not associated with any of the “EIRD” variables for FAccT, while there is an inverse correlation for CHI. ρ is the Pearson correlation coefficient.

Variable	FAccT		CHI (Linxen et al.)	
	ρ	95% CI	ρ	95% CI
Educated	-0.06	[-0.24, 0.10]	-0.22	[-0.30, -0.14]
Industrialized	-0.07	[-0.25, 0.10]	-0.21	[-0.29, -0.14]
Rich	-0.06	[-0.25, 0.10]	-0.23	[-0.30, -0.16]
Democratic	0.03	[-0.13, 0.20]	-0.10	[-0.17, -0.03]

bias by repeatedly using participants from the same country. However, there are also widely-used datasets that include samples from multiple countries such as the drug consumption [10] and the Vox-Celeb 1 [32]. Additionally, while the majority of participants in FAccT papers are from the U.S. due to the use of datasets such as the Adult, the COMPAS recidivism, and the ACS, these studies typically focus on specific subgroups within these datasets such as native Americans, Hispanic, and Black communities [7, 16, 41]. Other off-the-shelf datasets used in FAccT papers also include participants from multiple countries such as FairFace [27] in [53] and Labeled Faces in the Wild [21] in [5, 15]. However, since the countries of origin for the participants in these datasets are not specified, these papers were not included in the analysis.

From the 128 analyzed papers, we found that 90 (70.3%) of them were only using off-the-shelf datasets, 35 (27.3%) were using author-collected datasets, and 3 (2.4%) were using a mix of both. The top three off-the-shelf datasets were the Adult Income [10] (23 papers, 18%), the COMPAS Recidivism [24] (21 papers, 16.4%), and the German credit [20] (13 papers, 10.2%). Our initial assumption was that there would be no difference between the two types of datasets in terms of WEIRD-ness (excluding the 3 papers with mixed datasets). To test the EIRD part, we randomly shuffled the labels from the two sets, and found no statistically significant difference (Table 6). A chi-square test of independence was performed to examine the relation between the Western score and the type of datasets (i.e., off-the-shelf vs. author-collected), and found no statistically significant association between the two variables ($\chi^2(2, N = 125) = 4.32, p = 0.16$). However, Figure 4 shows that the author-collected datasets are slightly less Western. Separating papers based on the majority of participants’ country of origin also shows subtle differences. FAccT authors mainly used off-the-shelf datasets from the U.S., while datasets collected by authors (e.g., interviews or surveys) have a more localized representation (Figure 5).

Author Affiliation Countries. We examined whether the geographical differences in authors’ affiliations explain the sample representativeness of research studies. We wanted to understand whether authors from multiple countries are associated with the research paper being less WEIRD (inspired by discussions and proposals made by Linxen et al. [30]). The geographical location of an author’s affiliation can influence the population being studied. For example, suppose a study is conducted by researchers affiliated with institutions in the U.S. In that case, it is more likely that the

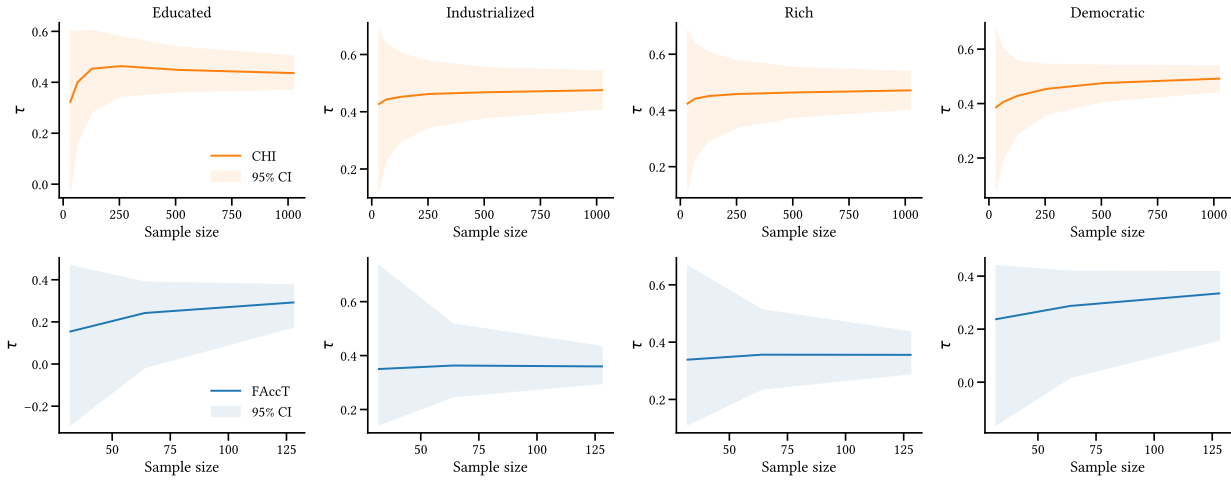


Figure 3: To see the effect of sample size, we did bootstrap sampling with different number of papers using a power of 2 (i.e., 16, 32, ..., maximum number of papers for the respective conference). The result suggests that FACCT is less “EIRD” than CHI.

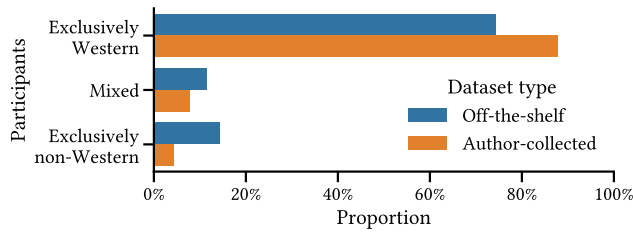


Figure 4: How types of datasets are associated with the proportion of sampled Western participants in FACCT.

Table 8: The Educated, Industrialized, Rich, and Democratic (EIRD) scores differences ($\Delta\tau$) between papers with authors from a single country and multiple countries. Positive values of $\Delta\tau$ indicate that papers with authors from a single country are sampling from more EIRD countries. By randomly shuffling the two types of papers and calculating the differences in scores, we found that the actual differences fall within the 95% confidence interval from the shuffling, indicating that these differences are not strong.

Variable	FACCT		CHI (Linxen et al.)	
	$\Delta\tau$	95% CI	$\Delta\tau$	95% CI
Educated	0.43	[-0.22, 0.59]	0.11	[-0.17, 0.12]
Industrialized	0.09	[-0.37, 0.23]	0.09	[-0.12, 0.18]
Rich	0.13	[-0.33, 0.25]	0.09	[-0.12, 0.18]
Democratic	0.28	[-0.17, 0.55]	0.02	[-0.07, 0.16]

sample will be composed of participants from the U.S. Conversely, if a study is conducted by researchers affiliated with institutions in Asia, it is more likely that the sample will be composed of participants from Asia. This is important to consider as it can lead to a lack of diversity in the sample and potential biases in the results.

Furthermore, the difference in countries can also influence the type of data collection methods used and the type of questions being asked, which can also impact the sample’s representativeness. Thus, we used Pearson correlation coefficients to measure the relationship between the number of countries that authors are affiliated with and the EIRD variables.

The first author manually extracted authors’ countries of affiliation from the 128 papers, using the affiliations reported in the papers. For authors with multiple affiliations, their first affiliation was used. Out of 519 authors, we were unable to determine the countries of 11 (2.2%), thus excluding them from this analysis. All analyzed papers ($n = 128$) have at least one author with their country listed, and the average number of authors is equal to 4.1 ($\sigma = 2.8$). We found that authors represent 20 unique countries (65% Western), and 75% of papers come from authors from a single country. Furthermore, we computed the Pearson correlation coefficients, one for each of the EIRD variables (Table 7). We found no significant association between the number of countries where the authors are affiliated with, and the EIRD variables.

To further explore the variations in EIRD scores, we divided the 128 papers into two mutually exclusive groups based on whether the authors were from one ($n = 96$) or multiple ($n = 32$) countries. We used a shuffling method to ensure the significance of the differences by re-arranging the paper labels 10,000 times, and calculating the differences in EIRD scores each time. Despite the noticeable differences in the Educated and Democratic scores, the results did not show statistical significance (Table 8). Given the number of analyzed papers that have been further divided into the two groups, these results may be due to a lack of statistical power (see §5.2). To partly address that, we conducted a similar analysis by using CHI data [30]. Despite coming from 70 unique countries, the authors of the 1076 papers tended to collaborate with authors in the same country; only 23% of papers were written by authors from at least two countries.

The Pearson correlation coefficients indicate that there is a negative correlation between the number of unique author affiliation

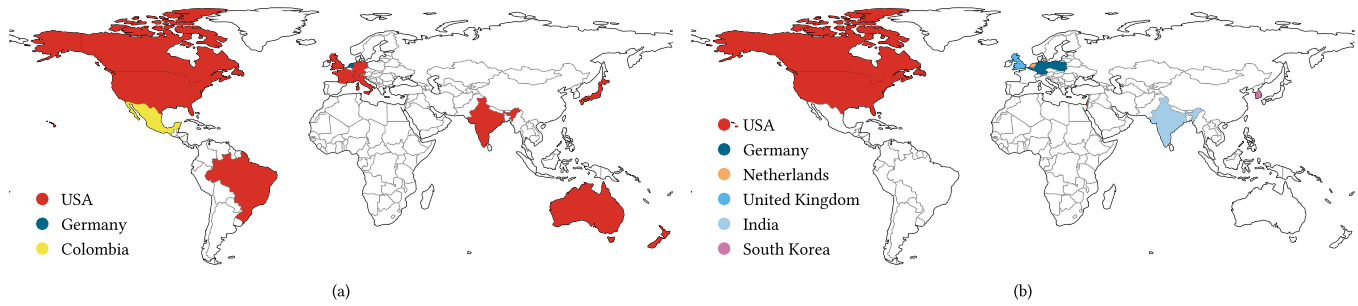


Figure 5: The maps are colored based on the majority of study participants’ country of origin. FAccT authors mainly used off-the-shelf datasets from the U.S (a), while datasets collected by the authors (e.g., interviews or surveys) have a more localized representation (b).

countries and the EIRD variables (Table 7). However, the differences in EIRD scores ($\Delta\tau$) observed using the shuffling method are only statistically significant ($p < 0.05$) for the Educated score (Table 8). This means that when CHI authors come from diverse countries, it is less likely for the resulting paper to be focused on Educated samples. We also examined the differences between FAccT and CHI for the Western variable based on the origin of the authors (i.e., whether they come from a single or multiple countries). Papers written by authors from multiple countries are less likely to have participants from exclusively Western countries (Figure 6). However, the difference in proportion is more apparent in CHI than in FAccT. This finding corroborates what was suggested by Linxen et al. [30] to “foster collaborations across Western and non-Western countries” to make CHI less WEIRD.

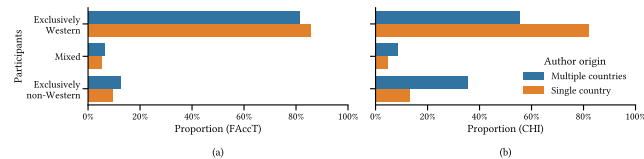


Figure 6: How author countries are associated with the proportion of sampled Western participants for (a) FAccT and (b) CHI. Papers produced by authors from multiple countries are more likely to sample from non-Western countries. The difference is more apparent on CHI, possibly due to the sample size difference between FAccT and CHI.

5 DISCUSSION

By analyzing 128 papers published between 2018 and 2022 at the FAccT conference, we studied the extent to which the conference relies on WEIRD samples and overused datasets. We found that 84% of FAccT papers are based exclusively on Western participants (representing less than 12% of the world’s population), 7% exclusively on non-Western participants, and 9% a mix of both. Compared to CHI, the leading conference on human factors in computing, FAccT papers draw less from Educated, Industrialized, Rich, and Democratic samples but more from Western ones. Comparing the WEIRD-ness between CHI and FAccT is crucial for identifying potential gaps in

research practices, methodologies, and datasets, given that FAccT is a relatively new conference with limited representation. While the lack of statistically significant differences does not refute the presence of WEIRD-ness, the comparison serves as an initial step towards recognizing disparities and promoting a more inclusive research community. Moreover, insights from the CHI conference offer foresight into future issues, providing valuable guidance for researchers toward more inclusive practices. Thus, the following subsections are based on results from both FAccT and CHI.

5.1 Implications

From a **theoretical standpoint**, our work contributes to the growing body of literature concerning WEIRD research, situated more broadly in the literature of Critical Computing [8]. We echo the criticisms raised by Laufer et al. [29] that the FAccT community is Western- and U.S.-centric, including the “*set of values and modes of discourse*.” A similar finding was also reported on van Berkel et al. [49]’s study comparing the same two conferences, albeit the study was conducted on a smaller sample size. Our study embarks on this strand of research by not only surfacing the WEIRD-ness of both conferences but also by uncovering their differences. In the broader Responsible AI context, questions remain to be quantitatively answered. Of prime importance is the so-called phenomenon of “*fair-washing*.” As Laufer et al. [29]’s described it, this phenomenon concerns the “*narrow notions of fairness [that are used] to condone existing practices*,” linking it with unregulated corporate influence on the scholarly discourse; industry interests may not necessarily be aligned with those of scholars, adding, or even amplifying biases, and minimally helping those affected by algorithmic systems. Waldman also points out that the discourse on privacy is heavily influenced by industry and its benefits [50].

From a **practical standpoint**, our findings speak to the need to take steps to make conferences less WEIRD. Here, we echo statements made by Sambasivan et al. [38] for “*incentivizing data excellence as a first-class citizen of AI, resulting in safer and more robust systems for all*.” What could be learned from best practices across conferences, and how organizing committees could embrace these practices? Next, we provide suggestions to alleviate FAccT’s WEIRD-ness:

Mandatory reproducibility statement: FAccT could prioritize reproducibility by promoting the replication and extension of findings and encouraging the sharing of data and materials. As a case in point, the Web (WWW) conference has enforced for reproducibility statement [1]. Such a statement not only allows for generalizable and reproducible findings but is also a way to help surface and mitigate the WEIRD-ness of a conference. Merely promoting research reproducibility may not adequately address WEIRD-ness if biased participant selection is the underlying issue. Nevertheless, when coupled with studies that highlight the importance of diverse participation in research and the provision of datasets containing non-WEIRD participants, it may result in the development of more varied off-the-shelf datasets for benchmarking purposes. With this approach, we foresee community-wide adoption of these datasets [52].

Mandatory data statement: FAccT could implement a mandatory data statement policy, requiring authors to report the geographic breadth of their participant samples. This also allows conference attendees to evaluate the generalizability of the research and understand its limitations. A **visualization tool** could do a real-time visualization of these points to provide a quick overview of datasets' characteristic to organizers, authors, and attendees.

Champion for author diversity across FAccT and other venues:

FAccT could prioritize diverse authorship by actively seeking and encouraging submissions from researchers from different countries and cultures. Limiting our analysis on FAccT may have resulted in missing this important point due to the small sample size. However, results from our CHI analysis (§4.2) show promises that such diversity can help make conferences less WEIRD (Figure 6). One idea could be to encourage collaboration using **speed dating** for attendees to encourage collaboration between authors that may not meet outside the conference. Attendees could randomly get assigned to tables for 5 minutes to introduce each other and give a short description of their to spark creativity and possible future collaborations.

Encourage cross-country online research methods to recruit from a diverse pool of participants: FAccT could create sessions or workshops focused on online research methods and their potential for studying geographically diverse samples. This could include training and support for researchers on effectively using these methods and overcoming any challenges they may face. FAccT could collaborate with other academic venues to incentivize data excellence by making data statements and datasheets for datasets mandatory for authors submitting their work.

5.2 Limitations and Future Work

Our work has three limitations that call for future research efforts. First, we limit our analyses to five years due to the FAccT's starting date, and the two sets of dates of the two conferences are not aligned because the CHI conference has a longer tenure. As the community grows, future studies could replicate our methodology, allowing for a broader comparison between the two communities (i.e., CHI and FAccT), accounting for more years of proceedings to be included. Along the same lines, the FAccT proceedings tend to attract a relatively larger number of studies concerned with datasets compared to CHI. Again, any differences attributed to that would

be possible to surface. The findings in §4.2 indicate insufficient evidence concluding that there is a significant difference in EIRD scores between datasets collected off-the-shelf and by the authors or between papers produced by authors from one country vs. multiple countries. However, these results should be interpreted with the caveat that the analysis was based on 128 papers, which had to be divided into these two groups. The limited sample size is, therefore, likely to decrease statistical power, leading to results with a low probability of accurately detecting a real effect or results influenced by both random and systematic errors. Nevertheless, by replicating Linxen et al. [30]'s dataset, we still observed statistically significant differences to some extent using a dataset of similar nature but nearly ten times larger in size. Second, we analyzed 30.8% of the papers due to the lack of relevant dataset information for the rest. We attribute this to not having a unified way of sharing datasets in FAccT. This speaks to the need for incentivizing data excellence and sharing, allowing for reproducibility (§5.1). Future research should also replicate this methodology in other conferences, establishing a single WEIRD metric (like impact factor for a conference's or journal's importance) for research diversity. By embracing such a metric, we foresee research diversification and focus on non-WEIRD populations. Third, this work considered only the WEIRD variables. Future research may well study datasets commonly used in the AI community through the lens of gender, class, and race, as well as a combination of demographics that can get lost when these factors are treated separately and as concrete groups [39]; a person could be part of multiple marginalized groups, which can result in compounding effect (e.g., black woman or an immigrant with a disability). Moreover, categorizing the world into uniform "civilizations" as defined by Huntington [22] is an oversimplification, ignoring the complex and nuanced nature of cultures that are evolving worldwide (including Western nations). Each culture is diverse and continually changing, shaped by unique combinations of historical, geographical, social, and economic factors. Recognizing and valuing the complexity and diversity of cultures is essential for inclusive and representative research. When examining the samples utilized in research studies, it may also appear that a significant proportion of them are from regions that are not WEIRD. However, upon closer inspection of factors such as the authors' background, institutions, methodologies, and results interpretation, it may still be apparent that a given research is heavily influenced by Western culture. Future studies could also focus on various perspectives, such as discussing counter-algorithm movements [4] or understanding the societal impact of algorithms [28].

6 CONCLUSION

Of 128 papers published at FAccT between 2018 and 2022, 84% were based on exclusively Western participants. Compared to the CHI conference, we found that FAccT is less EIRD. We did not find conclusive answers by exploring several hypotheses about the reasons for such differences. These differences are neither attributed to the types of datasets used (off-the-shelf vs. author-collected) nor to cross-affiliation collaborations. We encourage the FAccT community to increase sample diversity and address potential biases in the findings presented at the conference.

REFERENCES

- [1] 2022. *The Web Conference 2022*. Retrieved January 2023 from <https://www2022.thewebconf.org/cfp/research/>
- [2] Alan Agresti. 2010. *Analysis of Ordinal Categorical Data*. Vol. 656. John Wiley & Sons.
- [3] Jeffrey J Arnett. 2016. The neglected 95%: why American psychology needs to become less American. (2016).
- [4] Garfield Benjamin. 2022. #FuckTheAlgorithm: Algorithmic Imaginaries and Political Resistance. In *2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*. ACM, 46–57. <https://doi.org/10.1145/3531146.3533072>
- [5] Francois Buet-Golfouse and Islam Utyagulov. 2022. Towards Fair Unsupervised Learning. In *2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*. ACM, 1399–1409. <https://doi.org/10.1145/3531146.3533197>
- [6] Karoline Busse, Mohammad Tahaei, Katharina Krombholz, Emanuel von Zeszschwitz, Matthew Smith, Jing Tian, and Wenyuan Xu. 2020. Cash, Cards or Cryptocurrencies? A Study of Payment Culture in Four Countries. In *2020 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. IEEE, 200–209. <https://doi.org/10.1109/EuroSPW51379.2020.00035>
- [7] L. Elisa Celis, Lingxiao Huang, Vijay Keswani, and Nisheeth K. Vishnoi. 2019. Classification with Fairness Constraints: A Meta-Algorithm with Provable Guarantees. In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT* '19)*. ACM, 319–328. <https://doi.org/10.1145/3287560.3287586>
- [8] Rob Comber, Shaowen Bardzell, Jeffrey Bardzell, Mike Hazas, and Michael Muller. 2020. Announcing a New CHI Subcommittee: Critical and Sustainable Computing. *Interactions* 27, 4 (Jul 2020), 101–103. <https://doi.org/10.1145/3407228>
- [9] Frances Ding, Moritz Hardt, John Miller, and Ludwig Schmidt. 2021. Retiring adult: New datasets for fair machine learning. *Advances in Neural Information Processing Systems* 34 (2021), 6478–6490.
- [10] Dheeru Dua and Casey Graff. 2017. UCI Machine Learning Repository. <http://archive.ics.uci.edu/ml>
- [11] Economist Intelligence Unit. 2022. *Democracy Index 2021: the China challenge*. Technical Report. The Economist Group. <https://www.eiu.com/n/campaigns/democracy-index-2021> [Online; accessed 25. Jan. 2023].
- [12] European Union. 2020. Country profiles. https://european-union.europa.eu/principles-countries-history/country-profiles_en
- [13] Freedom House. 2022. Countries and territories. <https://freedomhouse.org/countries/freedom-world/scores>
- [14] Hana Frluckaj, Laura Dabbish, David Gray Widder, Huilian Sophie Qiu, and James D. Herbsleb. 2022. Gender and Participation in Open Source Software Development. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (Nov 2022), 31 pages. <https://doi.org/10.1145/3555190>
- [15] Mehrdad Ghadiri, Samira Samadi, and Santosh Vempala. 2021. Socially Fair K-Means Clustering. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21)*. ACM, 438–448. <https://doi.org/10.1145/3442188.3445906>
- [16] Przemyslaw A. Grabowicz, Nicholas Perello, and Aarshee Mishra. 2022. Marrying Fairness and Explainability in Supervised Learning. In *2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*. ACM, 1905–1916. <https://doi.org/10.1145/3531146.3533236>
- [17] David Harrison Jr and Daniel L Rubinfeld. 1978. Hedonic Housing Prices and the Demand for Clean Air. *Journal of Environmental Economics and Management* 5, 1 (1978), 81–102.
- [18] Lucy Havens, Melissa Terras, Benjamin Bach, and Beatrice Alex. 2020. Situated Data, Situated Systems: A Methodology to Engage with Power Relations in Natural Language Processing Research. In *Proceedings of the Second Workshop on Gender Bias in Natural Language Processing*. Association for Computational Linguistics, 107–124. <https://aclanthology.org/2020.gebnlp-1.10>
- [19] Joseph Henrich, Steven J Heine, and Ara Norenzayan. 2010. The Weirdest People in the World? *Behavioral and Brain Sciences* 33, 2-3 (2010), 61–83.
- [20] Hans Hofmann. 1994. German Credit Dataset. [https://archive.ics.uci.edu/ml/datasets/statlog+\(german+credit+data\)](https://archive.ics.uci.edu/ml/datasets/statlog+(german+credit+data))
- [21] Gary B. Huang, Manu Ramesh, Tamara Berg, and Erik Learned-Miller. 2007. *Labeled Faces in the Wild: A Database for Studying Face Recognition in Unconstrained Environments*. Technical Report 07-49. University of Massachusetts, Amherst.
- [22] Samuel P Huntington. 2000. The Clash of Civilizations? In *Culture and Politics*. Springer, 99–118.
- [23] Gareth James, Daniela Witten, Trevor Hastie, and Robert Tibshirani. 2013. *An Introduction to Statistical Learning*. Vol. 112. Springer.
- [24] Julia Angwin Jeff Larson. 2016. How We Analyzed the COMPAS Recidivism Algorithm. <https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm> [Online; accessed 20. Jan. 2023].
- [25] Anna Jobin, Marcello Lenca, and Effy Vayena. 2019. The global landscape of AI ethics guidelines. *Nature Machine Intelligence* 1, 9 (Sept. 2019), 389–399. <https://doi.org/10.1038/s42256-019-0088-2>
- [26] Josh Elliot. 2021. *The State Of AI Decision Making*. Forbes. Retrieved January 2023 from <https://www.forbes.com/sites/forbestechcouncil/2021/06/07/the-state-of-ai-decision-making>
- [27] Kimmo Karkkainen and Jungseock Joo. 2021. FairFace: Face Attribute Dataset for Balanced Race, Gender, and Age for Bias Measurement and Mitigation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. 1548–1558.
- [28] P. M. Krafft, Meg Young, Michael Katell, Jennifer E. Lee, Shankar Narayan, Micah Epstein, Dharma Dailey, Bernease Herman, Aaron Tam, Vivian Guetler, Corinne Bintz, Daniella Raz, Pa Ousman Jobe, Franziska Putz, Brian Robick, and Bissan Barghouti. 2021. An Action-Oriented AI Policy Toolkit for Technology Audits by Community Advocates and Activists. In *2021 ACM Conference on Fairness, Accountability, and Transparency*. ACM. <https://doi.org/10.1145/3442188.3445938>
- [29] Benjamin Laufer, Sameer Jain, A. Feder Cooper, Jon Kleinberg, and Hoda Heidari. 2022. Four Years of FAccT: A Reflexive, Mixed-Methods Analysis of Research Contributions, Shortcomings, and Future Prospects. In *2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*. ACM, 401–426. <https://doi.org/10.1145/3531146.3533107>
- [30] Sebastian Linxen, Christian Sturm, Florian Brühlmann, Vincent Cassau, Klaus Opwis, and Katharina Reinecke. 2021. How WEIRD is CHI?. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)*. ACM, 14 pages. <https://doi.org/10.1145/3411764.3445488>
- [31] Yong Ma, Yomna Abdelrahman, Barbarella Petz, Heiko Drewes, Florian Alt, Heinrich Hussmann, and Andreas Butz. 2022. Enthusiasts, Pragmatists, and Skeptics: Investigating Users' Attitudes Towards Emotion- and Personality-Aware Voice Assistants across Cultures. In *Proceedings of Mensch Und Computer 2022 (MuC '22)*. ACM, 308–322. <https://doi.org/10.1145/3543758.3543776>
- [32] Arsha Nagrani, Joon Son Chung, and Andrew Senior. 2017. VoxCeleb: A Large-Scale Speaker Identification Dataset. In *Proc. Interspeech 2017*. 2616–2620. <https://doi.org/10.21437/Interspeech.2017-950>
- [33] Jasmin Niess, Paweł W. Woźniak, Yomna Abdelrahman, Passant ElAgroudy, Yasmeen Abdrabou, Caroline Eckerth, Sarah Diefenbach, and Kristina Knaving. 2021. 'I Don't Need a Goal': Attitudes and Practices in Fitness Tracking beyond WEIRD User Groups. In *Proceedings of the 23rd International Conference on Mobile Human-Computer Interaction (MobileHCI '21)*. ACM, 14 pages. <https://doi.org/10.1145/3447526.3472062>
- [34] Laleh Nourian, Kristen Shinohara, and Garreth W. Tigwell. 2022. Digital Accessibility in Iran: An Investigation Focusing on Iran's National Policies on Accessibility and Disability Support. In *Proceedings of the 24th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '22)*. ACM, 5 pages. <https://doi.org/10.1145/3517428.3550385>
- [35] Alannah Oleson, Benjamin Xie, Jean Salac, Jayne Everson, F. Megumi Kivuva, and Amy J. Ko. 2022. A Decade of Demographics in Computing Education Research: A Critical Review of Trends in Collection, Reporting, and Use. In *Proceedings of the 2022 ACM Conference on International Computing Education Research - Volume 1 (ICER '22)*. ACM, 323–343. <https://doi.org/10.1145/3501385.3543967>
- [36] UNDP (United Nations Development Programme). 2022. Human Development Report 2021–22. *UNDP (United Nations Development Programme)* (2022). <http://report.hdr.undp.org>
- [37] Nithya Sambasivan, Garen Checkley, Amna Batool, Nova Ahmed, David Nemer, Laura Sanelly Gaytán-Lugo, Tara Matthews, Sunny Consolvo, and Elizabeth Churchill. 2018. "Privacy is Not for Me, It's for Those Rich Women": Performative Privacy Practices on Mobile Phones by Women in South Asia. In *Proceedings of the Fourteenth USENIX Conference on Usable Privacy and Security (SOUPS '18)*. USENIX Association, 127–142. <https://www.usenix.org/conference/soups2018/presentation/sambasivan>
- [38] Nithya Sambasivan, Shivani Kapania, Hannah Highfill, Diana Akrong, Praveen Paritosh, and Lora M Aroyo. 2021. "Everyone wants to do the model work, not the data work": Data Cascades in High-Stakes AI. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [39] Ari Schlesinger, W. Keith Edwards, and Rebecca E. Grinter. 2017. Intersectional HCI: Engaging Identity through Gender, Race, and Class. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, 5412–5427. <https://doi.org/10.1145/3025453.3025766>
- [40] Nandana Sengupta, Vidya Subramanian, Anwesh Mukhopadhyay, and Arul George Scaria. 2023. A Global South perspective for ethical algorithms and the State. *Nature Machine Intelligence* 5, 3 (2023), 184–186. <https://doi.org/10.1038/s42256-023-00621-9>
- [41] Sandipan Sikdar, Florian Lemmerich, and Markus Strohmaier. 2022. GetFair: Generalized Fairness Tuning of Classification Models. In *2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*. ACM, 289–299. <https://doi.org/10.1145/3531146.3533094>
- [42] Ryan Steed and Aylin Caliskan. 2021. Image Representations Learned With Unsupervised Pre-Training Contain Human-like Biases. In *2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21)*. ACM, 701–713. <https://doi.org/10.1145/3442188.3445932>
- [43] David Tedjopurnomo, Zhifeng Bao, Farhana Choudhury, Hui Luo, and A. K. Qin. 2022. Equitable Public Bus Network Optimization for Social Good: A Case Study of Singapore. In *2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*. ACM, 278–288. <https://doi.org/10.1145/3531146.3533092>

- [44] The World Bank Group. 2022. GDP per capita, PPP (current international \$). <https://data.worldbank.org/indicator/NY.GDP.PCAP.PP.CD>
- [45] The World Bank Group. 2022. GNI per capita, PPP (current international \$). <https://data.worldbank.org/indicator/NY.GNP.PCAP.PP.CD>
- [46] Valentin Todorov. 2020. *Competitive Industrial Performance Index 2020: Country Profiles (Report)*. Technical Report. United Nations Industrial Development Organization. <https://stat.unido.org/content/publications/competitive-industrial-performance-index-2020%253a-country-profiles> [Online; accessed 25. Jan. 2023].
- [47] Department of Economic United Nations and Population Division Social Affairs. 2022. World Population Prospects 2022: Data Sources. (UN DESA/POP/2022/DC/NO. 9). <https://population.un.org/wpp/> [Online; accessed 14. Dec. 2022].
- [48] United States Census Bureau. 2021. *American Community Survey (ACS)*. United States Census Bureau. Retrieved January 2023 from <https://www.census.gov/programs-surveys/acs/>
- [49] Niels van Berkel, Zhanna Sarsenbayeva, and Jorge Goncalves. 2023. The methodology of studying fairness perceptions in Artificial Intelligence: Contrasting CHI and FAccT. *International Journal of Human-Computer Studies* 170 (2023), 102954.
- [50] Ari Ezra Waldman. 2021. *Industry Unbound: The Inside Story of Privacy, Data, and Corporate Power*. Cambridge University Press. <https://doi.org/10.1017/9781108591386>
- [51] Darcia Wilkinson and Bart Knijnenburg. 2022. Many Islands, Many Problems: An Empirical Examination of Online Safety Behaviors in the Caribbean. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI '22)*. ACM, 25 pages. <https://doi.org/10.1145/3491102.3517643>
- [52] Mark D. Wilkinson, Michel Dumontier, IJsbrand Jan Aalbersberg, Gabrielle Appleton, Myles Axton, Arie Baak, Niklas Blomberg, Jan-Willem Boiten, Luiz Bonino da Silva Santos, Philip E. Bourne, Jildau Bouwman, Anthony J. Brookes, Tim Clark, Mercè Crosas, Ingrid Dillo, Olivier Dumon, Scott Edmunds, Chris T. Evelo, Richard Finkers, Alejandra Gonzalez-Beltran, Alasdair J.G. Gray, Paul Groth, Carole Goble, Jeffrey S. Grethe, Jaap Heringa, Peter A.C. 't Hoen, Rob Hooft, Tobias Kuhn, Ruben Kok, Joost Kok, Scott J. Lusher, Maryann E. Martone, Albert Mons, Abel L. Packer, Bengt Persson, Philippe Rocca-Serra, Marco Roos, Rene van Schaik, Susanna-Assunta Sansone, Erik Schultes, Thierry Sengstag, Ted Slater, George Strawn, Morris A. Swertz, Mark Thompson, Johan van der Lei, Erik van Mulligen, Jan Velterop, Andra Waagmeester, Peter Wittenburg, Katherine Wolstencroft, Jun Zhao, and Barend Mons. 2016. The FAIR Guiding Principles for scientific data management and stewardship. *Scientific Data* 3, 1 (March 2016). <https://doi.org/10.1038/sdata.2016.18>
- [53] Robert Wolfe and Aylin Caliskan. 2022. Markedness in Visual Semantic AI. In *2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*. ACM, 1269–1279. <https://doi.org/10.1145/3531146.3533183>
- [54] Kaiyu Yang, Klint Qinami, Li Fei-Fei, Jia Deng, and Olga Russakovsky. 2020. Towards Fairer Datasets: Filtering and Balancing the Distribution of the People Subtree in the ImageNet Hierarchy. In *2020 ACM Conference on Fairness, Accountability, and Transparency (FAT* '20)*. ACM, 547–558. <https://doi.org/10.1145/3351095.3375709>