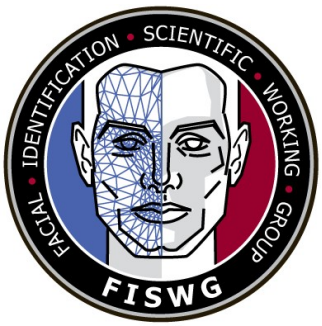

Disclaimer:

As a condition to the use of this document and the information contained herein, the Facial Identification Scientific Working Group (FISWG) requests notification by e-mail before or contemporaneously to the introduction of this document, or any portion thereof, as a marked exhibit offered for or moved into evidence in any judicial, administrative, legislative, or adjudicatory hearing or other proceeding (including discovery proceedings) in the United States or any foreign country. Such notification shall include: 1) the formal name of the proceeding, including docket number or similar identifier; 2) the name and location of the body conducting the hearing or proceeding; and 3) the name, mailing address (if available) and contact information of the party offering or moving the document into evidence. Subsequent to the use of this document in a formal proceeding, it is requested that FISWG be notified as to its use and the outcome of the proceeding. Notifications should be sent to: chair@fiswg.org

Redistribution Policy:

FISWG grants permission for redistribution and use of all publicly posted documents created by FISWG, provided that the following conditions are met:

Redistributions of documents, or parts of documents, must retain the FISWG cover page containing the disclaimer.

Neither the name of FISWG, nor the names of its contributors, may be used to endorse or promote products derived from its documents.

Any reference or quote from a FISWG document must include the version number (or creation date) of the document and mention if the document is in a draft status.

Understanding and Testing for Face Recognition Systems Operation Assurance: Identity Ground Truth

Purpose

This document provides guidelines and techniques to help administrators of automated face recognition systems (FRS) produce advanced and accurate recognition statistics on the face recognition systems.

The intended audience of this document is system owners, system users, and system administrators of existing automated face recognition systems. Outside the scope of this document include, but not necessarily limited to, system setup, system tuning, workflow management and improvement, and proof of concept pilots.

This document is a follow on from the FISWG document: “Understanding and Testing for Face Recognition Systems Operation Assurance” (version 1.0, 2020.12.11)

The issues presented in this document form a base for other considerations and advanced topics when testing (e.g., system setup and tuning) which will be covered in future FISWG documents.

Scope

The scope of this document is to provide a detailed process and examples of testing and repairing identity ground truth in facial data sets which is a critical initial step before recognition statistics are created and reviewed. This document does not address facial accuracy but is focused solely on testing and correcting identity ground truth. Any facial algorithm can be used with these processes. It is assumed that all facial images create proper templates. The facial data set used in this document is the “Labeled Faces in the Wild” (LFW) but conceptually any other facial data set with identity ground truth can be used.

Background

When doing accuracy profiling, there is always one key aspect which must be addressed first: what is the identity ground truth within the images? All data sets will potentially have some corruption in the identity ground truth with the data. Detecting and correcting this so pristine results can be reviewed is always a critical portion of profiling.

This type of identity ground truth verification will potentially exist in all data sets no matter where the data sets come from. This is an iterative process as the agency learns the algorithms, the data, and how the two interact with each other. If proper care is not given in these early stages, then incorrect assumptions on the outcomes will be made. It’s critical to understand this process with an investigative mindset before the

agency gets to the operational data sets which may have identity corruption and image quality issues that may be large but not uncommon in operational deployments. If the agency gets to the operational data set without a firm awareness and knowledge base on the how the core algorithms work with verified data, then the agency could be incorrectly assessing and measuring performance of the FRS. Agencies need to lay the groundwork to know and trust the algorithms before they get to the possibly unmanaged and unknown operational data.

Most of the work in these processes is on creating the testing frameworks and understanding how to repeatedly run tests, make corrections, and do retesting with what has been learned. Once the frameworks and the processing are understood, then the agency can make diligent progress, but it takes time and focus. The outcomes are worth the time spent as you begin to understand how the data interacts with the algorithms which give the agency the ability to trust the solution and not just assume the data is invalid.

Setting up frameworks to do enrollment and searching while recording results is fairly mechanical as you learn the facial algorithms and the data sets to develop proper profiling. Understanding the data and building frameworks to analytically qualify the results is not trivial but must be done so effective operational metrics can be derived and applied.

Before doing this analysis on operational data, it is recommended that the agency develop and test the framework on experimental datasets. After some experience is gained in this process and confidence that the process is correct, one could then assess the operational dataset.

This document describes procedures to assess an experimental dataset, which can be replicated before assessing operational datasets.

Data Set

Care should be taken in selecting data sets to profile experimentally. It is recommended to select data sets which:

- Have operational relevancy
- Have consistent image quality aspects: type of capture, size of images, subject poses, etc.
- Have sufficient identities and images to test with. This decision will be agency specific.
- Includes associated identity ground truth information which links each image to a unique identity

The data set used for this document is the LFW data set available at: http://vis-www.cs.umass.edu/lfw/ See Appendix 2: "**LFW Data Set Information**" for more details. Conceptually any other facial data set with identity ground truth can be used.

LFW is a widely used open source data set which will work well for this specific document serving as an introductory data set. Information on this data set includes:

- Has smaller but consistent image sizes and file formats
- Has over 5,700 identities and over 13000 images
- Has a wide range of subjects: sex, pose, lighting, etc
- Has stated identity ground truth errors

A key point in the LFW errata is that there are known errors in the LFW data set. While the LFW URL addresses these errors, this document will show how to locate them and give examples on how determining identity ground truths in relating images to identities is critical and needs to be addressed in any operational testing scenario.

Ground Truth Process

Step 1: Enroll the facial images into a facial gallery for searching.

Step 2: Search the facial images against the facial gallery. The number of candidates returned for this document was 100. This number may vary with agency specifics and the biometric algorithm deployed. It is recommended to test with a larger number of candidates than what may be operationally used so that deeper accuracy investigations can be analyzed. Do not use scoring thresholds.

Step 3: Analyze the scoring to delineate every candidate in all 1:N searches:

- Image file name

- Image identity
- Score
- Rank
- Mate scoring
- Imposter scoring

When this was done for this document the following accuracy curves were obtained:

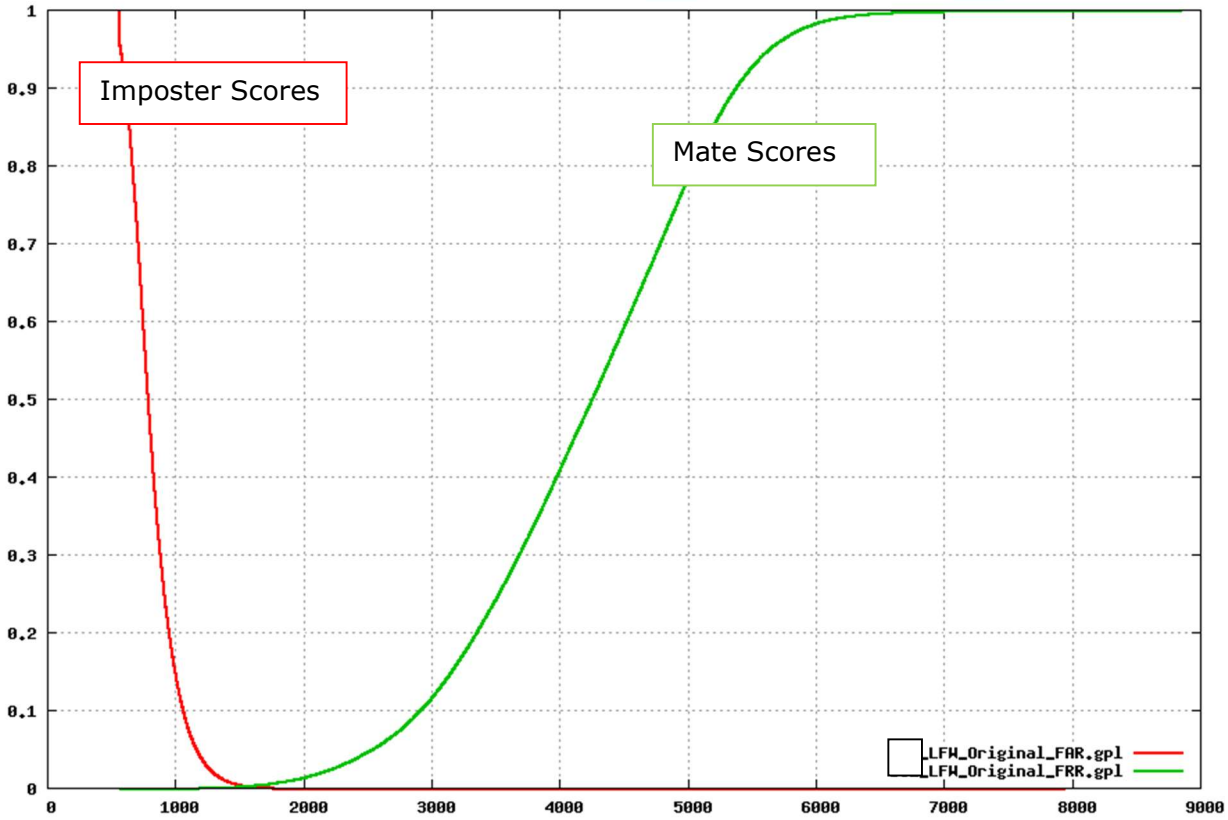

FIG. 1 FAR and FRR from uncorrected imagery

In Figure 1 the FAR and FRR scores are presented on a linear Y-axis. This can tend to hide the identity errors.

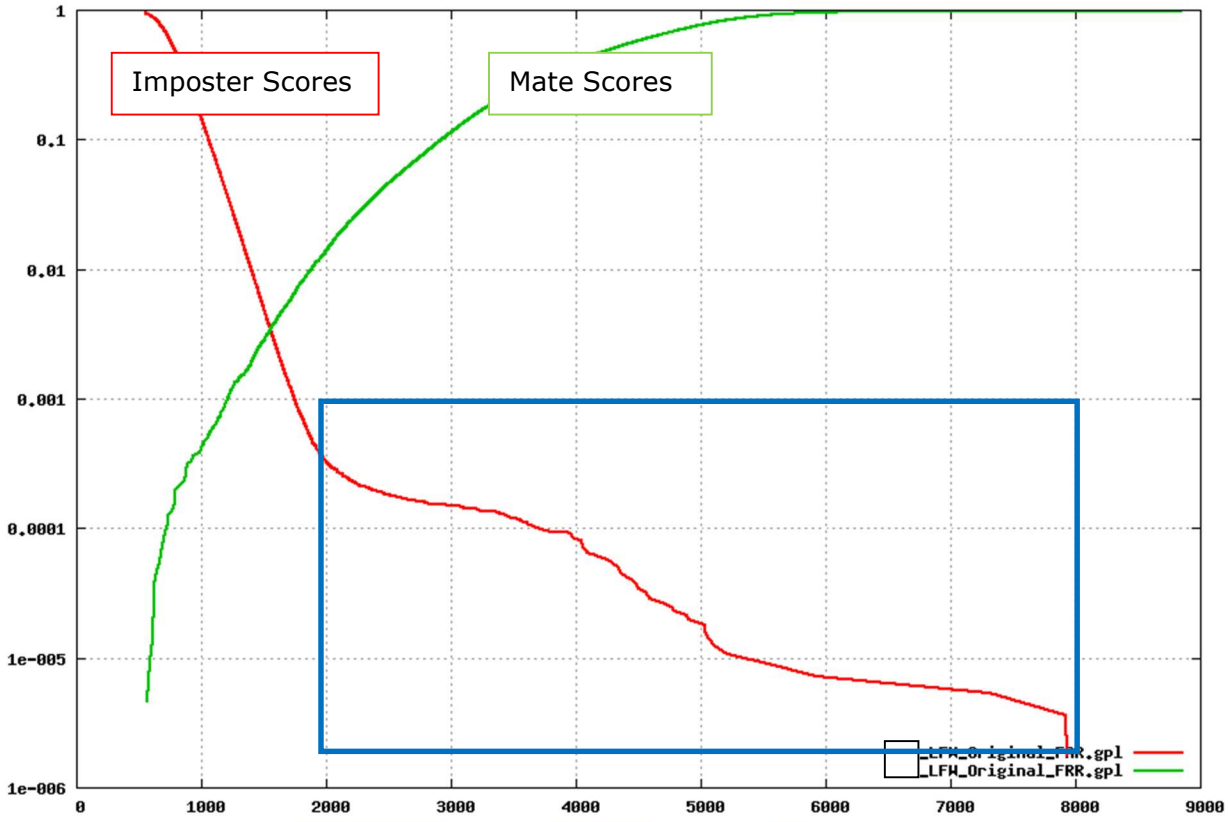

FIG. 2 FAR and FRR from uncorrected imagery

In Figure 2 the FAR and FRR scores are presented on a logarithmic Y-axis. This shows that there may be identity errors in the high scoring FAR scores (blue area). The sudden increase in FAR scores indicates potential mates that are incorrectly labeled imposters.

Other methods to select incorrectly labeled imposters are publicly available. See Appendix 1: "**Alternative Methods**".

Step 4: Analyze the scoring to resolve high scoring imposters to see if identity errors do exist in the data.

Step 5: Iterate between Steps 1-4 as long as identity errors are suspected. The expectation is that several passes will be needed to achieve an objective measure of correct ground truth.

NOTE: In the LFW data set several iterations were done and approximately 90 identities were modified. After these corrections were made the following accuracy plots were derived.

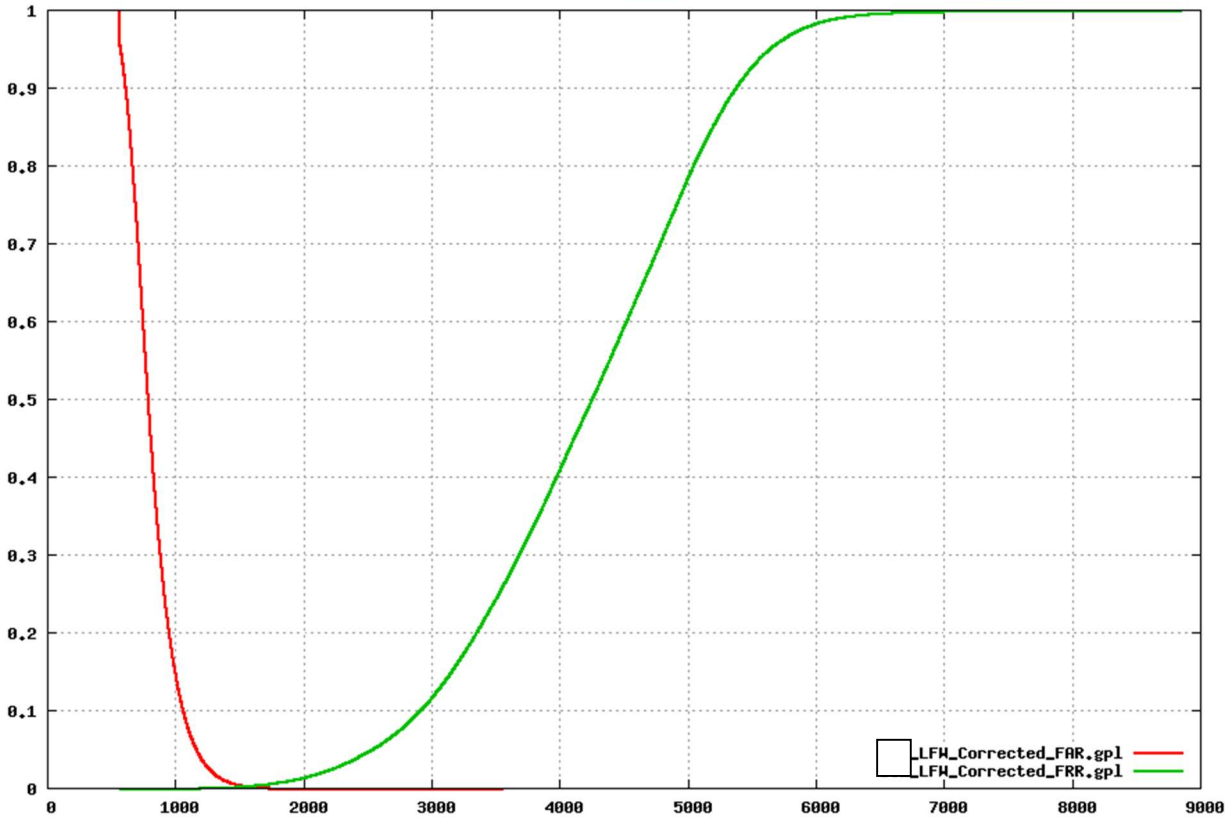

FIG. 3 FAR and FRR from corrected imagery

In Figure 3 the FAR and FRR scores are presented on a linear Y-axis. This can tend to hide the identity errors.

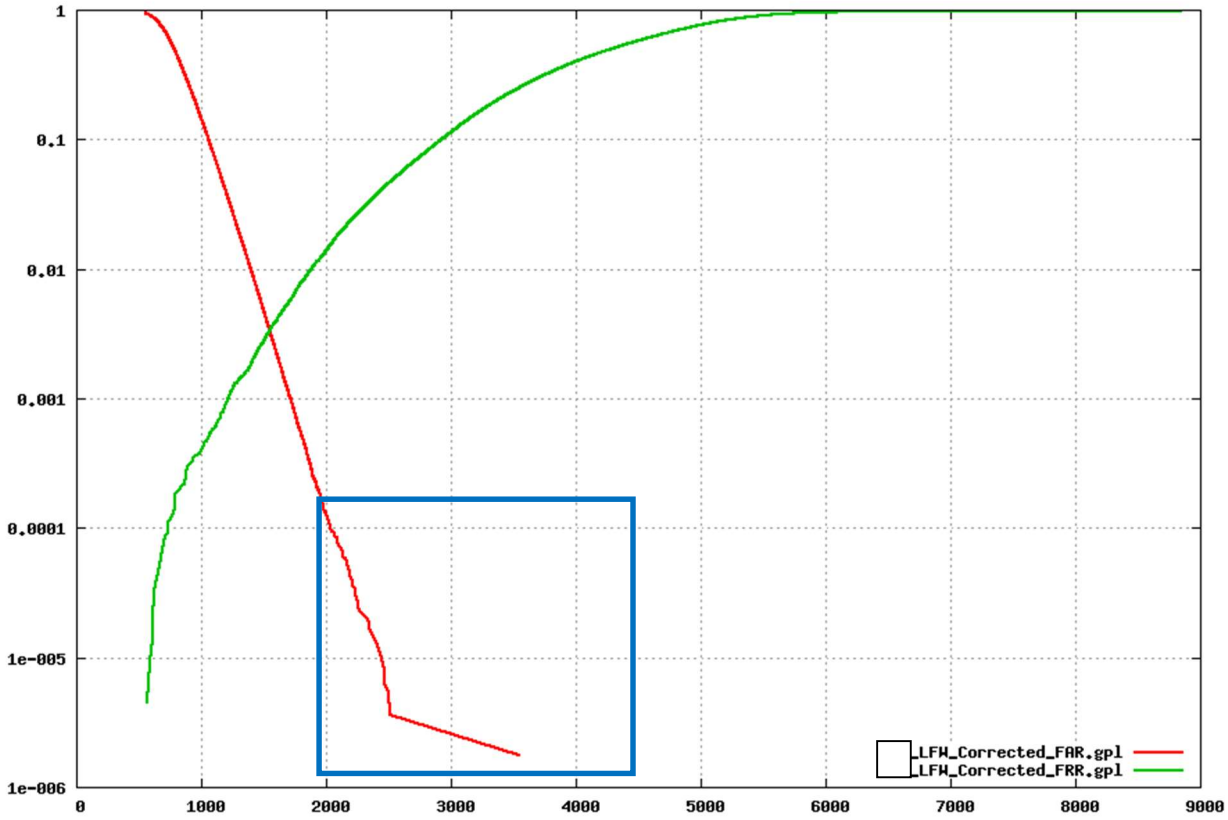

FIG. 4 FAR and FRR from uncorrected imagery

In Figure 4 the FAR and FRR scores are presented on a logarithmic Y-axis. This shows there may be several more identity errors in the FAR scoring (blue area).

155 One can find the following images by investigating the high score imposters.

Twins	
	
	

156
157

158


The images above are siblings


The images above are close appearances


The images above are close appearances

As these high scoring imposters were investigated, it became apparent that the high scores in the corrected FAR were heavily influenced by:

- Twins
- Siblings
- What are referred to as “doppelgangers”

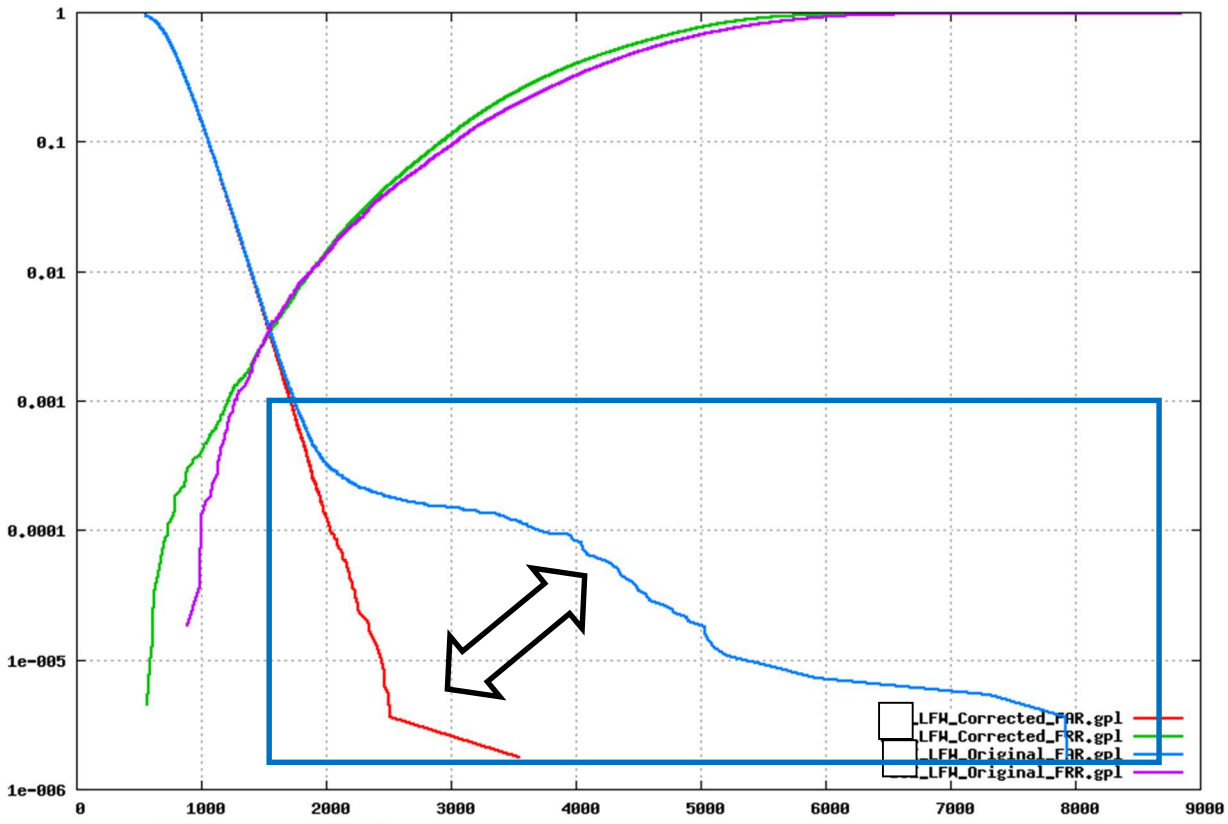

FIG. 5 FAR and FRR from original and corrected LFW imagery

In Figure 5 the FAR and FRR from original vs. corrected imagery are presented on a logarithmic Y-axis. This shows the scoring variances after the identity errors were corrected (blue area). This shows how a few improper identities can dramatically affect the FAR and FRR scoring profiles.

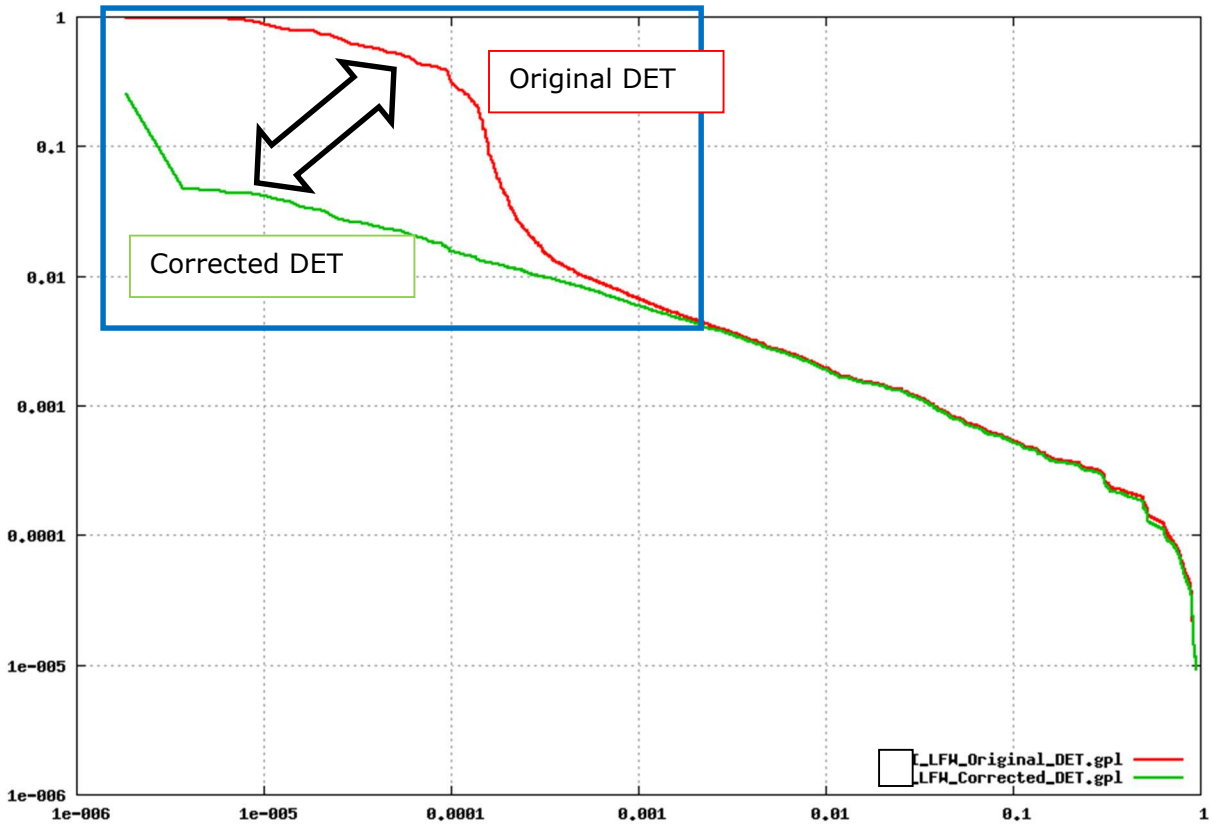

FIG 6 DET comparison between original and corrected LFW imagery

In Figure 6 the DET curve from original vs. corrected imagery are presented. This shows the scoring variances after the identity errors were corrected (blue area). This again shows how a few improper identities can dramatically affect the FAR and FRR scoring profiles.

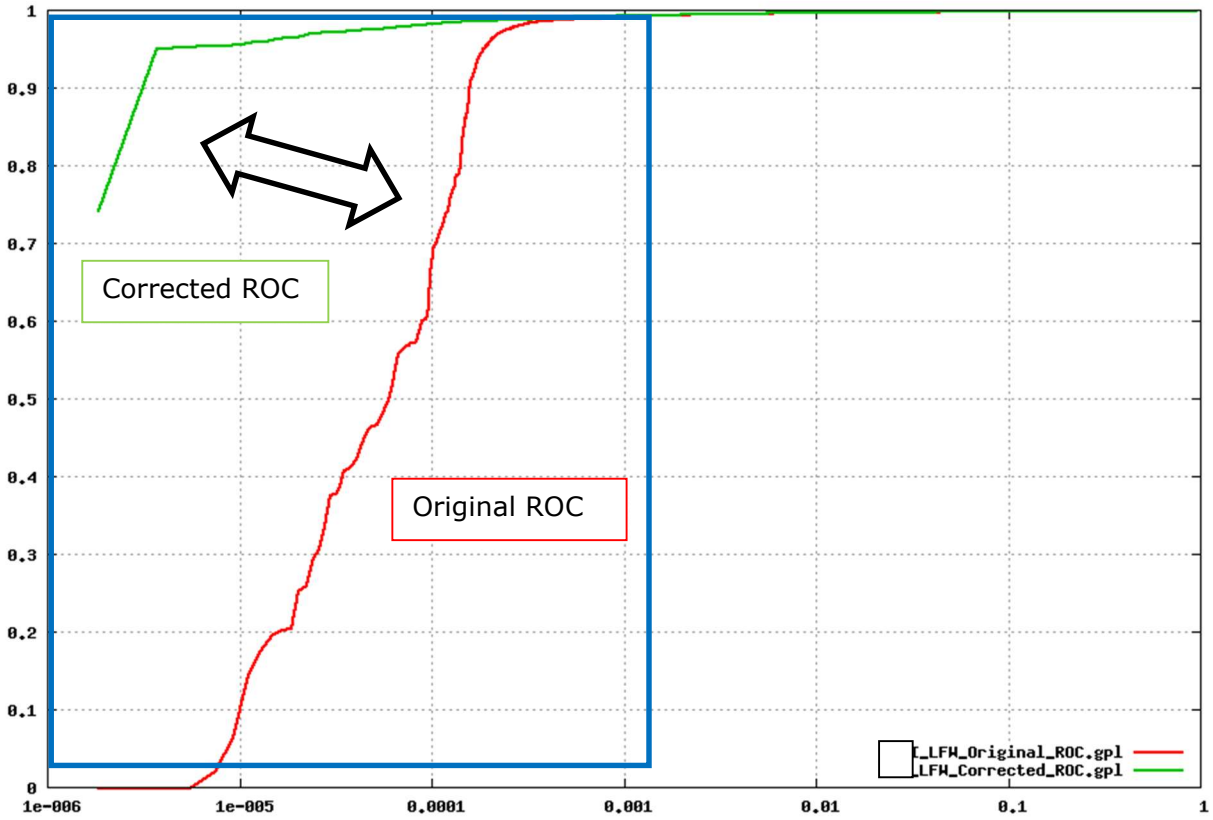

FIG. 7 ROC comparison between original and corrected imagery

In Figure 7 the ROC curve from original vs. corrected imagery are presented with a logarithmic X-axis. This shows the scoring variances after the identity errors were corrected (blue area).

184 Investigating low scoring mates - the following images were confirmed identities but
185 scored very low.

Low Scoring Mates

Outcomes

Based on this data set and the testing process documented here:

- The LFW data set tested had identity errors which need to be adjusted to ensure proper scoring analysis could be done. From the default data set downloaded ~90 identities should have been corrected.
- Iterative processes to properly analyze and modify this data set were required which focused on using the FAR/FRR curves to locate the identity errors.

- Correcting the identity errors improved the FAR/FRR scoring analysis
- Twins, siblings, and doppelgangers did affect the scoring analysis. Two twins should have been located with several siblings and doppelgangers causing the highest FAR scores.
- FAR, FRR, DET and CMC curves were utilized in these processes.
- Critical scoring analysis required the presentation of the scoring in both linear and logarithmic presentations to see the imposter scores which had identity errors.
- A variety of methods can be used to resolve identity errors in facial datasets. FRS administrators should be aware of the advantages and disadvantages of each method before selecting and applying a method, especially on operational datasets.

Glossary

- FR: face recognition
- FRS: face recognition systems
- ROC: Receiver Operating Characteristics
- DET: Detection error tradeoff
- FAR (False Accept Rate): the measure of the probability that the biometric system will incorrectly accept an access attempt by an unauthorized user
- FRR (False Reject Rate): the measure of the probability that the biometric system will incorrectly reject an access attempt by an authorized user
- Doppelganger: an apparition or double of a living person.

Appendix 1: Alternative Methods

General/Initial assessment

The initial assessment objectives are twofold: first, to acquire a general sense of how the FRS interacts with the data, and second, to select the identities or images for detailed review.

The initial assessment assumes that similarity scores, both genuine and impostor, are distributed according to predictable and smooth monomodal curves and that errors in the dataset are expected to introduce "unnatural" multimodal variations in the curves obtained from the sets of scores.

Because the specific shape of this curve (e.g., Gaussian, exponential, log-normal) varies from system to system and also depends on the dataset itself, it is important to develop a general sense of the interaction between the FRS and the data.

It is important to note that the effectiveness of these procedures, which are ultimately based on the distribution of scores, is fundamentally dependent on the accuracy of algorithm used to generate these scores. Algorithms whose scores are highly discriminative are more effective in identifying outliers.

An initial assessment could then be based on curves derived from genuine and impostor scores distributions as, e.g., the FAR/FRR curve, or directly inspecting both distributions of scores.

232 FAR/FRR curve plots both of these rates as a function of threshold. An example of

233 such a curve for the LFW dataset is shown below.

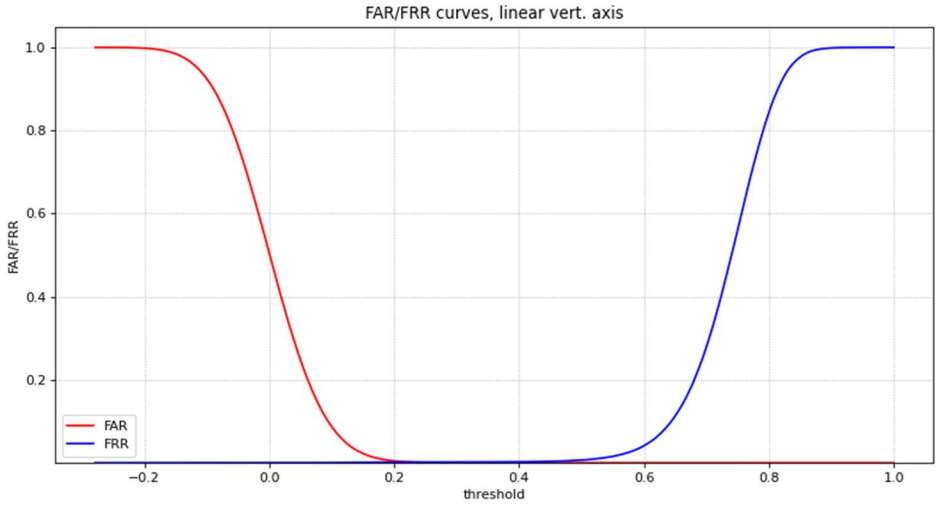

234
235 **FIG. A1.1 FAR and FRR from uncorrected imagery**

236 In this plot, the vertical axis is linear, but plotting the same data with the vertical axis

237 in a logarithmic scale facilitates the inspection of the curves in the lowest values.

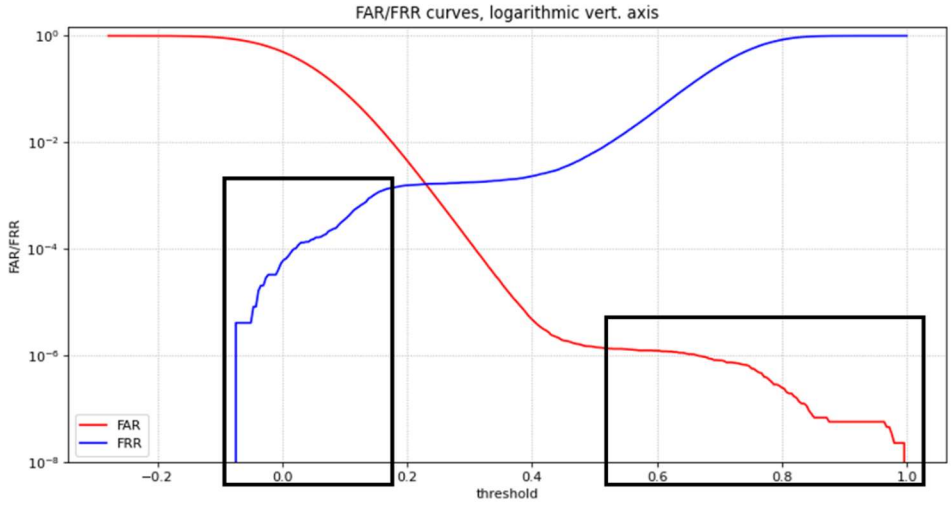

238
239 **FIG. A1.2 FAR and FRR from uncorrected imagery (Logr axis)**

In this plot, it is now possible to observe variations on the shape of the FAR and FRR curves that violate the assumption of smoothness of a monomodal distribution curve (framed portions). This raises suspicions that errors exist in the dataset, resulting in more impostor scores with high values and more genuine scores with low values.

This assessment can be further explored using the distributions of scores directly. The plot below is based on the same data.

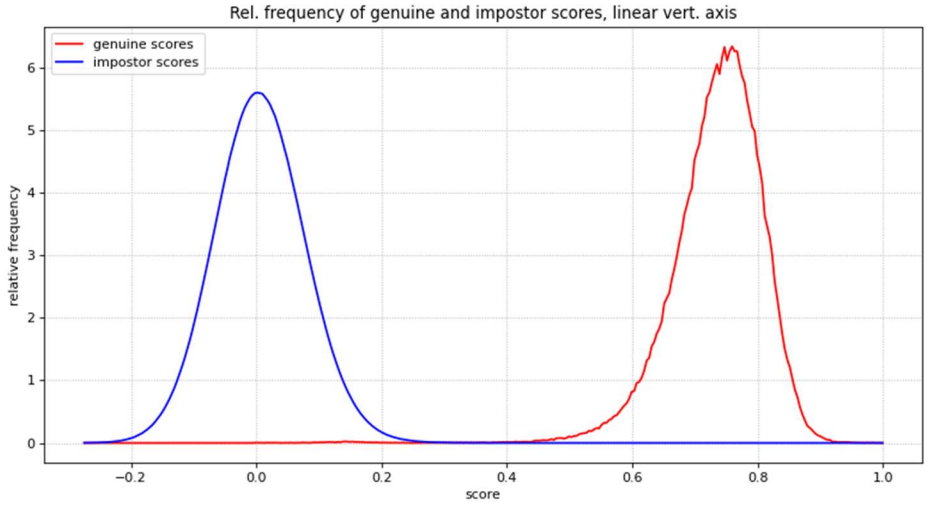

FIG. A1.3 Genuine and Imposter Scoring Histograms uncorrected imagery

The vertical axis's linear scale again makes it challenging to inspect the lower values, but using a logarithmic vertical axis allows one to observe a violation of the assumption of smoothness of a monomodal distribution curve (framed portions).

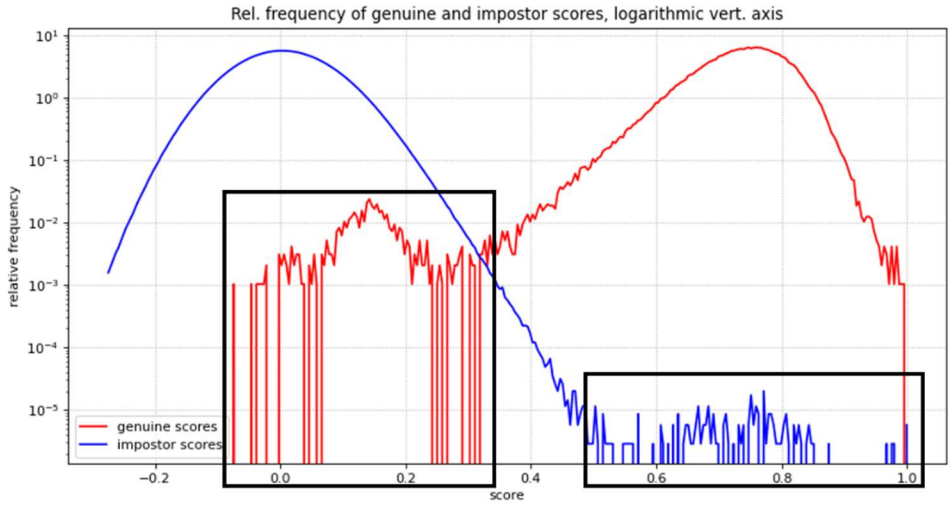

FIG. A1.4 Genuine and Imposter Scoring Histograms uncorrected imagery (Logr axis)

Both the higher impostor scores and the lower genuine scores could be selected from this plot to a detailed review of the images involved. Impostor scores should be selected for further inspection in highest to lowest order, and the opposite for genuine scores.

Apart from the visual assessment of the curves derived from the distribution of scores, statistical methods for identifying outliers in each set of scores can be used.

One such method involves selecting outliers based on multiples of the median absolute deviation from median (MAD) of each set of scores. In univariate statistics, the MAD is a robust dispersion measure in the presence of outliers, contrary to the more commonly used standard deviation around the mean. [1]

[1] Leys, C., et al., Detecting outliers: Do not use standard deviation around the mean, use absolute deviation around the median, Journal of Experimental Social Psychology (2013), http://dx.doi.org/10.1016/j.jesp.2013.03.013,

In practical terms, for genuine scores, those scores that lie below a multiple of MAD from the median should be considered as outliers and then manually verified. For impostor scores, the outliers are those higher than a multiple of MAD from the median. Larger multiples of MAD will result in smaller sets of outliers to be manually verified, while smaller multiples of MAD will result in a larger number of scores being considered as outliers.

It is essential to notice that both the MAD and median statistics should be calculated independently for each set of scores, genuine and impostors.

The specific multiple of MAD for each set should be determined experimentally, on a case by case basis. For the worked example, it was determined through a controlled experiment with the algorithm and the LFW dataset.

For this specific purpose, controlled errors were introduced in the dataset, which added 764 errors in the genuine scores and the same quantity in the impostor scores. These experimental sets were verified using the criteria of multiples of MAD, which resulted in the number of six MADs as being adequate to identify the controlled errors as outliers in both genuine and impostor scores.

The remainder of the analysis was conducted in the original LFW dataset, without the artificially introduced errors.

The figure below illustrates the procedure for selecting outliers in the impostor scores, in the original LFW dataset, after the appropriate multiple of MAD was determined. The procedure for genuine scores is analogous.

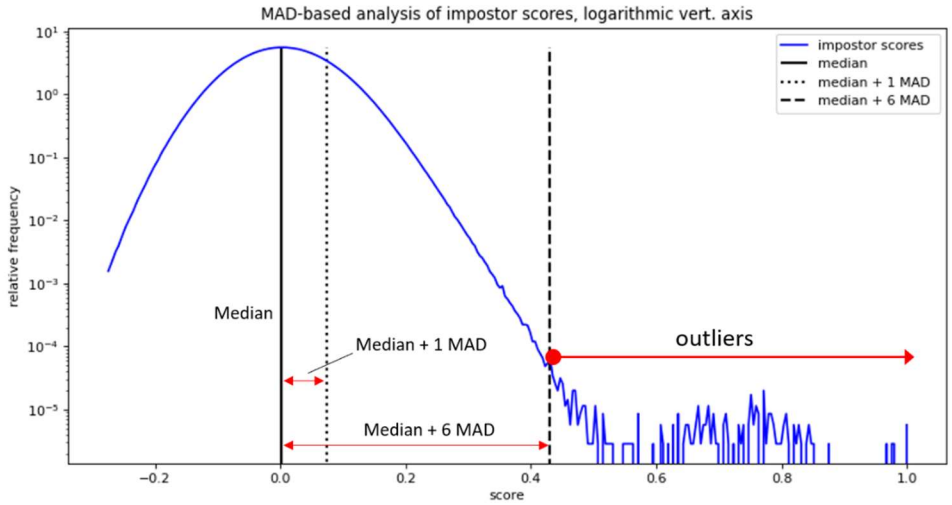

FIG. A1.5 MAD Based Imposter Scores uncorrected imagery (Logr axis)

In the case of the original LFW dataset, without prior corrections of the known errors, the number of six MADs for the genuine and impostor scores sets resulted in the selection of 474 and 220 image pairs as outliers in each set, respectively.

The amount of detailed inspection is subject to the available workforce. Ideally, it should be done until all image pairs identified as outliers are reviewed, and the FRS administrator is assured of the dataset integrity and correctness.

Specific Assessments

The general/initial assessment will result in two sets of image pairs (allegedly genuine and allegedly impostors) that must be manually reviewed to check for identity or other kinds of errors.

Depending on the number of image pairs, a direct assessment of all pairs can be a viable approach. However, if the number of pairs is relatively large, some strategies can be employed to optimize this reviewing process.

One such strategy is to verify if some images appear more frequently in each of the sets selected for review. These images should be inspected first.

Results

The remaining pairs of images were inspected, and all the known errors published on the LFW website were identified.

Some of the high impostor scores were confirmed to be from different persons. In most cases, this was caused by the presence of twins, siblings, or close appearances.

For the low genuine scores, some were verified as being the same person, but variations in pose, make-up, facial expression, and age were factors the negatively affected the scores.

After correcting the errors, the FAR/FRR and the scores distributions plots show a smoother monomodal behavior.

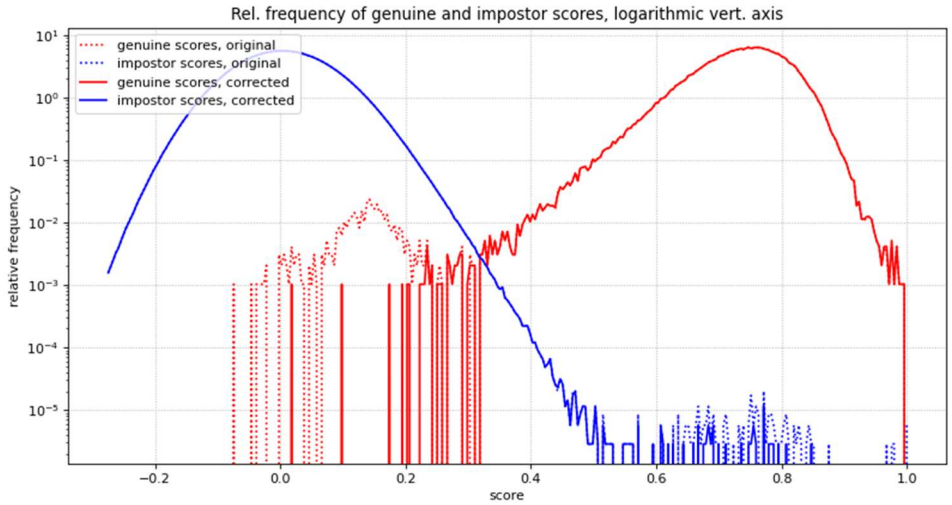

311
312 **FIG. A1.6 Genuine and Imposter Scoring Histograms corrected imagery (Logr axis)**

313 The distributions of scores improved noticeably, with fewer outliers both in the
314 genuine and impostor sets after the corrections were applied.

315 This improvement can also be observed in the FAR/FRR plot, which is directly related
316 to the scores' distributions. After the corrections, EER improved by an order of
317 magnitude.

318

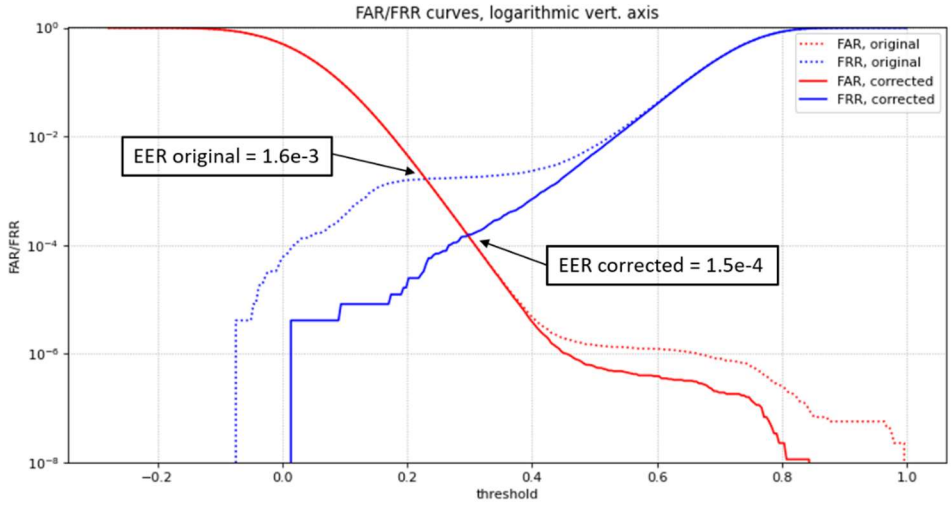

FIG. A1.7 FAR and FRR from uncorrected/corrected imagery (Logr axis)

Finally, the DET plot also reveals improvement in all operational points.

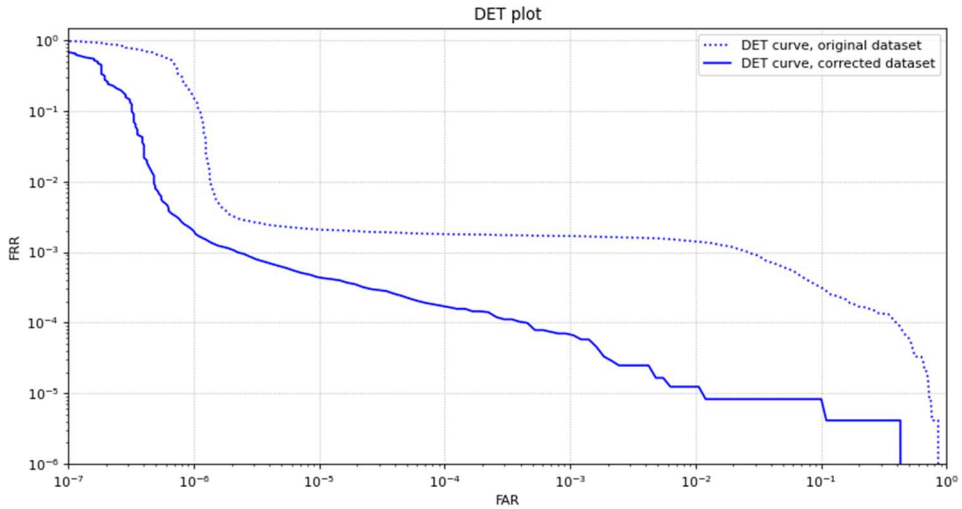

FIG. A1.8 DET from uncorrected/corrected imagery

The whole process can be repeated until the FRS administrator is assured of the dataset integrity and correctness or while there are resources available.

Other Methods

The verification of integrity and correctness of biometric datasets is an evolving area in the scientific literature. In this section, some methods and approaches that can be used to tackle this problem are referenced.

The chapter on Exploratory Data Analysis (EDA) of the NIST/SEMATECH e-Handbook of Statistical Methods [2] is recommended. The book describes EDA as "(…) an approach/philosophy for data analysis that employs a variety of techniques (mostly graphical) to (i) maximize insight into a data set; (…) (iv) detect outliers and anomalies; (…)". Although not specific to inspecting facial datasets, many of the concepts presented in this chapter can be used for this purpose.

In [3], the authors describe a method specifically designed to sort out identity errors in large facial datasets, similar to the MAD analysis described in the worked example. Their method is based on a two-layered thresholding process to select outliers. First, identity outliers are selected, and, for each selected identity, images considered as outliers are manually reviewed.

[2] NIST/SEMATECH e-Handbook of Statistical Methods, https://doi.org/10.18434/M32189

[3] V. Varkarakis and P. Corcoran, "Dataset Cleaning — A Cross Validation Methodology for Large Facial Datasets using Face Recognition," 2020 Twelfth International Conference on Quality of Multimedia Experience (QoMEX), Athlone, Ireland, 2020, pp. 1-6, doi: 10.1109/QoMEX48832.2020.9123123.

In [4], the authors present the concept of Biometric Menagerie and propose a framework for evaluating biometric systems. This framework is based on the relationship between a person's genuine and impostor scores and could be explored to select identities for further inspection.

Apart from methods that aid in selecting identities or images for manual review, some methods are proposed in the literature to automatically clean large facial datasets, with a large number of images assigned to each identity.

In one such approach, presented in [5], a clustering algorithm, Density-Based Spatial Clustering of Applications with Noise – DBSCAN, clusters similar images in each identity set, retaining the cluster with most images.

[4] Yager, Neil and Ted Dunstone. "The Biometric Menagerie." IEEE Transactions on Pattern Analysis and Machine Intelligence 32 (2010): 220-230.

[5] Gallo, Ignazio & Nawaz, Shah & Calefati, Alessandro & Piccoli, Gabriele & Zamberletti, Alessandro. (2018). A Pipeline to Improve Face Recognition Datasets and Applications. 10.1109/IVCNZ.2018.8634724.

Another automatic method [6] employs the community detection algorithm to identify and delete mislabeled images while preserving diversity in each identity's images.

The reader should be aware that these fully automated methods will remove some images from the dataset without human reviewing.

Although there is a limited number of commercially available software specialized for this task, most scientific literature methods can be implemented in software devoted to statistics and mathematics.

[6] Chi Jin, Ruochun Jin, Kai Chen, Yong Dou, "A Community Detection Approach to Cleaning Extremely Large Face Database", Computational Intelligence and Neuroscience, vol. 2018, Article ID 4512473, 10 pages, 2018. https://doi.org/10.1155/2018/4512473

Appendix 2: LFW Data Set Information

This is widely used open source data set which will work well for this specific document. Information on this data set includes:

DISCLAIMER:

Labeled Faces in the Wild is a public benchmark for face verification, also known as pair matching. No matter what the performance of an algorithm on LFW, it should not be used to conclude that an algorithm is suitable for any commercial purpose. There are many reasons for this. Here is a non-exhaustive list:

- Face verification and other forms of face recognition are very different problems. For example, it is very difficult to extrapolate from performance on verification to performance on 1:N recognition.
- Many groups are not well represented in LFW. For example, there are very few children, no babies, very few people over the age of 80, and a relatively small proportion of women. In addition, many ethnicities have very minor representation or none at all.
- While theoretically LFW could be used to assess performance for certain subgroups, the database was not designed to have enough data for strong statistical conclusions about subgroups. Simply put, LFW is not large enough

to provide evidence that a particular piece of software has been thoroughly tested.

- Additional conditions, such as poor lighting, extreme pose, strong occlusions, low resolution, and other important factors do not constitute a major part of LFW. These are important areas of evaluation, especially for algorithms designed to recognize images “in the wild”.

For all of these reasons, we would like to emphasize that LFW was published to help the research community make advances in face verification, not to provide a thorough vetting of commercial algorithms before deployment.

While there are many resources available for assessing face recognition algorithms, such as the Face Recognition Vendor Tests run by the USA National Institute of Standards and Technology (NIST), the understanding of how to best test face recognition algorithms for commercial use is a rapidly evolving area. Some of us are actively involved in developing these new standards and will continue to make them publicly available when they are ready.

Welcome to Labeled Faces in the Wild, a database of face photographs designed for studying the problem of unconstrained face recognition. The data set contains more than 13,000 images of faces collected from the web. Each face has been labeled with the name of the person pictured. 1680 of the people pictured have two or more

distinct photos in the data set. The only constraint on these faces is that they were detected by the Viola-Jones face detector. More details can be found in the technical report below.

Information:

- 13233 images
- 5749 people
- 1680 people with two or more images

Reference:

Please cite as:

Gary B. Huang, Manu Ramesh, Tamara Berg, and Erik Learned-Miller.

Labeled Faces in the Wild: A Database for Studying Face Recognition in Unconstrained Environments.

University of Massachusetts, Amherst, Technical Report 07-49, October, 2007.

[pdf]

BibTeX entry:

```
@TechReport{LFWTech,
  author =      {Gary B. Huang and Manu Ramesh and Tamara Berg and
                 Erik Learned-Miller},
```

title = {Labeled Faces in the Wild: A Database for Studying Face Recognition in Unconstrained Environments},

institution = {University of Massachusetts, Amherst},

year = 2007,

number = {07-49},

month = {October}}

Gary B. Huang and Erik Learned-Miller.

Labeled Faces in the Wild: Updates and New Reporting Procedures.

University of Massachusetts, Amherst, Technical Report UM-CS-2014-003, May, 2014. [pdf]

@TechReport{LFWTechUpdate,

author = {Gary B. Huang Erik Learned-Miller},

title = {Labeled Faces in the Wild: Updates and New Reporting Procedures},

institution = {University of Massachusetts, Amherst},

year = 2014,

number = {UM-CS-2014-003},

month = {May}}

367

368 Critical to this document's purpose is the errata found on the LFW URL:

> **Errata**:
>
> The following is a list of known errors in LFW. Due to the small number of such errors, the database will be left as is (without corrections) to avoid confusion.
>
> It is important that users of the database provide their algorithms with the database as is, i.e. without correcting the errors below, since previous results published for the database did not have the advantage of correcting for these errors.
>
> Currently, there are six incorrectly labeled matched pairs in View 2. While we do not believe this should have a significant effect on accuracy, we do encourage researchers to be aware of these errors when producing any visualizations (e.g. matched pairs most confidently predicted as mismatched, as the matched pair may actually be mismatched).

369 A list of errors can be viewed at: http://vis-www.cs.umass.edu/lfw/

Reference List

[1] ANSI/NIST-ITL Standard Homepage:

http://www.nist.gov/itl/iad/ig/ansi_standard.cfm

[2] P. Grother, M. Ngan, K. Hanaoka “NISTIR 8271 DRAFT SUPPLEMENT Face Recognition Vendor Test (FRVT) Part 2: Identification”

https://pages.nist.gov/frvt/reports/1N/frvt_1N_report.pdf

FISWG documents can be found at: www.FISWG.org