

THE IJB-A CHALLENGE

Recognition of unconstrained faces:
Overview and Participation

Patrick Grother + Mei Ngan

FaCE
Face Recognition Challenges and Evaluations

Four steps to participation

- Request the IJB-A image dataset
 - http://www.nist.gov/itl/iad/ig/ijba_request.cfm
 - We respond to your email with credentials to a download website.
- Execute your face recognition algorithms on one or more of the two defined challenges:
 - 1:1 face verification; 1:N face search.
- Format the recognition algorithm outputs
 - See the examples that accompany the image dataset
 - Tar.gz or zip the directory tree
- Email those recognition outputs to NIST indicating whether your participation is anonymous or not.
 - NIST will post performance results to our leaderboard
 - <http://www.nist.gov/itl/iad/ig/facechallenges.cfm>
 - Optionally, please submit citations [PDF, BIBTEX, URL] to relevant papers

Background

- The design of the IJB-A test is described in Klare et al. *Pushing the Frontiers of Unconstrained Face Detection and Recognition: IARPA Janus Benchmark A*, CVPR, June 2015.
- The design is differentiated from other benchmarks:
 - Evaluate **1:1 verification** in a style similar to the LFW benchmark but with a large number of impostor comparisons to support $FMR \leq 0.001$
 - Evaluate **1:N search** algorithms without assuming an algorithm would execute N 1:1 scores. This affords indexing or tree-based search strategies, for example.
 - Evaluate **open-universe** 1:N search, where not every search has an enrolled identity in the gallery.
 - Always assumes developers have used **external training data**, but includes internal training partitions also.

Task descriptions

- IJB-A 1:1
 - Description: Using "in-the-wild" imagery of 500 persons, this challenge involves comparing genuine and impostor samples for one-to-one verification. Each sample is a composite of still images and video frames, each annotated with face location and landmarks. Each comparison results in a similarity score.
- IJB-A 1:N
 - Description: This is a one-to-many identification task. Using "in-the-wild" imagery of 500 persons, this challenge involves searching samples against an enrolled gallery. This task is open-set: a search sample may or may not have an enrolled mate. Each gallery entry and each search sample is a composite of still and video frames, each annotated with face location and landmarks. Each search results in a list of 20 candidate identities with similarity scores, sorted in descending order.

Required formats for 1:1 submissions

- IJB-A 1:1
 - The required scores are submitted as entries in text files.
 - These reside in a tree, with 10 subdirectories corresponding to the 10 splits (partitions) of the input imagery.
 - Follow the format given in the file [IJB-A_11_output.tar.gz](#) which is present alongside the image distribution.
- NOTES
 - Similarity scores should be non-negative scalars.
 - Similarity scores should not be heavily quantized; enough unique scores should be returned to support fine-grained thresholds and ROC computation.
 - If your algorithm fails to extract features from a sample, or if it fails to complete the comparison, the submitted results should nevertheless include a line representing the attempt and this should include a similarity score set to -1.
 - Template size should be the total size of all features extracted across all input images/frames for each sample.
 - Developers may use external training data. However, individuals who appear in the test sets must be excluded from all training processes.

Required formats for 1:N submissions

- IJB-A 1:1
 - The required candidate lists are submitted as entries in text files.
 - These reside in a tree, with 10 subdirectories corresponding to the 10 splits (partitions) of the input imagery.
 - Follow the format given in the file [IJB-A_1N_output.tar.gz](#) which is present alongside the image distribution.
- NOTES:
 - The candidate list length is 20. If your algorithm produces fewer than 20 candidates, zero-pad the candidate list to a length of 20 (i.e., set all fields to 0 for each unset row)
 - The candidate list for each search should be sorted in decreasing order of similarity score - highest match first.
 - Similarity scores should be non-negative scalars.
 - Similarity scores should not be heavily quantized; enough unique scores should be returned to support fine-grained thresholds and ROC computation.
 - Template size should be the total size of all features extracted across all input images/frames for each input sample.
 - Developers may use external training data. However, individuals who appear in the test sets must be excluded from all training processes. The list of individuals present in the IJB-A set appears in the file [IJB-A_subjectnames_by_id.csv](#) which is present alongside the image distribution.

Submission of algorithm outputs to NIST

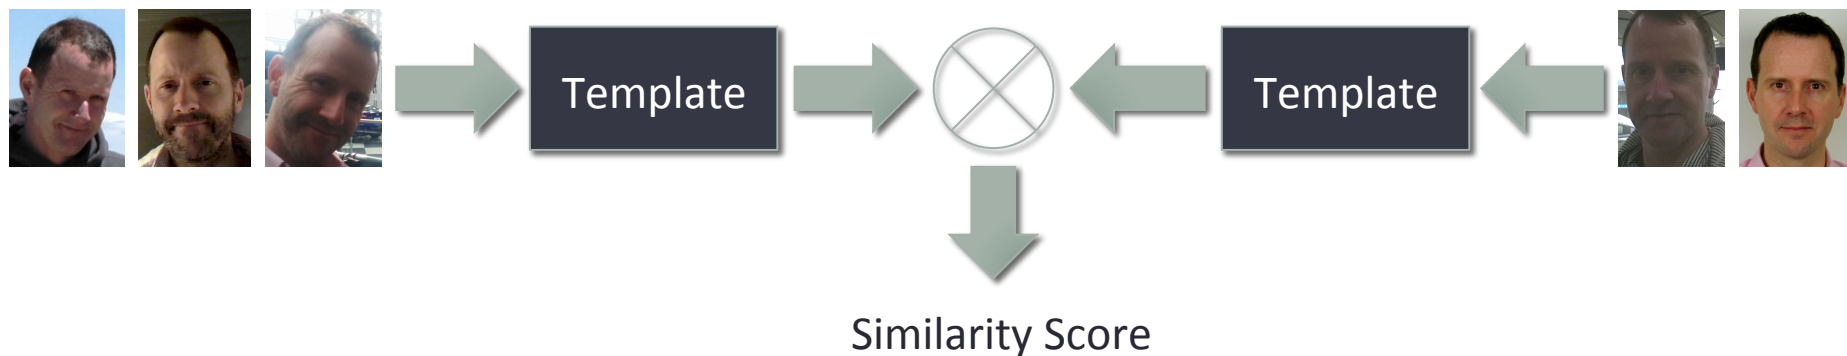
- Email FaceSubmissions@nist.gov with
 - Subject: IJB-A 1:1 submission
 - Attach a file named as follows
 - IJB-A_11_<organization>_<algorithm>_<version>_<date>.tar.gz
 - Indicate whether the results should be attributed to
 - A named organization (e.g. corporation or university)
 - Anonymously
- Subject: IJB-A 1:N submission
- Attach a file named as follows
 - IJB-A_1N_<organization>_<algorithm>_<version>_<date>.tar.gz
- Indicate whether the results should be attributed
 - To a named organization (e.g. corporation or university)
 - Anonymously

NIST publication of results

- Accuracy and other performance data will be posted to NIST's website
 - <http://www.nist.gov/itl/iad/ig/facechallenges.cfm>
- With specific links as follows
 - Verification:
http://nigos.nist.gov:8080/facechallenges/results/IJBA_11_report.pdf
 - Identification:
http://nigos.nist.gov:8080/facechallenges/results/IJBA_1N_report.pdf
- NIST will include any submitted citations to peer-reviewed papers

Use of multiple images

- In IJB-A, a sample or “TEMPLATE” is constructed from $K \geq 1$ images containing the face of the a person.
- Algorithms must extract information from K faces to make a representation of the face that can be used in verification or identification transactions.
 - EXAMPLE: The figures shows production of two templates, one from three images, the other from two. This is followed by the comparison operation which takes two templates and produces a single scalar score.



- How the algorithm integrates information from multiple images is implementation-defined, and is a research topic.

Contact

- Please send questions or comments to
 - FaceChallenges@nist.gov
- Subscribe to news
 - Subscribe to the Face Challenges mailing list to receive emails when announcements or updates are made:
 - <mailto:facechallengesinfo-request@nist.gov?subject=subscribe>