

IARPA Janus Benchmark C

December 11, 2017

Contents

1	Legal Notice	2
2	Notes	3
2.1	IJB-C Changelog	3
2.1.1	General	3
2.1.2	STR Feedback	3
2.1.3	UMD Feedback	4
2.1.4	Additional Noblis Annotation Review and Correction	5
2.2	CS4.1.1 Changelog	7
2.3	CS4.1 Changelog	7
2.4	CS4.0	8
3	Getting Started	9
3.1	Unpacking Instructions	9
3.2	Directory Structure	10
4	Dataset Summary	11
4.1	Media Statistics	11
4.2	Protocol Descriptions	11
4.2.1	Test 1: 1:1 Verification	11
4.2.2	Test 2: 1:1 Covariate Verification	12
4.2.3	Test 4: 1:N Mixed Search	12
4.2.4	Test 6: Face Detection	12
4.2.5	Test 7: Clustering	12
4.2.6	Test 9: Wild Probe with Still Images and Frames	12
4.2.7	Test 10: Wild Probe with Full Motion Video	13
4.2.8	Test 11: Wild Probe Mixed	14
4.3	Metadata Description	16
4.3.1	Enumerated Metadata	16
4.3.2	Face Occlusion	17
5	Cluster Evaluation	19

1 Legal Notice

IARPA Janus Benchmark C (IJB-C) (the Database), is made available under different Creative Commons license variants found in the `license/` directory A Creative Commons (CC) license is one of several public copyright licenses that enable the free distribution of an otherwise copyrighted work.

The Database consists of still images, frames and videos of celebrities and Internet personalities collected from the web. There are 3531 subjects in the Database and the Database is designed to have no overlap with the popular face recognition benchmarks, such as University of Oxford's VGG-Face dataset and the CASIA WebFace dataset.

Users and prospective users of the database will:

1. Comply with the licensing terms of the Database detailed in `license/*.csv` and <https://creativecommons.org/licenses/>
2. Comply with best practices for attribution found here: https://wiki.creativecommons.org/wiki/best_practices_for_attribution
3. Understand that the Database is subject to change at any time.
4. Understand that Noblis makes no guarantee or warranty of non-infringement with respect to the Database.

2 Notes

2.1 IJB-C Changelog

2.1.1 General

- Added janus_degrade_images utility for synthetic image degradation
- Added crop script
- Fixed minor issue with face detection protocol
- Modified synthetic degradation script for rotate-90. The function now includes a fixed-seed PRNG to guarantee deterministic output
- Fixed filename (CS4 vs. IJB-C) issues within the synthetic degradation scripts

2.1.2 STR Feedback

- Rotated bounding boxes for media:
 - frames/5302.png
 - frames/5303.png
 - frames/5304.png
 - frames/5305.png
 - frames/5306.png
 - frames/5309.png
 - frames/5310.png
 - frames/5311.png
 - frames/5313.png
 - frames/5316.png
 - frames/5317.png
 - frames/5318.png
 - frames/5359.png
 - frames/153399.png
- Corrected bounding boxes for media:
 - img/443329.jpg
 - img/444921.jpg
 - img/445657.jpg
 - img/447774.jpg
 - img/448257.jpg
 - img/450481.jpg
 - img/450773.jpg
 - img/451521.jpg
- Corrected bounding boxes for templates:
 - 171018, 179504, 178672, 184101, 181982, 171173, 181985, 180305, 180212, 180818, 228327, 262157, 209356, 214616, 209284, 246692, 95074, 91829, 227775, 229250, 267211, 228327, 214003, 255872, 204171, 265154, 227775, 30111, 223296, 204615, 250427, 217769, 262947, 217767, 221019, 170092, 170092, 170307, 170307, 170405, 170405, 170748, 170748, 170748, 171173, 171173, 171173, 171428, 171662, 178672, 178832, 178832, 180214, 180818, 180818, 181982, 181982, 181985, 181986

- Removed invalid nonface images:
 - nonfaces/228687.jpg
 - nonfaces/228667.jpg
 - nonfaces/229299.jpg
 - nonfaces/231928.jpg
- Fixed occlusion bitfield for img/752.jpg
- Merged subject:
 - 13599 (Maxima of the Netherlands) → 12016 (Queen Maxima of the Netherlands)
- Removed incorrect identity from Kim Tae-young
- Removed chimeric subject Amina Sato
- Rotated bounding box and image for img/101446.jpg
- Corrected/updated media in template 1734 to point to appropriate subject
- Corrected/updated media in template 1829 to point to appropriate subject
- Removed mislabelled templates: 96320, 96322
- Removed incorrect identity from Muhammad Ali (frames/132890.png)
- Removed incorrect identity from Alexandra Soldatova (frames/155428.png, frames/155699.png, frames/155819.png, frames/156190.png, frames/156239.png, frames/156867.png)
- Removed incorrect identity from Andrea Petkovic (frames/173022.png, frames/173858.png, frames/174805.png)
- Removed incorrect identity from Lilia Podkopyeva (img/122288.jpg)
- Removed incorrect identity from Daria Kondakova (img/124295.jpg)
- Removed incorrect identity from Alessia Russo (img/143526.jpg)
- Removed incorrect identity from Katie Ledecky (frames/397985.png)

2.1.3 UMD Feedback

- Merged subjects:
 - 13270 (Amarinder Singh) → 13116 (Amrinder Singh)
 - 12521 (Kim Hyun-a) → 1270 (Hyuna)
- Removed chimeric subject Deepika Singh
- Removed chimeric subject Lech Kaczynski
- Removed chimeric subject Jarosaw Kaczynski
- Removed chimeric subject Aditi Sharma
- Removed chimeric subject Shraddha Pandit
- Removed chimeric subject Kashmira Shah
- Removed chimeric subject Mai Haruna
- Removed chimeric subject Miho Miyazaki

- Removed chimeric subject Tomomi Kasai
- Removed incorrect identity from Andreea Banica
- Remove templates 80452, 18290, 18292 to resolve Akshay Kumar/Diljit Dosanjh labelling issue ,
- Fixed bounding box orientation in templates 170650, 181982, 181984, 181985, 181986, 171173, 184959, 184961, 184962, 184963, 184964, 184965, 171018, 180212, 180214, 180215, 180216, 170748, 178672, 178674, 178675, 178676, 171572, 187205, 187209, 187210, 187211, 180305, 180307, 180308, 180309

2.1.4 Additional Noblis Annotation Review and Correction

- Removed incorrect identity from Sukhbir Singh Badal media (frames/170766.png)
- Removed incorrect identity from Alina Maksymenko (frames/432428.png, frames/361901.png)
- Removed incorrect identity from David Kobia (img/440590.jpg)
- Removed incorrect identity from Hendra Setiawan (frames/377140.png)
- Removed incorrect identity from Javier Aquino (frames/431904.png, frames/396258.png, frames/382652.png, frames/418435.png, frames/353559.png, frames/429971.png, frames/438457.png, frames/403021.png, frames/344361.png)
- Removed incorrect identity from Michael Eric Dyson (frames/401097.png, frames/355559.png)
- Removed incorrect identity from Ral Jimnez (frames/341530.png, frames/398038.png, frames/387441.png)
- Removed incorrect identity from Vladimir Putin (frames/6629.png, frames/7389.png, frames/7390.png, frames/7391.png)
- Removed incorrect identity from Dilma Rousseff (frames/73901.png)
- Removed incorrect identity from Tony Blair (frames/17516.png, frames/12690.png, frames/12698.png, frames/12711.png, frames/12713.png, frames/12714.png, frames/12722.png, frames/12724.png, frames/12727.png, frames/12729.png)
- Removed incorrect identity from Saddam Hussein (frames/23565.png)
- Removed incorrect identity from Kevin Rose (frames/126502.png)
- Removed incorrect identity from Ma Maestro (frames/129221.png)
- Removed incorrect identity from Johnny Cash (frames/216741.png, frames/216792.png)
- Added frame to Dan Rather (frames/23565.png)
- Corrected bounding box for Anglica Rivera (img/2058.jpg)
- Corrected bounding box for Phumzile Mlambo-Ngcuka (img/124445.jpg)
- Corrected bounding box for Saad Lamjarred (img/124487.jpg)
- Corrected bounding box for Abhisit Vejjajiva (img/4929.jpg)
- Corrected bounding box for Naoto Kan (frames/10905.png)
- Corrected bounding box for Barack Obama (frames/15385.png)
- Corrected bounding boxes for Vladimir Putin (img/2969.jpg, img/1416.jpg)
- Corrected bounding box for Bill Clinton (frames/18650.png)
- Corrected bounding box for Mohamed Morsi (frames/15851.png)

- Corrected bounding boxes for Jacques Chirac (frames/9923.png, frames/9918.png)
- Corrected bounding box for Silvio Berlusconi (frames/18663.png)
- Corrected bounding box for Gary Locke (frames/13170.png)
- Corrected bounding boxes for Colin Powell (frames/13485.png, frames/13487.png)
- Corrected bounding box for Larry Ellison (img/4582.jpg)
- Corrected bounding boxes for Gene Robinson (frames/18756.png, frames/18764.png)
- Corrected bounding box for Kevin Rudd (frames/10016.png)
- Corrected bounding box for Gursharan Kaur (img/150.jpg)
- Corrected bounding box for Thabo Mbeki (img/933.jpg)
- Corrected bounding boxes for Michelle Obama (frames/5294.png, frames/5299.png)
- Corrected bounding box for Khaleda Zia (frames/16083.png)
- Corrected bounding box for Nigel Farage (frames/5182.png)
- Corrected bounding box for Anand Sharma (img/695.jpg)
- Corrected bounding box for Paul Bremer (img/752.jpg)
- Corrected bounding box for Nagma (frames/127090.png)
- Corrected bounding box for Kofi Kingston (img/116318.jpg)
- Corrected bounding box for Aecio Nevesa (img/3880.jpg)
- Corrected bounding box for Emily Haines (frames/210036.png)
- Corrected bounding box for Oxmo Puccino (img/120783.jpg)
- Corrected bounding box for Deepa Sannidhi (img/112037.jpg)
- Removed duplicate image for Hyuna (img/5076.jpg)
- Removed chimeric subject Adam Price
- Merged subjects and removed duplicate pieces of media from template:
 - 2447 (Rached Ghannouchi) → 404 (Rashid al-Ghannushi)
 - 27540 (Femi Kutu) → 13325 (Fela Kutu)
- Removed bad bounding boxes from face detection protocols from media:
 - img/124711.jpg
 - img/4484.jpg
 - img/448725.jpg
 - img/440838.jpg
 - img/450277.jpg
 - img/442193.jpg
 - img/2568.jpg
 - frames/53937.png
- Updated licensing file and removed orphaned media

2.2 CS4.1.1 Changelog

- Cleaned and corrected wild protocol descriptions in sections 4.2.6, 4.2.7, and 4.2.8

2.3 CS4.1 Changelog

- Rotated bounding boxes for media:
 - `img/107689.jpg` ← Media also rotated
 - `img/4280.jpg`
- Removed bad box from `frames/157752.png`
- Removed extra box from `frames/14062.png`
- Added wild protocols (tests 9 , 10, and 11)
- Updated bounding boxes for small subset of media
- Renamed Galleries $S\{1, 2\} \rightarrow G\{1, 2\}$
- Added test 7 cluster evaluation script.
- Removed subjects from G1 and relocated corresponding media to G2:
 - 12498
 - 12666
 - 12790
- Updated frame number for all frames
- Removed small subset of redundant videos and frames from protocols
- Removed templates from 1:1 Verification matches file:
 - 170111
 - 170113
 - 170599
- Removed 1:1 Covariate Verification templates:
 - 68350
 - 76145
 - 81890
 - 90201
 - 29251
- Removed clustering test 7 templates:
 - 188230
 - 167972
 - 239608
 - 217718
 - 21402
- Removed 1:N mixed probe templates:
 - 187230
 - 187233

2.4 CS4.0

- An EXIF parsing bug exists in OpenCV 3.1.0 which causes `imread()` to hang indefinitely while decoding some JPEG images. See: <https://github.com/opencv/opencv/pull/6800>
- All coordinates are integers
- Reduced number of test 7 clustering protocols
- Removed clustering test 8 (now only for internal T&E)
- Removed 1:N Image protocol
- Face detection protocol now includes an ignore flag. This flag is set to 1 if a Turker noted any errors (extra box, missing box, incorrect box) with the bounding box(es) for a given piece of media.

3 Getting Started

3.1 Unpacking Instructions

The following will first extract all IJB-C media followed by the IJB-C protocols.

```
$ #IJB-C
$ cd /path/to/IJBC-1.0.tar.XXX
$
$ #Extraction of media
$ cat IJBC-1.0.tar.* | tar xvf - -C /output/path
$
$ #Extraction of IJB-C protocols
$ tar xvfz /path/to/IJBC.tar.gz -C /output/path
```

3.2 Directory Structure

```
IJB-C/
|-- ijbc_subject_names.csv          # Map of Subject ID --> Subject Name
|-- frames/                        # IJB-C Frames
|-- img/                           # IJB-C Images
|-- license/                       # CS{3,4} Licensing and Attribution Information
|   |-- cs3_media.csv
|   |-- cs4_media.csv
|   '-- ijbc_license_info.csv
|-- media.md5                      # frames/* img/* video/* MD5 hashes
|-- nonfaces/                     # IJB-C Nonface Images
|-- protocols/
|   |-- clustering/
|   |   |-- ijbc_clustering_1021_hint_10000.csv
|   |   |-- ijbc_clustering_1839_hint_10000.csv
|   |   |-- ijbc_clustering_32_hint_100.csv
|   |   |-- ijbc_clustering_3531_hint_10000.csv
|   |   '-- ijbc_clustering_ground_truth.csv
|   |-- ijbc_11_covariate_matches.csv
|   |-- ijbc_11_covariate_probe_reference.csv
|   |-- ijbc_11_G1_G2_matches.csv
|   |-- ijbc_1N_gallery_G1.csv
|   |-- ijbc_1N_gallery_G2.csv
|   |-- ijbc_1N_probe_mixed.csv
|   |-- ijbc_face_detection.csv
|   |-- ijbc_face_detection_ground_truth.csv
|   |-- ijbc_metadata.csv
|   |-- ijbc_wild_test10.csv
|   |-- ijbc_wild_test11.csv
|   '-- ijbc_wild_test9.csv
|-- README.pdf
|-- utilities/
|   |-- cluster_eval/
|   |   '-- modified_bcubed.py
|   |-- ijbb_crop/
|   |   '-- crop
|   '-- janus_degrade_images/
|       |-- CMakeLists.txt
|       |-- degrade_operations.cpp
|       |-- degrade_operations.h
|       |-- fast-cpp-csv-parser/
|       |   |-- csv.h
|       |   |-- LICENSE
|       |   '-- README.md
|       |-- LICENSE.txt
|       |-- main.cpp
|       '-- README
'-- video/                        # IJB-C Videos
```

4 Dataset Summary

4.1 Media Statistics

# Images	21294
# Frames	117542
# Videos	11779
# Nonfaces	10040

Table 1: Media Counts

4.2 Protocol Descriptions

Test	Mode	Purpose	”Curation”
1	1:1 Verification	Templates are comprised of mixed media (frames and stills)	All templates
2	1:1 Covariate Verification	Single still-image comparisons to perform failure analysis. <i>e.g. yaw difference</i>	All templates
4	1:N Mixed Search	Open set 1:N protocol using mixed media (frames, stills) as probe	G1/G2
6	Face Detection	Pure detection, without recognition. Includes pile of non-face imagery	–
7	Clustering	Cluster cropped faces. Bounding box and landmark annotations are provided to delineate the subject of interest.	–
9	Wild Probe Still/Frames	Identify subjects of interest from a collection of still images and frames.	G1/G2
10	Wild Probe Video	Identify subjects of interest from a collection of videos	G1/G2
11	Wild Probe Mixed	Identify subjects of interest from a collection of still images and videos. Templates are created from subject clusters and then searched against galleries G1 and G2.	G1/G2

Table 2: Summary of Protocol Descriptions.

4.2.1 Test 1: 1:1 Verification

See Tables 2 and 3.

# Subjects	3531
# Templates	23124
# Genuine Matches	19557
# Impostor Matches	15638932

Table 3: 1:1 Verification Statistics

4.2.2 Test 2: 1:1 Covariate Verification

See Tables 2 and 4.

# Subjects	3531
# Templates	140739
# Genuine Matches	7819362
# Impostor Matches	39584639

Table 4: 1:1 Covariate Verification Statistics

4.2.3 Test 4: 1:N Mixed Search

See Tables 2 and 5.

# Templates	19593
# Files	127152

Table 5: 1:N Mixed Search Statistics

4.2.4 Test 6: Face Detection

See Table 2.

4.2.5 Test 7: Clustering

The clustering protocol is designed to test an algorithm’s ability to identify multiple instances of the same subject from a collection of various pieces of media. In experiment 7, there are seven clustering sub-protocols, each with a different number of subjects and hint. See Table 6.

Subjects	32	1021	1839	3531
Hints	100	10000	10000	10000
# Templates	955	41074	71392	140623

Table 6: Test 7 Sub-Protocol Subjects and Hints

The hints serve as a coarse upper bound on the number of true clusters within each sub-protocol. The hint is calculated by the following equation: $10^{\lceil \log_{10} |\text{subjects}| \rceil}$. For each sub-protocol, all imagery for each selected subject in IJB-C is used (still images and video frames) and is a superset of the previous sub-protocol (e.g., Clustering-1021 contains all media from Clustering-32 plus media from 989 additional subjects). The input in this protocol is a collection of face images and bounding boxes that delineate the face of interest. The faces are then treated as individual items to be clustered. If an FTE occurs, the cluster number should be set to -1. BCubed [1] precision and recall, and F-measure are used as the primary evaluation metric. FTEs, although not clustered, still count towards evaluation with 0 precision and 0 recall.

4.2.6 Test 9: Wild Probe with Still Images and Frames

Janus performers are tasked with identifying subjects of interest (IJB-C subjects) from a collection of still images and frames. Performers will detect faces and search each detected face against the IJB-C galleries G1 and G2. This experiment resembles the operational workflow performed in practice today. **NOTE:** Does not involve subject clustering.

1. Each single-face template will be matched to the IJB-C curated galleries (G1 and G2).
2. CMC and DET/IET curves will be computed based on the matching performance of each individual template.

The API calls, given in the order in which they are called, are:

1. `janus_initialize(const std::string &sdk_path,
const std::string &temp_path
const std::string &algorithm,
const int gpu_dev)`
2. `janus_create_template(const janus_media &media,
const janus_template_role role,
std::vector<janus_template> &templates,
std::vector<janus_track> &tracks)`

This produces a list of templates and the bounding boxes used to create them. The template role is IDENTIFICATION. This is called repeatedly for each image or frame in the protocol.

3. `janus_search(const janus_template &probe,
const janus_gallery &gallery,
const size_t num_requested_returns,
std::vector<janus_template_id> &template_ids,
std::vector<double> &similarities)`

Templates constructed in step 2 are searched against curated galleries G1 and G2.

4. `janus_finalize()`

Four output files are to be generated for datacalls:

1. A metadata file containing detections/tracks with the following header:
TEMPLATE_ID,FILENAME,FRAME_NUM,FACE_X,FACE_Y,FACE_WIDTH,FACE_HEIGHT,CONFIDENCE,DETECTION_TIME
2. Two candidate lists (searching G1, G2) with the following header:
SEARCH_TEMPLATE_ID,RANK,ERROR_CODE,GALLERY_TEMPLATE_ID,SCORE,SEARCH_TIME
3. A template creation file with the following header:
TEMPLATE_ID,TEMPLATE_ROLE,ERROR_CODE,TEMPLATE_CREATION_TIME,TEMPLATE_SIZE

For the evaluation, T&E will associate the bounding boxes for the detected faces with the ground truth bounding boxes of the IJB-C identities.

4.2.7 Test 10: Wild Probe with Full Motion Video

Same as experiment 9 but with full motion videos used as input. **NOTE:** Does not involve subject clustering.

1. A template is created from each face track and then searched against the curated IJB-C galleries (G1 and G2).
2. CMC and DET/IET curves will be computed based on the matching performance of each individual template.

The API calls, given in the order in which they are called, are:

1. `janus_initialize(const std::string &sdk_path,
const std::string &temp_path
const std::string &algorithm,
const int gpu_dev)`

```
2. janus_create_template(const janus_media &media,
                        const janus_template_role role,
                        std::vector<janus_template> &templates,
                        std::vector<janus_track> &tracks)
```

This produces a list of templates and the bounding boxes used to create them. The template role is IDENTIFICATION. This is called repeatedly for each image or frame in the protocol.

```
3. janus_search(const janus_template &probe,
               const janus_gallery &gallery,
               const size_t num_requested_returns,
               std::vector<janus_template_id> &template_ids,
               std::vector<double> &similarities)
```

Templates constructed in step 2 are searched against curated galleries G1 and G2.

```
4. janus_finalize()
```

Four output files are to be generated for datacalls:

1. A metadata file containing detections/tracks with the following header:
 TEMPLATE_ID,FILENAME,FRAME_NUM,FACE_X,FACE_Y,FACE_WIDTH,FACE_HEIGHT,CONFIDENCE,DETECTION_TIME
2. Two candidate lists (searching G1, G2) with the following header:
 SEARCH_TEMPLATE_ID,RANK,ERROR_CODE,GALLERY_TEMPLATE_ID,SCORE,SEARCH_TIME
3. A template creation file with the following header:
 TEMPLATE_ID,TEMPLATE_ROLE,ERROR_CODE,TEMPLATE_CREATION_TIME,TEMPLATE_SIZE

For the evaluation, T&E will associate the bounding boxes for the detected faces with the ground truth bounding boxes of the IJB-C identities.

4.2.8 Test 11: Wild Probe Mixed

The wild probe mixed protocol is designed to test an algorithm's ability to create and match templates from a collection of media (still images and full-motion video) to a curated gallery. Performers are tasked with building templates for all identities determined to be within the collection of media. Performers will create templates from single images or videos, cluster those templates, and then create multi-media templates from the clustering output. Noblis will then perform a 1:N search test between the multi-media templates and the curated galleries (G1 and G2).

Noblis is developing evaluation scripts for this protocols and will release them upon completion.

The API calls for all options, given in the order in which they are called, are:

1. A template is created from each face track and then searched against the curated IJB-C galleries (G1 and G2).
2. CMC and DET/IET curves will be computed based on the matching performance of each individual template.

The API calls, given in the order in which they are called, are:

```
1. janus_initialize(const std::string &sdk_path,
                  const std::string &temp_path,
                  const std::string &algorithm,
                  const int gpu_dev)
```

2. `janus_create_template(const janus_media &media,
 const janus_template_role role,
 std::vector<janus_template> &templates,
 std::vector<janus_track> &tracks)`

This produces a list of templates and the bounding boxes used to create them. The template role is CLUSTERING. This is called repeatedly for each image or video in the protocol.

3. `janus_cluster(const std::vector<janus_template> &templates,
 const size_t hint,
 std::vector<cluster_pair> &clusters)`

`janus_cluster()` is called with all templates constructed in step 2.

4. `janus_create_template(std::vector<janus_association> &associations,
 const janus_template_role role,
 janus_template &template_)`

`janus_create_template()` will be called for each CLUSTER_INDEX_ID with role set to IDENTIFICATION.

5. `janus_search(const janus_template &probe,
 const janus_gallery &gallery,
 const size_t num_requested_returns,
 std::vector<janus_template_id> &template_ids,
 std::vector<double> &similarities)`

Templates constructed in step 4 are searched against curated galleries G1 and G2.

6. `janus_finalize()`

NOTE: API changes are under review that will enable the ability to directly cluster media instead of templates.

Seven output files are to be generated for datacalls:

1. After detection, a metadata file is generated with the following header:
TEMPLATE_ID,FILENAME,FRAME_NUM,FACE_X,FACE_Y,FACE_WIDTH,FACE_HEIGHT,CONFIDENCE,DETECTION_TIME
NOTE: FRAME_NUM=NaN for images
2. After clustering, a metadata file is generated with the following header:
TEMPLATE_ID,CLUSTER_INDEX_ID
3. Two candidate list files (searching G1, G2) with the following header:
CLUSTER_INDEX_ID,RANK,ERROR_CODE,GALLERY_TEMPLATE_ID,SCORE,SEARCH_TIME
4. A clustering duration file with the following header:
CLUSTER_TIME
5. A template creation file for clustering:
TEMPLATE_ID,TEMPLATE_ROLE,ERROR_CODE,TEMPLATE_CREATION_TIME,TEMPLATE_SIZE
6. A template creation file for search:
CLUSTER_INDEX_ID,TEMPLATE_ROLE,ERROR_CODE,TEMPLATE_CREATION_TIME,TEMPLATE_SIZE

For the evaluation, T&E will associate the bounding boxes for the detected faces with the ground truth bounding boxes of the IJB-C identities.

4.3 Metadata Description

Attributes collected for images in the Janus program were annotated using crowd sourcing. As a result, attributes are defined in layman’s terms, instead of strict scientific definitions. For more information, see <http://libjanus.org>

4.3.1 Enumerated Metadata

0-19	20-34	35-49	50-64	65+	Unknown	No Annotation
1	2	3	4	5	0	NaN

Table 7: Age Enumeration

Outdoor	Indoor	No Annotation
0	1	NaN

Table 8: Indoor/Outdoor

Light Pink	1
Light Yellow	2
Medium Pink/Brown	3
Medium Yellow/Brown	4
Medium-Dark Brown	5
Dark Brown	6
No Annotation	NaN

Table 9: Skin Tone Enumeration

No Facial Hair	Moustache	Goatee	Beard	No Annotation
0	1	2	3	NaN

Table 10: Facial Hair Enumeration

Female	Male	No Annotation
0	1	NaN

Table 11: Gender

Roll		Yaw	
Min	Max	Min	Max
-135°	62°	-87°	78°

Table 12: Face Orientation

4.3.2 Face Occlusion

Occlusion annotations for IJB-C media are summarized by the grid shown in Figure 1.




Figure 1: Occlusion Grid

Some examples of occlusion grid annotations are shown below:




Figure 2: Full Eye Occlusion




Figure 3: Single Eye Occlusion




Figure 4: Full Mouth Occlusion




Figure 5: Partial Mouth Occlusion




Figure 6: Full Forehead Occlusion




Figure 7: Full Forehead Occlusion – Bangs




Figure 8: Full Forehead Occlusion – Side Bangs



Figure 9: Full Forehead Occlusion – Short Bangs

5 Cluster Evaluation

Two scripts are provided for cluster evaluation in the `cluster_eval` directory which contains an implementation of the BCubed evaluation metric relevant to IJB-C test 7.

- `modified_bcubed.py`: This is used to evaluate clustering Test 7 (i.e. IJB-C subject of interests when given the face bounding box). The script can be called as such:


```
./modified_bcubed.py cluster_file ground_truth [-no_fte] [-h]
```

 - `cluster_file` is the output of Test 7 and should have a header of:
`TEMPLATE_ID,FILENAME,CLUSTER_INDEX,CONFIDENCE`
 - `ground_truth` file contains the mapping between template IDs and filenames to subject IDs. It should have a header of:
`TEMPLATE_ID,SUBJECT_ID,FILENAME`
 - `[-no_fte]` If set, all FTEs with a cluster assignment set to -1 will be ignored for all calculations.
 - `[-h]` Help

The provided implementation of the BCubed metric has been modified according to the following guidelines:

- For test 7, any face that is a Failure To Enroll (FTE) will be counted as an event and evaluated with 0 precision and 0 recall
 - BR events clustered with IJB-C SOIs are used in the calculation of precision
 - Pure BR events will be ignored completely during evaluation

References

- [1] Enrique Amigo, Julio Gonzalo, Javier Artiles, and Felisa Verdejo. A comparison of extrinsic clustering evaluation metrics based on formal constraints. *Inf. Retr.*, 12(4):461–486, August 2009.