



A Comparison of Differential Performance Metrics for the Evaluation of Automatic Speaker Verification Fairness

Oubaïda Chouchane¹, Christoph Busch^{2,3}, Chiara Galdi¹,
Nicholas Evans¹, Massimiliano Todisco¹

¹EURECOM, France

²Hochschule Darmstadt, Germany

³Norwegian University of Science and Technology (NTNU), Norway

¹`firstname[dot]lastname[at]eurecom[dot]fr`

²`firstname[dot]lastname[at]h-da[dot]de`

Abstract

When decisions are made and when personal data is treated by automated processes, there is an expectation of fairness – that members of different demographic groups receive equitable treatment. This expectation applies to biometric systems such as automatic speaker verification (ASV).

We present a comparison of three candidate fairness metrics and extend previous work performed for face recognition, by examining differential performance across a range of different ASV operating points. Results show that the Gini Aggregation Rate for Biometric Equitability (GARBE) is the only one which meets three functional fairness measure criteria.

Furthermore, a comprehensive evaluation of the fairness and verification performance of five *state-of-the-art* ASV systems is also presented. Our findings reveal a nuanced trade-off between fairness and verification accuracy underscoring the complex interplay between system design, demographic inclusiveness, and verification reliability.

Index Terms: automatic speaker verification, bias, fairness metrics

1. Introduction

The successful deployment of biometric systems is dependent upon acceptability in the target population. Acceptability requires that data subjects have confidence that they are treated in a fair manner by the biometric algorithm. The stratospheric rise in the use of artificial intelligence (AI) and its use for biometric person recognition has brought to prominence the question of whether or not it provides for the fair treatment of different demographic groups. The potential for certain demographic groups to be disadvantaged by the use of biometric person recognition systems was highlighted in the 2020 documentary film *Coded Bias*¹ which illustrates how a face recognition algorithm, which works well for individuals with lighter skin tone, fails to deliver the same experience to individuals with darker skin tones. Fairness is nowadays a fundamental requirement for biometric recognition algorithms.²

The NISTIR 8280³ investigated the extent of demographic-

related differential performance in the case of a face recognition system. Differential performance is the difference in score distributions for mated and non-mated comparisons.⁴ Demographic variables can be categorical,⁵ e.g. gender categories (male, female or neutral), or continuous⁶ e.g. an individual's skin color. There is potential for differential performance in the case of an automatic speaker verification (ASV) systems too, which can exhibit bias towards specific demographic groups, e.g. native (or non-native) speakers of the target language. An overview of algorithmic bias in biometric systems and a survey of the recent literature can be found in [3]. The reasons for bias are manifold and range from the use of unbalanced training datasets to systematic effects in the training procedures [4, 5, 6].

On the path to reach fair biometric systems, a fairness testing methodology is needed. Recent proposals for fairness measures [5, 7] are now under consideration to become the testing methodology in the draft International Standard ISO/IEC DIS 19795-10 [1]. To date, most of the scientific articles addressing the issue of fairness in biometric recognition systems focus on face recognition (FR). For instance, 39 out of 54 works reviewed in [3] involve face recognition. Thus, an open question remains as to whether the criteria for selecting metrics for fairness assessment, such as those suggested in [5], retain the same characteristics when applied to different biometric traits. In this article, we present a detailed study of the most relevant metrics currently proposed for the analysis of fairness in the context of ASV.

The remainder of the paper is structured to first discuss fairness metrics and criteria in biometric recognition systems. It then presents an overview of candidate metrics: the fairness discrepancy rate [1, 8]; the inequity rate [1]; the Gini Aggregation Rate for Biometric Equitability [1, 5]. We explain the computational methods of each, and key differences. Subse-

⁴Following the International Standard ISO/IEC 19795-1 [2] biometric performance is reported in terms of mated and non mated trials and in the form of the false match rate (FMR) and the false non-match rate (FNMR). In the language of the speaker verification community, these are target and non-target/impostor trials, the false alarm rate and the miss rate.

⁵The draft International Standard ISO/IEC DIS 19795-10 [1] defines categorical demographic variable as “demographic characteristic of an individual that is nominally or ordinally described”

⁶The draft International Standard ISO/IEC DIS 19795-10 [1] defines continuous demographic variable as “demographic characteristic of an individual that is observable, measurable, and that is not necessarily constrained to discrete categories”

¹The documentary film is described at: https://en.wikipedia.org/wiki/Coded_Bias

²<https://gdpr-info.eu/art-5-gdpr/>

³The draft International Standard ISO/IEC DIS 19795-10 [1] defines differential performances as “difference in biometric system metrics across different demographic groups”.

quent sections cover the experimental setup, including ASV systems, databases, and the fairness evaluation procedure. We then present experimental results and discussions, focusing on the evaluation of fairness metrics in the case of ASV systems, and conclude with a comprehensive assessment of system fairness, offering insights into performance and implications for equitable biometric recognition.

2. Fairness Metrics and Criteria

This section introduces fairness metrics initially proposed for the evaluation of face recognition systems, alongside essential criteria that such metrics must meet to effectively assess fairness. These guidelines ensure that the metrics are not only theoretically robust but also practically relevant in real-world applications.

2.1. Fairness Discrepancy Rate

The fairness discrepancy rate (FDR) focuses on the balance between the false match rate and the false non-match rate in assessing demographic-related differential performance [1, 8]. The measure introduces two components, namely the false positive differential (FPD) and the false negative differential (FND), which represent the maximum discrepancy in FMR and FNMR, respectively, between any two demographic groups d_i and d_j , belonging to a set D , at a specific discrimination threshold τ . The FDR quantifies these discrepancies, modulated by risk parameters α and $1 - \alpha$, to weight the relative importance of FMR and FNMR differences according to the security needs of a given application. High-risk situations, for instance, require a lower FMR to minimise security breaches. The FDR value ranges from 0 to 1, with 1 indicating full fairness and 0 indicating full unfairness. The FDR is calculated according to:

$$\text{FPD}(\tau) = \max(|\text{FMR}_{d_i}(\tau) - \text{FMR}_{d_j}(\tau)|) \quad (1)$$

$$\text{FND}(\tau) = \max(|\text{FNMR}_{d_i}(\tau) - \text{FNMR}_{d_j}(\tau)|) \quad (2)$$

for any $d_i, d_j \in D$.

$$\text{FDR}(\tau, \alpha) = 1 - (\alpha \text{FPD}(\tau) + (1 - \alpha) \text{FND}(\tau)) \quad (3)$$

In contrast to the goal in this paper, namely the comparison of different candidate metrics, the FDR metric has been used previously for the evaluation of ASV system fairness [9, 10, 11].

2.2. Inequity Rate

The inequity rate (IR) assesses fairness by examining the ratio of the boundary (minimum and maximum) FMR and FNMR values across different demographic groups d_i and d_j [1]. This is accomplished by evaluating the ratio of the highest FMR and FNMR to the lowest FMR and FNMR across all groups. Risk parameters α and $1 - \alpha$ are again used to scale the ratios before their aggregation. Importantly, in contrast to the the FDR, lower IR values indicate greater fairness. The IR is calculated according to:

$$\text{FPD}(\tau) = \frac{\max_{d_i} \text{FMR}_{d_i}(\tau)}{\min_{d_j} \text{FMR}_{d_j}(\tau)} \quad \forall d_i, d_j \in D \quad (4)$$

$$\text{FND}(\tau) = \frac{\max_{d_i} \text{FNMR}_{d_i}(\tau)}{\min_{d_j} \text{FNMR}_{d_j}(\tau)} \quad \forall d_i, d_j \in D \quad (5)$$

$$\text{IR}(\tau, \alpha) = \text{FPD}(\tau)^\alpha \cdot \text{FND}(\tau)^{(1-\alpha)} \quad (6)$$

2.3. The Gini Aggregation Rate for Biometric Equitability

The Gini aggregation rate for biometric equitability (GARBE) is derived from the Gini index, a measure for inequality [1, 5]. The GARBE metric uses a normalised Gini coefficient for n demographic groups. The normalization by $\frac{n}{n-1}$ is proposed in [12] to correct the downward bias when the number of samples (demographic groups) is small.

The Gini coefficient related to the FMR, is defined by:

$$G_{\text{FMR}}(\tau) = \frac{n}{n-1} \left(\frac{\sum_{i=1}^n \sum_{j=1}^n |\text{FMR}_{d_i}(\tau) - \text{FMR}_{d_j}(\tau)|}{2n^2 \overline{\text{FMR}}(\tau)} \right) \quad (7)$$

where $\overline{\text{FMR}}$ is the mean value. Similarly, the Gini coefficient related to the FNMR is defined by:

$$G_{\text{FNMR}}(\tau) = \frac{n}{n-1} \left(\frac{\sum_{i=1}^n \sum_{j=1}^n |\text{FNMR}_{d_i}(\tau) - \text{FNMR}_{d_j}(\tau)|}{2n^2 \overline{\text{FNMR}}(\tau)} \right) \quad (8)$$

for any $d_i, d_j \in D$.

In adapting to be consistent with the notation above, the pair of Gini coefficients are combined according to:

$$\text{FPD}(\tau) = G_{\text{FMR}}, \quad \text{FND}(\tau) = G_{\text{FNMR}} \quad (9)$$

$$\text{GARBE}(\tau, \alpha) = \alpha \text{FPD}(\tau) + (1 - \alpha) \text{FND}(\tau) \quad (10)$$

GARBE values range from 0 to 1, with 0 indicating full fairness and 1 indicating full unfairness.

2.4. Functional Fairness Measure Criteria

The fundamental aim of fairness metrics is to evaluate and identify the most equitable classification algorithms. Howard et al. [13] define the requirement for essential attributes of such metrics as the functional fairness measure criteria (FFMC), which reflect their interpretability and applicability:

1. **FFMC.1:** The contributions of FMR and FNMR to the fairness metric should be intuitive across typical risk parameters and operationally relevant error rates.
2. **FFMC.2:** The metric needs defined boundaries, with minimum and maximum values, to establish clear benchmarks.
3. **FFMC.3:** The metric must remain computable for demographic groups with no observed errors, a condition becoming more common with ever-more accurate biometric algorithms.

3. Experimental setup

In this section we present the ASV systems used for the fairness metrics evaluation, database and the fairness evaluation procedure.

3.1. Speaker verification systems

We use five different ASV systems in our assessment of the fairness metrics, with each system manifesting unique structural and functional characteristics. The ECAPA system [14] employs a standard ECAPA-TDNN [15] structure with 3 SE-Res2Block modules to extract a 192-dimensional speaker embedding. As a back-end, ECAPA uses the cosine similarity. ResNetSE34L [16] is a slimmed-down version of the original ResNet-34 [17]. It uses self-attentive pooling (SAP) [18] for aggregating frame-level features into utterance-level features,

Table 1: ASV performance in terms of pooled EER and FMR/FNMR at the threshold corresponding to the pooled EER, across nine groups of different nationalities. Color background transitions from better (green) to mid (yellow) to lower (red) performances.

Nationality	ASV Systems									
	ERes2Net		CAM++		ECAPA		ResNetSE34V2		ResNetSE34L	
	FMR (%)	FNMR (%)	FMR (%)	FNMR (%)	FMR (%)	FNMR (%)	FMR (%)	FNMR (%)	FMR (%)	FNMR (%)
USA	1.22	1.04	1.40	1.54	1.13	1.45	1.36	1.68	2.76	2.08
UK	0.68	0.45	0.41	0.14	0.72	0.23	0.50	0.50	2.08	1.90
Germany	0.59	2.81	0.63	4.26	2.49	6.34	2.31	6.34	4.17	8.11
Australia	0.68	0.27	1.99	0.14	1.90	0.27	0.86	0.77	1.54	1.72
Italy	1.95	2.58	1.40	2.72	2.58	3.85	3.26	2.54	3.53	4.08
India	2.31	0.09	1.90	0.14	2.76	1.00	6.11	0.00	5.53	0.09
Ireland	0.18	2.04	0.82	2.31	0.77	1.45	0.45	2.36	0.86	4.21
New Zealand	1.86	0.27	1.86	0.18	2.17	0.27	0.95	1.18	1.86	2.67
Canada	1.13	1.31	1.40	0.86	1.31	1.22	1.13	1.63	5.30	2.17
Pooled EER (%)	1.18		1.37		1.79		1.88		3.01	

which focuses on the most informative frames. ResNetSE34L uses the squared Euclidean distance as a distance metric. ResNetSE34V2 [19] is a performance-optimised variant of ResNet-34. The stride is removed at the first convolutional layer to reduce computational cost. ResNetSE34V2 adopts attentive statistics pooling (ASP) [20] for temporal frames aggregation. ERes2Net [21] addresses the limitations of the Res2Net structure by integrating local and global feature fusion, capturing both detailed and holistic patterns in the input signal. Last, the CAM++ [22] architecture mainly consist of a front-end convolution module and a densely connected time delay neural network (D-TDNN) backbone. In each D-TDNN layer, an improved context-aware masking (CAM) module is included. The CAM++ system employs multi-granularity pooling to capture discriminative speaker characteristics.

3.2. Databases

The pre-trained models^{7,8,9} of the ASV systems used in our experiments are trained using the development set of the VoxCeleb2 database [23]. It contains utterances collected from 5,994 speakers. Despite VoxCeleb2 being a multilingual dataset, the predominance of English speakers, and consequently also the English language, creates significant imbalance. All evaluation work is performed using the VoxCeleb1 development and test sets [24] combined. We created a balanced protocol to assess the utility of the ASV systems. This protocol was created from a selection of speakers from nine different nationalities: USA, UK, Germany, Australia, Italy, India, Ireland, New Zealand, and Canada. From each of these nationality groups, eight speakers were chosen at random, resulting in a total of 72 speakers. For each speaker, 24 utterances were selected. The ASV protocol¹⁰ of the pooled data consists of 39,744 comparison trials, with a balanced number of 2,208 mated and 2,208 non-mated combinations for each nationality. Due to limitations in the available data, we could not include more speakers in the study. We aimed to cover as many nationalities as possible while maintaining an equal number of speakers for each nationality.

⁷ <https://github.com/TaoRuijie/ECAPA-TDNN>

⁸ https://github.com/clovaai/voxceleb_trainer

⁹ <https://github.com/alibaba-damo-academy/3D-Speaker/tree/3dspeaker>

¹⁰ https://github.com/OubaidaOubaida/FairnessMetricsEvaluation/blob/main/pooled_data.txt

3.3. Fairness evaluation procedure

This section is dedicated to the fairness procedures and assessment of ASV systems. Presented in Table 1 are results for a preliminary analysis of ASV performance in terms of pooled EER and FMR/FNMR at the threshold corresponding to the pooled EER, across nine groups of different nationalities. As expected, pooled EER results show that different ASV systems exhibit varying levels of performance. The comparison also reveals consistent differences in the FMR and FNMR among nationality groups. For instance, results for the ERes2Net system are best for the UK group, with low FMR and FNMR. In contrast, while the same system is similarly secure for the German group (similar FMR), it does not offer the same level of convenience (higher FNMR). Interestingly, the opposite behaviour is observed for the Indian group. Similar diverging results are observed for other nationality groups. This analysis underscores the importance of a *single* measure which reflects fairness across the *full set* of groups. Such an approach is crucial for ensuring that ASV systems are equitable and do not disproportionately disadvantage any particular group based on nationality or other demographic factors.

Noting that the EER is not suited to the assessment of *any* binary classifier in the case that a particular application calls for the prioritisation of a lower rate of FMR or FNMR [8, 10], it is similarly unsuitable as a measure of fairness. Since the use of different operating points necessarily entails a trade-off between the FMR and the FNMR, and as argued in [8], any measure of fairness must take difference in both into account. While one can average FMR and FNMR rates across groups, this leads to two measures which still need further interpretation if they are to provide a measure of fairness. This is exactly what each of the three candidate metrics delivers.

Our methodology to evaluate the proposed fairness metrics is aligned with that in [13]. To this end, we adopt a benchmark threshold that corresponds to an FMR of 0.1% for the initial evaluation. In addition, and going beyond the work in [13], we follow the requirements of ISO/IEC DIS 19795-10 [1] by investigating a range of thresholds, specifically within an FMR interval of 0.1%-10%, alongside a comprehensive risk parameter range of [0,1]. The rationale behind the selection of a 0.1%-10% FMR interval is twofold: first, though it depends on the application, today's most advanced ASV systems achieve acceptable levels of FNMR at FMRs in the order of 0.1%; second, ASV systems (or indeed any other biometric system) with an FMR surpassing 10% might not limited use. Our objective is to examine any variation in fairness at a range of representa-




Figure 1: FDR values using 5 automatic speaker verification systems at a threshold corresponding to FMR = 0.1%




Figure 2: FDR values using 5 automatic speaker verification systems at a range of thresholds corresponding to a FMR varying from 0.1% to 10%

tive operating points, for five distinct systems, thereby gaining a holistic understanding of metric behaviour. This examination is crucial as, in real-world scenarios, each system operates optimally at a threshold that is contingent upon the specific application it serves.

4. Experimental results and discussion

In this section, we present an assessment of the fairness metrics presented in Section 2, first for a fixed threshold which provides an FMR = 0.1% and, second, for a range of thresholds from 0.1% to 10%. Last, we evaluate each fairness metric in terms of the three FFCMs described in Section 2.4.

4.1. Metrics evaluation results at a fixed threshold

We evaluate the fairness metrics performance across five different ASV systems. We aggregate the evaluation results to understand how these metrics behave under different systems.

For each metric we present three plots. The first is the metric density distribution when the importance of the FMR/FNMR differentials are equal ($\alpha = 0.5$), such as presented in the reference work [13]. The second is to show the density distribution of the metric for all alpha values in $[0,1]$. These two plots show the spread of the metrics values to ascertain whether the metric offers an intuitive comparison between the systems, as well as their behavior under varying risk parameter alpha. The third

graph is the distribution of the FPD and FND terms for alpha in $[0,1]$ to study the scales of the error rates.

To align our work with that in [13], we evaluated the fairness metric at a fixed threshold. Experiments were performed with decision thresholds τ set so as to produce a FMR of 0.1%

4.1.1. FDR evaluation

Our assessment starts with the calculation of the FDR metric described in Section 2.1. The FDR values are visually depicted in Figure 1. Figure 1(a) and 1(b) show FDR values concentrated in the range $[0.82-1]$. It is not intuitive to compare which system is more fair and to assess the impact of the risk parameter α , especially in the case of mostly fair systems. Figure 1(c) illustrates that the differential terms (FPD, related to the FMR, and FND), related to the FNMR, operate on significantly disparate scales. Aggregating these terms presents a challenge in accurately configuring them with the term α . This complexity underscores the difficulty of intuitively measuring the contributions of FMR and FNMR within the FDR metric. Consequently, this approach fails to satisfy the criteria established by the first FFCM principle (Section 2.4).

4.1.2. IR evaluation

The evaluation of the IR metric presented in Section 2.2 revealed minFMR values of 0 for certain subgroups. This renders




Figure 3: GARBE values using 5 automatic speaker verification systems at a threshold corresponding to FMR = 0.1%




Figure 4: GARBE values using 5 automatic speaker verification systems at a range of thresholds corresponding to a FMR varying from 0.1% to 10%

the computation of the FPD term of the IR equation (4) unfeasible, subsequently leading to the inability to calculate the IR. Among the five evaluated ASV systems, the IR metric is only computable for the ResNetSEV2 system, with the value being 13.35. This shows that the IR metric does not fulfill the third FFCM (Section 2.4). Consequently, this casts doubt on the suitability of IR as a reliable metric for fairness assessment in such contexts. Additionally, the ratio nature of the IR introduces another complexity: there is no upper limit to the IR values. This means that the IR can vary greatly and potentially reach extremely high numbers, further complicating its interpretation.

4.1.3. GARBE evaluation

We now turn our attention to the GARBE metric described in Section 2.3. The results illustrated in Figure 3(a) and 3(b) exhibit a wider range compared to the FDR values. This range extends from 0.19 to 0.61 for α in $[0,1]$, which is about half of the theoretical range. This provides a more intuitive comparison between the systems and assessment of the impact of α . An additional critical finding is related to the FPD and FND terms, depicted in Figure 3(c). These terms are scaled to a comparable magnitude. Specifically, the median value for the FPD term is identified at 0.55, while the median for the FND term is observed at 0.29. The normalised Gini coefficient computation reduced the discrepancy of the differential terms scale. The

impact of α is then more pronounced.

4.2. Metrics evaluation results at different thresholds

In an extension of the previous work, we expand the assessment to include not just different systems, but also different systems with different operating points. We expanded the range of our analysis by varying the threshold of the five ASV systems from that corresponding to a 0.1% to 10% FMR. This approach ensures an assessment reflective of real-world scenarios.

For each metric we present three plots as in the previous section. The only difference this time is the FMR of the plots range from 0.1% to 10%.

4.2.1. FDR evaluation

Despite adjusting both the threshold and α values, the FDR values display a consistent trend, with values concentrated between 0.72 and 1, as shown in Figures 2(a) and 2(b). It is still not intuitive to compare which system is more fair as most of the FDR values almost overlap for all α s. Moreover, the scale disparity between the FPD and FND terms is still persistent, as depicted in Figure 2(c). This implies the intuitiveness of the contributions of FMR and FNMR to the FDR metric. Hence, the FDR metric does not meet the FFCM.1.




Figure 5: IR values using 5 automatic speaker verification systems at a range of thresholds corresponding to a FMR varying from 0.1% to 10%

4.2.2. IR evaluation

Results for the IR metric confirm the impossibility of its computation in certain scenarios, even with a variable threshold range. The FFMC.3 is still not fulfilled. The selective representation of 8.6% of computable values in Figures 5(a) and 5(b) illustrates the extensive range of the IR, reaching up to 200 in certain instances. This confirms that the IR metric is not bounded, therefore, it does not meet FFMC.2. However, the FPD and FND terms are on a similar scale, as demonstrated in Figure 5(c). Use of a ratio-based approach means that the IR metric offers a more balanced comparison between terms. Thus, the FFMC.1 is fulfilled.

4.2.3. GARBE evaluation

The GARBE metric is effective in avoiding the limitations seen in the previous metrics effectively. As Figure 4(a) and 4(b) show, GARBE values consistently span the full theoretical range from 0 to 1. Furthermore, Figure 4(b) demonstrates the sensitivity of GARBE to α changes. The FPD and FND terms, presented in Figure 3(c), with median values of 0.21 and 0.88 respectively, are on the same scale. This ensures that both terms contribute significantly to the fairness assessment, meeting FFMC.1.

Analysis of the boxplots shows a swap in the boxes of the FPD and FND terms between Figures 1(c) and 2(c), as well as between 3(c) and 4(c). The swap occurs because the FPD is higher or lower than the FND before reaching certain thresholds, and vice versa. More precisely, within the FDR metric analysis, for thresholds that yield an FMR lower than 0.9%, the FND term surpasses the FPD term. Conversely, for thresholds resulting in an FMR above 0.9%, the FPD term is higher. Regarding the GARBE metric, thresholds that produce an FMR below 0.4% result in a higher FPD than FND. This explains the observed variation in the boxplot positions when the FMR is set at 0.1% and for FMR ranges from 0.1% to 10%.

4.3. Summary of the Fairness Metrics Criteria

Our evaluation of fairness metrics for ASV within the framework of the Functional Fairness Measure Criteria (Section 2.4) presents varied results for the FDR, IR, and GARBE metrics. These outcomes are summarised in Table 2.

The FDR metric is inherently bounded therefore satisfies

Table 2: Summary of Fairness Measures Criteria for ASV

FFMC Criteria	FDR	IR	GARBE
FFMC.1		✓	✓
FFMC.2	✓		✓
FFMC.3	✓		✓

the FFMC.2 criterion. This quality provides an interpretable measure of fairness with a value of 1 to a perfectly fair system and a value of 0 to a perfectly unfair system. The FDR metric also maintains computability when FNMR or FMR is zero, adhering to FFMC.3. However, the differential terms FPD and FND exist at vastly different scales when using a normal range of risk parameters (α in $[0,1]$) and operationally relevant error rates (FMR in $[0.1\%,10\%]$). This renders the interpretation of the contribution of the FMR and FNMR in the computation of the FDR metric challenging. Hence, this metric does not meet FFMC.1.

The IR metric fulfills FFMC.1 by taking a ratio approach which better balances the contributions of FPD and FND terms. Nevertheless, it faces limitations in meeting FFMC.2 and FFMC.3 due to its unbounded nature, which makes it challenging to set benchmarks and renders it incalculable when FNMR or FMR reach zero.

GARBE emerges as the most robust metric, fully meeting all FFMC criteria. By employing the Gini coefficient, as specified in Equations 7 and 8, the FPD and FND terms are converted to a same scale before their aggregation. This normalisation is key to fulfilling FFMC.1, as it makes the contribution of the FMR and FNMR to the calculation of the GARBE metric intuitive. It ensures a nuanced and balanced representation of the metric measures across varying α values. Second, GARBE meets FFMC.2 by maintaining set boundaries, which allows for the establishment of clear benchmarks. Furthermore, it remains calculable even where error rates are zero, thereby meeting FFMC.3.

Our analysis validates and extends the findings of the study of fairness for face recognition reported in [13], reinforcing the conclusion that the GARBE metric stands out as the most suitable for evaluating fairness in biometric systems. This congruence illustrates that GARBE are not confined to face recognition systems but extend also to studies of fairness in automatic speaker verification.




Figure 6: GARBE of different ASV systems for different decision $\tau = FMR_x$ where x varies from 0.1% to 10% and for $\alpha = 0.5$.




Figure 7: Detection error tradeoff (DET) curve of different ASV systems.

5. Fairness and ASV assessment

While it is not the primary objective of this work, our attention now shifts to an analysis of the ASV systems themselves. We report an assessment of the five ASV systems in terms of verification performance (FMR vs. FNMR) and fairness, here assessed using only the GARBE metric.

Figure 6 shows a plot of GARBE values over an FMR range of between 0.1% and 10%, for $\alpha = 0.5$. Figure 7 shows a corresponding detection error trade-off (DET) plot. The ResNetSE34L system shows the lowest GARBE values across all FMR thresholds, indicating that the system is the fairest. However, in terms of verification performance, the same ResNetSE34L system is the worst. Improved fairness (lower differential performance across groups) in this case comes at the cost of degraded verification performance.

GARBE values for the least fair system, ResNetSE34V2, are highest across a majority of operating points, indicating significant differential performance across groups. Verification ac-

curacy is, however, better than that for the ResNetSE34L system, another indication of an apparent trade-off between performance and fairness. The ERes2Net and CAM++ systems are not only comparably fair, albeit not *the* fairest, especially for FMRs under 1%, they are among the best performing.

The ECAPA model exhibits different behaviour. Verification performance is average, but there is substantial variation in GARBE values between lower and higher FMRs, more so than for other systems. This pronounced fluctuation suggests a particular sensitivity for the ECAPA model, a distinct aspect of its operational characteristics.

6. Conclusions

We present a comparison of three different fairness metrics — FDR, IR, and GARBE — in the context of automatic speaker verification (ASV) alongside an analysis of fairness and verification performance for five *state-of-the-art* ASV systems. Results show that the Gini Aggregation Rate for Biometric Equitability (GARBE) metric stands out as that which best fulfills the Functional Fairness Measure Criteria (FFMCs).

Our analysis reveals a delicate balance between fairness and accuracy. The fairest system has the worst verification performance. The system with the best verification performance has average fairness. These findings show the challenge to strike a balance between fairness and verification performance. Given the requirement for fairness, it should be evaluated during the development of ASV systems, as it should be during the development of any biometric system; there is no guarantee that a focus on raw verification performance alone will produce equitable solutions. We recommend as future work to develop *fairness by design* concepts.

7. Acknowledgements

This work is supported by the TRSPAS-ETN project funded by the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 860813.

8. References

- [1] ISO/IEC JTC1 SC37 Biometrics, *ISO/IEC DIS 19795-10. Information Technology – Biometric Performance Testing and Reporting – Part 10: Quantifying biometric system performance variation across demographic groups*, International Organization for Standardization, 2023.
- [2] ISO/IEC JTC1 SC37 Biometrics, *ISO/IEC 19795-1:2021. Information Technology – Biometric Performance Testing and Reporting – Part 1: Principles and Framework*, International Organization for Standardization, June 2021.
- [3] Pawel Drozdowski, Christian Rathgeb, Antitza Dantcheva, Naser Damer, and Christoph Busch, “Demographic bias in biometrics: A survey on an emerging challenge,” *Trans. on Technology and Society (TTS)*, vol. 1, no. 2, pp. 89–103, June 2020.
- [4] Pawel Drozdowski, Christian Rathgeb, and Christoph Busch, “The watchlist imbalance effect in biometric face identification: Comparing theoretical estimates and empiric measurements,” in *Intl. Conf. on Computer Vision Workshops (ICCVW)*, New York, October 2021, pp. 1–9, IEEE.

- [5] John Howard, Eli Laird and Yevgeniy Sirotnin, Rebecca Rubin, Jerry Tipton, and Arun Vemury, “Evaluating proposed fairness models for face recognition algorithms,” in *Proc. Intl. Conf. on Pattern Recognition*, 2022.
- [6] Christian Rathgeb, Pawel Drozdowski, Dinusha C. Frings, Naser Damer, and Christoph Busch, “Demographic fairness in biometric systems: What do the experts say?,” *IEEE Technology and Society Magazine*, vol. 41, pp. 71–82, December 2022.
- [7] Ketan Kotwal and Sébastien Marcel, “Fairness index measures to evaluate bias in biometric recognition,” in *Proc. Intl. Conf. on Pattern Recognition*, 2022.
- [8] Tiago de Freitas Pereira and Sébastien Marcel, “Fairness in biometrics: A figure of merit to assess biometric verification systems,” *IEEE Transactions on Biometrics, Behavior, and Identity Science*, vol. 4, no. 1, pp. 19–29, 2022.
- [9] Mariel Estevez and Luciana Ferrer, “Study on the fairness of speaker verification systems on underrepresented accents in english,” *arXiv preprint arXiv:2204.12649*, 2022.
- [10] Raghuv eer Peri, Krishna Somandepalli, and Shrikanth Narayanan, “A study of bias mitigation strategies for speaker recognition,” *Computer Speech & Language*, vol. 79, pp. 101481, 2023.
- [11] Oubaïda Chouchane, Michele Panariello, Chiara Galdi, Massimiliano Todisco, and Nicholas Evans, “Fairness and privacy in voice biometrics: A study of gender influences using wav2vec 2.0,” in *2023 International Conference of the Biometrics Special Interest Group (BIOSIG)*. IEEE, 2023, pp. 1–7.
- [12] George Deltas, “The small-sample bias of the gini coefficient: results and implications for empirical research,” *Review of economics and statistics*, vol. 85, no. 1, pp. 226–234, 2003.
- [13] John J. Howard, Eli J. Laird, Rebecca E. Rubin, Yevgeniy B. Sirotnin, Jerry L. Tipton, and Arun R. Vemury, “Evaluating proposed fairness models for face recognition algorithms,” in *International Conference on Pattern Recognition*. Springer, 2022, pp. 431–447.
- [14] Rohan Kumar Das, Ruijie Tao, and Haizhou Li, “Hlt-nus submission for 2020 nist conversational telephone speech sre,” *arXiv preprint arXiv:2111.06671*, 2021.
- [15] Brecht Desplanques, Jenthe Thienpondt, and Kris Demuyne, “ECAPA-TDNN: Emphasized channel attention, propagation and aggregation in tdnn based speaker verification,” in *Proc. INTERSPEECH 2020*, 2020, pp. 3830–3834.
- [16] Joon Son Chung, Jaesung Huh, Seongkyu Mun, Minjae Lee, Hee Soo Heo, Soyeon Choe, Chiheon Ham, Sunghwan Jung, Bong-Jin Lee, and Icksang Han, “In defence of metric learning for speaker recognition,” in *Proc. INTERSPEECH 2020*, 2020, pp. 2977–2981.
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [18] Weicheng Cai, Jinkun Chen, and Ming Li, “Exploring the encoding layer and loss function in end-to-end speaker and language recognition system,” in *Proc. The Speaker and Language Recognition Workshop (Odyssey 2018)*, 2018, pp. 74–81.
- [19] Yoochwan Kwon, Hee-Soo Heo, Bong-Jin Lee, and Joon Son Chung, “The ins and outs of speaker recognition: lessons from voxsrc 2020,” in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 5809–5813.
- [20] Koji Okabe, Takafumi Koshinaka, and Koichi Shinoda, “Attentive statistics pooling for deep speaker embedding,” in *Proc. INTERSPEECH 2018*, 2018, pp. 2252–2256.
- [21] Yafeng Chen, Siqi Zheng, Hui Wang, Luyao Cheng, Qian Chen, and Jiajun Qi, “An enhanced Res2Net with local and global feature fusion for speaker verification,” in *Proc. INTERSPEECH 2023*, 2023, pp. 2228–2232.
- [22] Hui Wang, Siqi Zheng, Yafeng Chen, Luyao Cheng, and Qian Chen, “Cam++: A fast and efficient network for speaker verification using context-aware masking,” in *Proc. INTERSPEECH 2023*, 2023, pp. 5301–5305.
- [23] Joon Son Chung, Arsha Nagrani, and Andrew Zisserman, “Voxceleb2: Deep speaker recognition,” in *Proc. INTERSPEECH 2018*, 2018, pp. 1086–1090.
- [24] Arsha Nagrani, Joon Son Chung, and Andrew Zisserman, “Voxceleb: a large-scale speaker identification dataset,” in *Proc. INTERSPEECH 2017*, 2017, pp. 2616–2620.