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 Facial Recognition Systems Operation Assurance

1. Scope

1.1. This document provides guidelines and techniques to help administrators of automated facial recognition systems (FRS) test them in an operational setting to provide assurance of facial recognition accuracy, system integrity, configuration, and data storage.

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

2. Referenced Documents

2.1. ISO/IEC (International Organization for Standardization/International Electrotechnical Commission). 2012. ISO/IEC 19795-6:20121 — Information

[1] For the referenced ANSI standard, visit the ANSI webstore at http://webstore.ansi.org/RecordDetail.aspx?sku=ISO/IEC+19795-6:2012.

technology — Biometric performance testing and reporting — Part 6: Testing methodologies for operational evaluation.

2.2. NIST/ITL (National Institute of Standards and Technology/ Information Technology Laboratory). 2011. NIST Special Database 32–Multiple Encounter Dataset (MEDS)[2].

2.3. NIST publication "Relating ROC and CMC Curves[3]," 2012, by Arun Ross.

2.4. "NIST Interagency Report 8271 DRAFT SUPPLEMENT Face Recognition Vendor Test (FRVT) Part 2: Identification,[4]" by P. Grother, M. Ngan, K. Hanaoka.

3. Terminology

3.1. Definitions:

3.1.1. *FR, acronym* for facial recognition.

3.1.2. *FRS, acronym* for facial recognition systems.

3.1.3. *CMC, acronym* for cumulative match characteristic.

3.1.4. *ROC*, *acronym* for Receiver Operating Characteristics.

3.1.5. *Exact binary copy, n* - two images that are byte-identical, i.e., they yield the same MD5 checksum. Note: An image encoded as a PNG is not an exact binary copy of the same image encoded as a JPG.

3.1.6. *Mate, n* - a separately captured image of the same subject

[2] For the referenced NIST document visit http://www.nist.gov/itl/iad/ig/sd32.cfm.

3 For the referenced NIST document visit https://www.nist.gov/system/files/documents/2016/12/06/12_ross_cmc-roc_ibpc2016.pdf.

4 For the referenced NIST document visit https://pages.nist.gov/frvt/reports/1N/frvt_1N_report.pdf,

3.1.7. *Retrieval rank, n* - the order with which a particular image in the gallery was retrieved respective to a probe and all other images in the gallery

4. Related Work

4.1. The overarching goal of ISO/IEC 19795-6:2012 — Information technology — Biometric performance testing and reporting — Part 6: Testing methodologies for operational evaluation[5] "to measure or monitor operational biometric system performance" parses into sub-goals of:

4.1.1. Determining if performance meets expectations or may be improved through system tuning or reconfiguring;

4.1.2. Predicting expected performance for increases in number of enrollments and/or systems;

4.1.3. Obtaining information that affects system performance (e.g., changes in target population and environmental parameters); and/or

4.1.4. Obtaining performance data from a pilot implementation or to benchmark future systems.

4.2. This document serves as a qualitative introduction to a subset of the topics and concepts that are covered in greater technical detail in ISO/IEC 19795-6:2012.

[5] *ISO/IEC (International Organization for Standardization/International Electrotechnical Commission). 2012. ISO/IEC 19795-6:2012 — Information technology — Biometric performance testing and reporting — Part 6: Testing methodologies for operational evaluation*. http://webstore.ansi.org/RecordDetail.aspx?sku=ISO/IEC+19795-6:2012.

5. Operational Testing Techniques

5.1. The operational testing techniques provided in this document can be performed manually, automatically, or semi-automatically, depending on operational factors and the administrator's technical knowledge.

5.2. Note: System administrators familiar with performance tests run by the National Institute of Standards and Technology (NIST) Information Technology Laboratory (ITL) like the Face Recognition Vendor Tests (FRVT) and the Multiple Biometrics Grand Challenge (MBGC) should be aware that biometric test results may not be applicable to other datasets and operational systems with different processing constraints. These test results should be reviewed within the context in which they are obtained and might not generalize to operational data.

6. Tier-based System Evaluation Guide

6.1. A tier-based approach to operational system evaluation is proposed which provides a tradeoff between the difficulty of implementing the testing strategy and the level of system assurance that can be achieved through such testing. Figure 1 illustrates the tier-based testing paradigm.

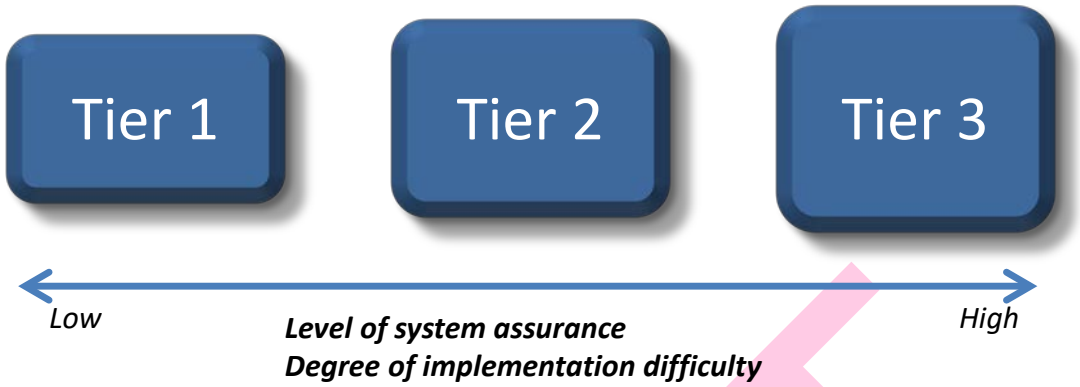

Figure 1. Recommendations for performing an operational evaluation of a facial recognition system are tiered, which allows administrators a tradeoff between the degree of testing difficulty and the level of system assurance

6.2. **Tier 1 – Basic:**

6.2.1. The simplest level of testing is Tier 1 testing, or self-identification, which is a basic check to ensure there are not egregious errors in the system. Tier 1 testing can be performed by the system administrator, but may also be performed by system users in certain circumstances.

6.2.2. Tier 1 testing is performed as follows:

6.2.2.1. Query the facial recognition system with images whose exact binary copies have already been enrolled in the database.

6.2.2.2. Ensure the Rank-1 match candidate is the same image as the probe.

6.2.2.3. Continue to perform the first two steps with additional images as often as availability of computing resources and human effort allow.

6.2.3. Tier 1 testing is performed to ensure:

6.2.3.1. Gallery images are properly enrolled and their corresponding

templates are both valid and accessible.

6.2.3.2. The Facial recognition (FR) system's network access is not interrupted for any distributed resources.

6.2.3.3. Software running on local and network resources is not exhibiting any failures.

6.2.4. Tier 1 testing provides minimal assurance regarding the recognition accuracy of the facial recognition system. It provides a basic level of assurance of the system integrity, configuration, and data storage.

6.3. **Tier 2 – Intermediate:**

6.3.1. Tier 2 testing provides intermediate retrieval accuracy statistics on the facial recognition system using test subjects whose mates are known to be in the gallery. Depending on operational factors and the administrator's technical knowledge, Tier 2 testing may be performed by the system administrator and/or the FR algorithm vendor/integrator.

6.3.2. Tier 2 testing must be undertaken using the operating or vendor recommended default threshold in order that the results are truly representative of expected operational performance.

6.3.3. Testing should be undertaken with operationally representative images. If the query images are from a number of different sources (e.g. passport, CCTV, social media) it may be necessary to categorize the images into different query sets based on these sources. It may be acceptable to adjust the threshold in order to determine the retrieval rate

for lower quality images. However, these tests should not be run with a 0 threshold.

6.3.4. **Tier 2 testing is performed as follows**:

6.3.4.1. Query system with images whose mates are known to be in the system. The query image should be from a separate encounter than at least one gallery mate. Query images should be representative of operational distributions. For example, for a database of 10,000 images comprised of 80% white males and 20% black females, the system administrator might query images of 80 white males and images of 20 black females.

6.3.4.2. Record the top rank in which the probe's mate was retrieved. If an image from same encounter as probe image is contained in the gallery (e.g., the exact binary copy), then do not consider same encounter candidate list in the retrieval results.

6.3.4.3. Continue to perform the first two steps with additional images as often as availability of computing resources and human effort allow.

6.3.4.4. Use recorded rank retrievals from all retrieval tests to generate Cumulative Match Characteristic (CMC) accuracies.

6.3.5. When testing on galleries that continually increase in size, the recorded accuracies for this test generally will decrease over time. However, for a properly functioning recognition system no major negative changes in

retrieval accuracy should occur for major system updates or high frequency fixed-interval tests.

6.3.6. How to generate CMC scores:

6.3.6.1. The CMC scores list what percentage of image queries had their mate returned at a particular retrieval rank or better. The CMC scores typically would store accuracies up to the first N ranks. The value of N would be determined based on how many retrievals are typically examined in the system's operational use. Thus, for an application where analysts examine the top 20 matches, N=20.

6.3.6.2. The CMC scores contain N values, which correspond to the Rank-1 accuracy, the Rank-2 accuracy, and all the way up to the Rank-N accuracy. These accuracies are interpreted as follows. The Rank-1 accuracy is what percentage of queries had their mate at the first retrieval rank (i.e, the closest match). The Rank-2 accuracy is what percentage of queries had the mate returned at the second retrieval rank, or better. Similarly, the Rank-3 accuracy is the percentage of queries that had their mate retrieved at the third rank, or better.

6.3.6.3. It is important to note the "or better" portion of the above description. If an image is matched at Rank-1, then it is also included in the Rank-2 (and beyond) CMC scores. Thus, when plotting the CMC scores respective the retrieval ranks, the

graph will not decrease.

6.3.6.4. When computing CMC scores in operational testing, the user would submit an image query and record the highest rank pertaining to a match. This process would be continued k times, where (as discussed above) k is determined based on available resources.

6.3.6.5. As an example, suppose k = 50 (i.e., 50 different images submitted to the system).For N = 7 ranks, the test may yield the results in Table 1:

Rank	**1**	**2**	**3**	**4**	**5**	**6**	**7**
Number at Rank	32	7	1	3	1	0	1

Table 1. Example of Candidates for Rank Order 1 - 7

The remaining 5 of the 50 submissions do not return within the designated threshold of 7 candidates, thus are not reported in the table. Given these results, the CMC scores are as shown in Table 2:

Rank	**1**	**2**	**3**	**4**	**5**	**6**	**7**
CMC score	32/50 =0.64	39/50 =0.78	40/50 =0.8	43/50 =0.86	44/50 =0.88	44/50 =0.88	45/50 =0.9

Table 2. Example of CMC Scores for Candidates for Rank Order 1 - 7

Thus, when performing Tier 2 testing, these CMC scores would be logged each time the test was performed. With a fixed set of query images, one should not expect major deviations from these CMC scores.

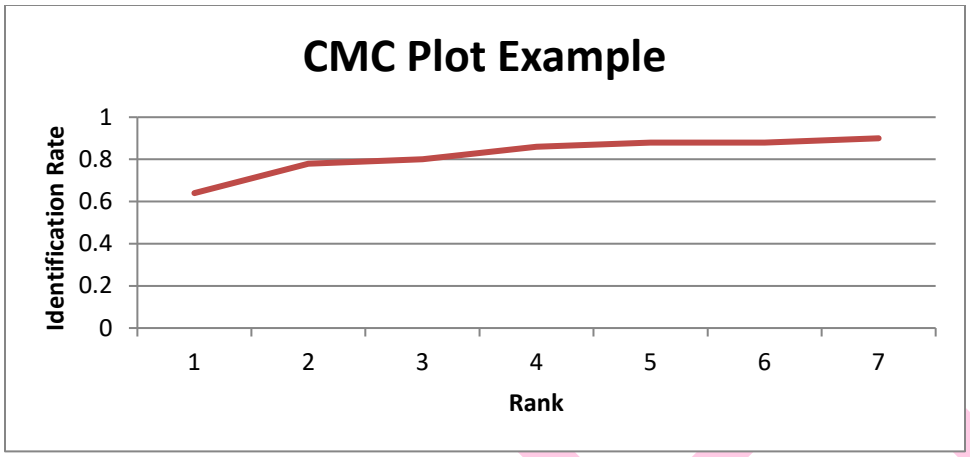

Figure 2. Example CMC Plot

Figure 2 is a plot of the CMC curve of the same example previously discussed. One can visualize certain aspects of CMC curves, and how they can be interpreted. One key point is that the Identification Rate is never decreasing with increasing Rank. When reading this plot, it can be interpreted as follows: at Rank=3, the rate is 0.8. Thus, in the test example, 80% of the time the subjects are matched within the top three ranks. Similarly, at Rank=7, the identification rate is 0.9. Thus, in the test example, 90% of the subjects are matched within the top seven ranks. Finally, because only results of the top seven ranks are recorded, the CMC curve[6] never reaches 100%

[6] For more detail about creating a CMC curve, readers should refer to the NIST publication "Relating ROC and CMC Curves ," 2012, by Arun Ross at https://www.nist.gov/system/files/documents/2016/12/06/12_ross_cmc-roc_ibpc2016.pdf.

accuracy. Measuring results up to a higher rank, such as 50, may or may not allow one to record an identification rate of 1.0.

6.4. **Tier 3 – Advanced:**

6.4.1. Tier 3 testing provides advanced recognition accuracy statistics on the facial recognition systems. This section introduces and briefly discusses the Receiver Operating Characteristics (ROC) measurement and how to interpret this result. For more in depth discussions on how to generate and interpret advanced results such as ROC, readers are referred to other documents that provide these details [3][4].

6.4.2. Generally, these accuracy measurements cannot be computed by the FR system administrator. Instead, these measurements will often be made available through software provided by the FR algorithm vendor/integrator. Additionally, these results can be computed in non-operational scenarios and offline environments using sufficient and representative ground truth operational data. For example, the National Institute of Standards and Technology routinely measures the accuracy of FR vendor algorithms on various image matching scenarios [4]. By understanding how to read advanced accuracy measurements, FR system administrators will have a better idea of which algorithms will best suit their respective organization's needs.

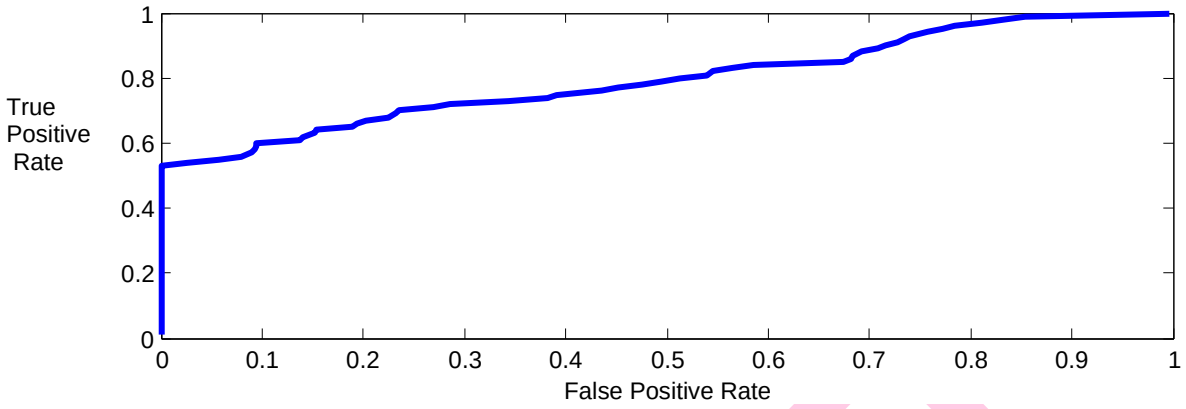

Figure 3 Example ROC curve

6.4.3. Figure 3 is an example of an ROC curve. A ROC curve plots on the x-axis the false positive rate (i.e., the percentage of impostor (or different subject) comparisons that exceed the match threshold) versus the true positive rate (i.e., the percentage of genuine (or same subject) comparisons that exceed the match threshold) on the y-axis. Any particular point on the ROC plot corresponds to the measured true positive and false positive rate at a particular match score threshold. The value of visualizing FR algorithm accuracy in the form of an ROC plot is that an organization (with the help of its integrator) can better determine which match threshold should be used as a function of how many false positive identifications will require human resource and effort to process, or the minimum true positive rate acceptable by its' mission.

6.4.4. Note that while these performance metrics can be computed for most FR systems, they may not be appropriate for certain FRS applications. For example, Rank based systems, used primarily in human adjudicated applications may not benefit from such testing. Further, when tracking

ROC accuracy over time in systems with highly variable gallery characteristics, it may be the case that results are not stable over time. By contrast, FRS applications that threshold match scores will generally have more use for such ROC analysis, as they can use these measurements to tune such decision thresholds over time.

7. Additional Considerations and Best Practices:

7.1. When to perform testing:

7.1.1. Operational testing should be performed before and after all major system updates (e.g., software, hardware, network) to ensure success of any such updates and measure changes in biometric matching accuracy. Such testing is imperative due to the volatility of system updates.

7.1.2. Fixed-interval testing should also be performed. If the system is susceptible to attacks (e.g., cyber-security related), or has experienced recent errors, then fixed-interval testing should occur at a higher frequency.

7.1.3. The system administrator should set operational and fixed-interval testing schedules based on the availability of computing resources and human effort. There is an expectation that Tier 1 and Tier 2 testing will be performed much more frequently than Tier 3 testing, which might be performed only before and after major system updates.

7.2. How to select test images:

7.2.1. Operational test images should target enrollments that span different

periods of time from the near term to the long term. This may help uncover any defects in system performance that result from template corruption, errors in updates, algorithm changes, or patches to extraction and/or matching algorithms.

7.2.2. For Tier 2 and Tier 3 testing, when selecting testing images, subjects contained in those images who are recidivate (i.e., enrolled frequently in the database), should be avoided, as they could reduce the consistency of the accuracy reports.

7.2.3. Subject to agency policies around retention of images, deceased subjects are ideal for enrolling into the database as test images, as deceased state mostly ensures no new encounters with test subjects. The MEDS database [2] contains deceased subjects, and is publicly available.

7.2.4. Data may be collected from an uncontrolled set of test subjects that are reflective of the system's target population or a test crew that is representative of the system's target population, but that has not been enrolled in the system.

7.3. Other**:**

7.3.1. If available, a non-operational system may be configured to use in an "evaluation mode" to collect information not available during normal system operation.

7.3.2. It is important that test data is representative of the operational database in terms of both subjects and image quality.

7.3.3. Due to the introductory nature in this document, there are advanced configurations and operational concerns not addressed which include:

7.3.3.1. Use of score thresholds when testing

7.3.3.2. Use of other accuracy statistics: Detection Error Tradeoff (DET), False Accept Rate (FAR), False Reject Rate (FRR)

7.3.3.3. Demographic differentials in FRS algorithms

7.3.3.4. Selection and continuity of imagery used for probe and database

7.3.3.5. How to address identity ground truth in mate/imposter imagery

7.3.4. These advanced issues and other considerations when testing (e.g., system setup and tuning) will be covered in future FISWG documents.

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