

1



2

3

4

5

6

7

8

Face In Video Evaluation (FIVE)
Concept, Evaluation Plan, and API
Version 0.4

9

10

11

12

13

14

Patrick Grother and Mei Ngan

Image Group

Information Access Division

Information Technology Laboratory



October 3, 2014

15

16

17

18

Timeline of the FIVE Evaluation

Phase	Date	External actions, deadlines
Phase 0	2014-07-15	Web site up, announce schedule
	2014-08-15	First draft Evaluation Plan and API
	2014-08-31	Public comments on first drafts due
	2014-10-01	Second draft Evaluation Plan and API
	2014-10-15	Public comments on second drafts due
	2014-10-30	Third draft Evaluation Plan and API. Draft five.h available.
	2014-11-07	Public comments on third drafts due
	2014-11-10	Final Evaluation Plan and API available. Final five.h available
	2014-11-10	FIVE validation package available
Phase 1	2014-11-17	Opening of Phase 1 submission period
	2015-02-08	Deadline for submission for inclusion of results in first interim report card
	2015-03-28	First interim report card released to submitting participants
Phase 2	2015-04-01	Opening of Phase 2 submission period
	2015-06-05	Deadline for submission for inclusion of results in second interim report card.
	2015-07-30	Second interim report card released to submitting participants
Phase 3	2015-08-01	Opening of Phase 3
	2015-10-05	Deadline for submission of algorithms to Phase 3

19

20

November 2014							December 2014							January 2015							February 2015							March 2015							April 2015							
Su	Mo	Tu	We	Th	Fr	Sa	Su	Mo	Tu	We	Th	Fr	Sa	Su	Mo	Tu	We	Th	Fr	Sa	Su	Mo	Tu	We	Th	Fr	Sa	Su	Mo	Tu	We	Th	Fr	Sa	Su	Mo	Tu	We	Th	Fr	Sa	
						1	1	2	3	4	5	6							1	2	3	1	2	3	4	5	6	7	1	2	3	4	5	6	7	1	2	3	4	5	6	7
2	3	4	5	6	7	8	7	8	9	10	11	12	13	4	5	6	7	8	9	10	8	9	10	11	12	13	14	8	9	10	11	12	13	14	5	6	7	8	9	10	11	
9	10	11	12	13	14	15	14	15	16	17	18	19	20	11	12	13	14	15	16	17	15	16	17	18	19	20	21	15	16	17	18	19	20	21	12	13	14	15	16	17	18	
16	17	18	19	20	21	22	21	22	23	24	25	26	27	18	19	20	21	22	23	24	22	23	24	25	26	27	28	22	23	24	25	26	27	28	19	20	21	22	23	24	25	
23	24	25	26	27	28	29	28	29	30	31	25	26	27	28	29	30	31								29	30	31	26	27	28	29	30										
30																																										
May 2015							June 2015							July 2015							August 2015							September 2015														
Su	Mo	Tu	We	Th	Fr	Sa	Su	Mo	Tu	We	Th	Fr	Sa	Su	Mo	Tu	We	Th	Fr	Sa	Su	Mo	Tu	We	Th	Fr	Sa	Su	Mo	Tu	We	Th	Fr	Sa								
						1	2	1	2	3	4	5	6	5	6	7	8	9	10	11	2	3	4	5	6	7	8	6	7	8	9	10	11	12								
3	4	5	6	7	8	9	7	8	9	10	11	12	13	12	13	14	15	16	17	18	9	10	11	12	13	14	15	13	14	15	16	17	18	19								
10	11	12	13	14	15	16	14	15	16	17	18	19	20	19	20	21	22	23	24	25	16	17	18	19	20	21	22	20	21	22	23	24	25	26								
17	18	19	20	21	22	23	21	22	23	24	25	26	27	26	27	28	29	30	31	23	24	25	26	27	28	29	27	28	29	30												
24	25	26	27	28	29	30	28	29	30					30	31	30	31	27	28	29	30																					
31																																										

21

22 Major API Changes since FRVT 2013 Class V

23 The header/source files for the API will be made available to implementers at <http://nigos.nist.gov:8080/five>.

- 24 — The structures ONEFACE (see Table 11) and MULTIFACE (see Table 12) have been changed to classes.
- 25 — The MULTIFACE class contains a new “description” member variable and valid values are specified in Table 10.
- 26 — The labels for describing types of still images have been updated (see Table 9).
- 27 — The ONEVIDEO (see Table 14) class contains a new “peopleDensity” member variable and valid values are specified in
- 28 Table 13.

29

30 **Table of Contents**

31 1. FIVE 6

32 1.1. Scope 6

33 1.2. Audience 6

34 1.3. Market drivers 7

35 1.4. Offline testing 7

36 1.5. Phased testing 7

37 1.6. Interim reports 7

38 1.7. Final reports..... 7

39 1.8. Application scenarios..... 8

40 1.9. Image source labels 8

41 1.10. Rules for participation 8

42 1.11. Number and schedule of submissions 8

43 1.12. Use of multiple images per person 9

44 1.13. Core accuracy metrics 9

45 1.14. Generalized accuracy metrics..... 10

46 1.15. Reporting template size..... 10

47 1.16. Reporting computational efficiency 10

48 1.17. Exploring the accuracy-speed trade-space 10

49 1.18. Hardware specification 10

50 1.19. Operating system, compilation, and linking environment 11

51 1.20. Software and documentation..... 11

52 1.21. Runtime behavior 12

53 1.22. Threaded computations 13

54 1.23. Time limits 13

55 1.24. Test datasets..... 14

56 1.25. Ground truth integrity 15

57 2. Data structures supporting the API 15

58 2.1. Overview..... 15

59 2.2. Requirement..... 15

60 2.3. File formats and data structures 15

61 2.4. File structures for enrolled template collection 20

62 3. API Specification 21

63 3.2. 1:N Identification..... 22

64 3.3. Interfaces..... 24

65 4. References 36

66 Annex A Submission of Implementations to the FIVE 37

67 A.1 Submission of implementations to NIST 37

68 A.2 How to participate..... 37

69 A.3 Implementation validation 38

70

71 **List of Figures**

72 Figure 1 – Organization and documentation of the FIVE 6

73 Figure 2 – Examples of pose angles and their encodings (yaw, pitch) 16

74

75 **List of Tables**

76 Table 1 – Abbreviations..... 5

77 Table 2 – Subtests supported under the FIVE activity..... 8

78 Table 3 – Cumulative total number of algorithms 8

79 Table 4 – Summary of accuracy metrics..... 9

80 Table 5 – Implementation library filename convention 11

81 Table 6 – Number of threads allowed for each function 13

FIVE

82 Table 7 – Processing time limits in milliseconds 13

83 Table 8 – Main video corpora (others will be used) 14

84 Table 10 – Labels describing types of images..... 15

85 Table 11 – Labels describing types of MULTIFACES..... 16

86 Table 12 – ONEFACE class 17

87 Table 13 – MULTIFACE class 17

88 Table 14 – Labels describing the density of people in the video frames 17

89 Table 15 – ONEVIDEO Class 17

90 Table 16 – EYEPAIR Class 18

91 Table 17 – PersonTrajectory typedef 18

92 Table 18 – PERSONREP Class 19

93 Table 19 – CANDIDATE Class 19

94 Table 20 – CANDIDATELIST typedef 19

95 Table 21 – ReturnCode class 20

96 Table 22 – Enrollment dataset template manifest 20

97 Table 23 – API implementation requirements for FIVE 21

98 Table 24 – Procedural overview of the identification test 22

99 Table 25 – VideoEnrollment::initialize 24

100 Table 26 – VideoEnrollment::generateEnrollmentTemplate 25

101 Table 27 – VideoFinalize::finalize 26

102 Table 28 – VideoFeatureExtraction::initialize 27

103 Table 29 – VideoFeatureExtraction::generateIdTemplate 27

104 Table 30 – VideoSearch::initialize 28

105 Table 31 – VideoSearch::identifyVideo and VideoSearch::identifyImage 29

106 Table 32 – ImageEnrollment::initialize 30

107 Table 33 – ImageEnrollment::generateEnrollmentTemplate..... 30

108 Table 34 – ImageFinalize::finalize..... 31

109 Table 35 – ImageFeatureExtraction::initialize 32

110 Table 36 – ImageFeatureExtraction::generateIdTemplate 33

111 Table 37 – ImageSearch::initialize 34

112 Table 38 – ImageSearch::identifyVideo..... 34

113

114

115 **Acknowledgements**

116 — The authors are grateful to the experts who made extensive comments on the first version of this document.

117 **Project History**

118 — 2012 – 2014 – The FRVT 2013 program included a video track (class V) that evaluated face recognition from video.
 119 The FIVE program supersedes the FRVT work but proceeds in an almost identical manner.

120 — August 15, 2014 - Release of first public draft of the Face In Video Evaluation (FIVE) – Concept, Evaluation Plan and
 121 API v0.1.

122 **Terms and definitions**

123 The abbreviations and acronyms of Table 1 are used in many parts of this document.

124 **Table 1 – Abbreviations**

FNIR	False negative identification rate
FPIR	False positive identification rate
FIVE	NIST’s Face In Video Evaluation program
FRVT	NIST’s Face Recognition Vendor Test program
FTA	Failure to acquire a search sample
FTE	Failure to extract features from an enrollment image
DET	Detection error tradeoff characteristic: For identification this is a plot of FNIR vs. FPIR.
INCITS	InterNational Committee on Information Technology Standards
ISO/IEC 19794	ISO/IEC 19794-5: Information technology — Biometric data interchange formats — Part 5:Face image data. First edition: 2005-06-15. (See Bibliography entry).
MBE	NIST’s Multiple Biometric Evaluation program
NIST	National Institute of Standards and Technology
SDK	The term Software Development Kit refers to any library software submitted to NIST. This is used synonymously with the terms "implementation" and "implementation under test".

125

126 **1. FIVE**

127 **1.1. Scope**

128 The Face In Video Evaluation (FIVE) is being conducted to assess the capability of face recognition algorithms to correctly
 129 identify or ignore persons appearing in video sequences – i.e. the open-set identification problem. Both comparative and
 130 absolute accuracy measures are of interest, given the goals to determine which algorithms are most effective and
 131 whether any are viable for the following primary operational use-cases:

- 132
- 133 1. High volume screening of persons in the crowded spaces (e.g. an airport)
- 134 2. Low volume forensic examination of footage from a crime scene (e.g. a convenience store)
- 135 3. Persons in business meetings (e.g. for video-conferencing)
- 136 4. Persons appearing in television footage

137

138 These applications differ in their tolerance of false positives, whether a human examiner will review outputs, the prior
 139 probabilities of mate vs. non-mate presence, and the cost of recognition errors.

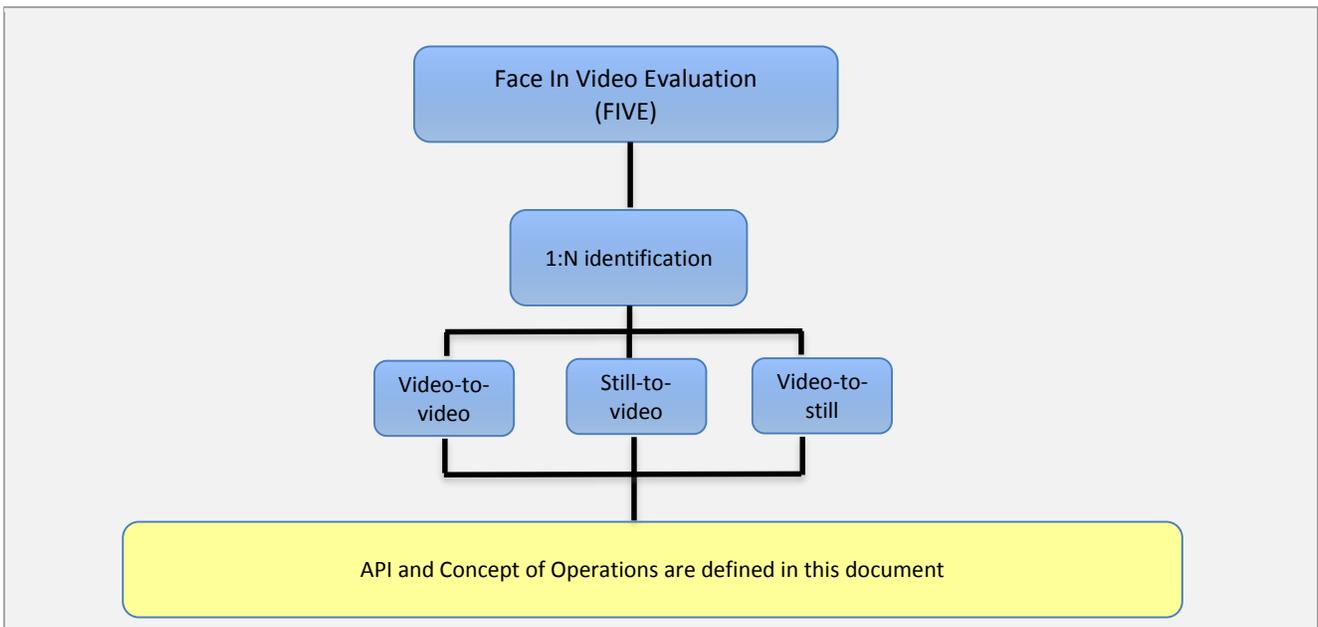
140

141 **Out of scope:** Areas that are out of scope for this evaluation and will not be studied include: gait, iris and voice
 142 recognition; recognition across multiple views (e.g. via stereoscopic techniques); tracking across sequential cameras (re-
 143 identification); anomaly detection; detection of evasion.

144

145 This document establishes a concept of operations and an application programming interface (API) for evaluation of face
 146 recognition in video implementations submitted to NIST's Face In Video Evaluation. See
 147 <http://www.nist.gov/itl/iad/ig/five.cfm> for all FIVE documentation.

148



149

Figure 1 – Organization and documentation of the FIVE

150 **1.2. Audience**

151 Universities and commercial entities with capabilities in detection and identification of faces in video sequences are
 152 invited to participate in the FIVE Video test.

153 Organizations will need to implement the API defined in this document. Participation is open worldwide. There is no
 154 charge for participation. While NIST intends to evaluate technologies that could be readily made operational, the test is
 155 also open to experimental, prototype and other technologies.

156 **1.3. Market drivers**

157 This test is intended to support a plural marketplace of face recognition in video systems. There is considerable interest
 158 in the potential use of face recognition for identification of persons in videos.

159 **1.4. Offline testing**

160 While this set of tests is intended as much as possible to mimic operational reality, this remains an offline test executed
 161 on databases of images. The intent is to assess the core algorithmic capability of face recognition in video algorithms. This
 162 test will be conducted purely offline - it does not include a live human-presents-to-camera component. Offline testing is
 163 attractive because it allows uniform, fair, repeatable, and efficient evaluation of the underlying technologies. Testing of
 164 implementations under a fixed API allows for a detailed set of performance related parameters to be measured.

165 **1.5. Phased testing**

166 To support research and development efforts, this testing activity will embed multiple rounds of testing. These test
 167 rounds are intended to support improved performance. Once the test commences, NIST will evaluate implementations
 168 on a first-come-first-served basis and will return results to providers as expeditiously as possible. Providers may submit
 169 revised SDKs to NIST only after NIST provides results for the prior SDK and invites further submission. The frequency with
 170 which a provider may submit SDKs to NIST will depend on the times needed for developer preparation, transmission to
 171 NIST, validation, execution and scoring at NIST, and developer review and decision processes.

172 For the schedule and number of SDKs of each class that may be submitted, see sections 1.10 and 1.11.

173 **1.6. Interim reports**

174 The performance of each SDK will be reported in a "score-card". This will be provided to the participant. While the score
 175 cards may be used by the provider for arbitrary purposes, they are intended to facilitate development. Score cards will

- 176 – be machine generated (i.e. scripted),
- 177 – be provided to participants with identification of their implementation,
- 178 – include timing, accuracy and other performance results,
- 179 – include results from other implementations, but will not identify the other providers,
- 180 – be expanded and modified as revised implementations are tested, and as analyses are implemented,
- 181 – be generated and released asynchronously with SDK submissions,
- 182 – be produced independently of the other status of other providers' implementations,
- 183 – be regenerated on-the-fly, usually whenever any implementation completes testing, or when new analysis is added.

184 NIST does not intend to release these interim test reports publicly. NIST may release such information to the U.S.
 185 Government test sponsors. While these reports are not intended to be made public, NIST can only request that agencies
 186 not release this content.

187 **1.7. Final reports**

188 NIST will publish one or more final public reports. NIST may also

- 189 – publish additional supplementary reports (typically as numbered NIST Interagency Reports),
- 190 – publish in other academic journals,
- 191 – present results at conferences and workshops (typically PowerPoint).

192 Our intention is that the final test reports will publish results for the best-performing implementation from each
 193 participant. Because “best” is ill-defined (accuracy vs. time vs. template size, for example), the published reports may
 194 include results for other implementations. The intention is to report results for the most capable implementations (see
 195 section 1.13, on metrics). Other results may be included (e.g. in appendices) to show, for example, examples of progress
 196 or tradeoffs. **IMPORTANT:** Results will be attributed to the providers.

197 **1.8. Application scenarios**

198 This test will include one-to-many identification tests for video sequences. As described in Table 2, the test is intended to
 199 represent identification applications for face recognition in video.

200 **Table 2 – Subtests supported under the FIVE activity**

#		Video-to-Video	Video-to-Still	Still-to-Video
1.	Aspect	1:N identification of video-to-video	1:N identification of video-to-still	1:N identification of still-to-video
2.	Enrollment dataset	N enrolled video sequences	N enrolled stills	N enrolled video sequences
3.	Prior NIST test references	Equivalent to 1 to N matching in [FRVT 2013]		
4.	Example application	Open-set identification against a central database, e.g. a search of a wanted criminal through a live-video surveillance system at an airport who may attempt to flee the country		
5.	Score or feature space normalization support	Any score or feature based statistical normalization techniques-are applied against enrollment database		
6.	Intended number of subjects	Expected $O(10^2)$ - $O(10^4)$		
7.	Number of images per individual	N/A	Variable, see section 1.12.	Variable, see section 1.12.

201 **1.9. Image source labels**

202 NIST may mix images from different sources in an enrollment set. For example, NIST could combine frontal images and
 203 images with varying poses into a single enrollment dataset. For this reason, in the data structure defined in clause 2.3.3,
 204 each image is accompanied by a "label" which identifies the set-membership images. Legal values for labels are in clause
 205 2.3.2.

206 **1.10. Rules for participation**

207 A participant must properly follow, complete and submit a participation agreement (see Annex A). This must be done
 208 once, not before November 17, 2014. It is not necessary to do this for each submitted SDK. All submitted SDKs must
 209 meet the API requirements as detailed in section 3.

210 **1.11. Number and schedule of submissions**

211 The test is conducted in three phases, as scheduled on page 2. The maximum total (i.e. cumulative) number of
 212 submissions is regulated in Table 3.

213 **Table 3 – Cumulative total number of algorithms**

#	Phase 1	Total over Phases 1 + 2	Total over Phases 1 + 2 + 3
Cumulative total number of submissions	2	3	5 if at least 1 was successfully executed by end Phase 2 2 if zero had been successfully executed by end Phase 2

214 The numbers above may be increased as resources allow.

215 NIST cannot conduct surveys over runtime parameters because NIST must limit the extent to which participants are able
 216 to train on the test data.

217 **1.12. Use of multiple images per person**

218 Some of the proposed datasets includes $K > 2$ images per person for some persons. For video-to-still recognition in this
 219 test, NIST will enroll $K \geq 1$ images under each identity. For still-to-video, the probe will consist of $K \geq 1$ images. Normally
 220 the probe will consist of a single image, but NIST may examine the case that it could consist of multiple images. The
 221 method by which the face recognition implementation exploits multiple images is not regulated: The test seeks to
 222 evaluate developer provided technology for multi-presentation fusion. This departs from some prior NIST tests in which
 223 NIST executed fusion algorithms (e.g. [FRVT2002b]), and sum score fusion, for example, [MINEX]).

224 This document defines a template to be the result of applying feature extraction to a set of $K \geq 1$ images or $K \geq 1$ video
 225 frames. That is, a template contains the features extracted from one or more images or video frames, not generally just
 226 one. An SDK might internally fuse K feature sets into a single representation or maintain them separately - In any case the
 227 resulting proprietary template is contained in a contiguous block of data. All identification functions operate on such
 228 multi-image or multi-frame templates.

229 The number of images per person will depend on the application area:

- 230 – In civil identity credentialing (e.g. passports, driving licenses) the images will be acquired approximately uniformly
 231 over time (e.g. five years for a Canadian passport). While the distribution of dates for such images of a person might
 232 be assumed uniform, a number of factors might undermine this assumption¹.
- 233 – In criminal applications the number of images would depend on the number of arrests². The distribution of dates for
 234 arrest records for a person (i.e. the recidivism distribution) has been modeled using the exponential distribution, but
 235 is recognized to be more complicated. NIST currently estimates that the number of images will never exceed 100.

236 **1.13. Core accuracy metrics**

237 For identification testing, the test will target open-universe applications such as benefits-fraud and watch-lists. It will not
 238 address the closed-set task because it is operationally uncommon.

239 While some one-to-many applications operate with purely rank-based metrics, this test will primarily target score-based
 240 identification metrics. Metrics are defined in Table 4. The analysis will survey over various rank and thresholds³. Plots of
 241 the two error rates, parametric on threshold, will be the primary reporting mechanism.

242 **Table 4 – Summary of accuracy metrics**

Application	Metric
1:N Identification (Video-to-Still)	FPIR = The rate at which unknown subjects are incorrectly associated with any of N enrolled identities. The association will be parameterized on a continuous threshold T .
	FNIR = The rate at which known subjects are incorrectly not associated with the correct enrolled identities. The association will be parameterized on a continuous threshold T , and a candidate rank, R .

243
 244 FPIR will be estimated using probe images or video clips for which there is no enrolled mate. The stability of FPIR at a
 245 fixed threshold under changes to image properties or demographics will be reported.

246 NIST will extend the analysis in other areas, with other metrics, and in response to the experimental data and results.

¹ For example, a person might skip applying for a passport for one cycle (letting it expire). In addition, a person might submit identical images (from the same photography session) to consecutive passport applications at five year intervals.

² A number of distributions have been considered to model recidivism, see "Random parameter stochastic process models of criminal careers." In Blumstein, Cohen, Roth & Visher (Eds.), Criminal Careers and Career Criminals, Washington, D.C.: National Academy of Sciences Press, 1986.

³ Threshold and rank limits are established operationally to limit human labor requirements: One the one side, in a low volume forensic application e.g. investigation of video collected in a convenience store hold-up, or in looking at videos of passengers dis-embarking flights to document an asylum claim, an examiner might be willing to adjudicate $R \gg 1$ candidates with threshold, $T = 0$. At the other end, a high volume watch-list application in which crowded airport concourses are surveilled for bad actors, a high threshold would be used to limit false positive outcomes. In that case, candidate lists will often have zero length. NIST will report metrics appropriate to the "human-automated" hybrid application, and the "lights-out" hits-are-rare use case.

247 **1.14. Generalized accuracy metrics**

248 Under the ISO/IEC 19795-1 biometric testing and reporting standard, a test must account for "failure to acquire" (FTA)
249 and "failure to enroll" (FTE) events (e.g. elective refusal to make a template, or fatal errors). The way these are treated is
250 application-dependent.

251 For identification, the appropriate metrics reported in FIVE will be generalized to include FTA and FTE events.

252 **1.15. Reporting template size**

253 Because template size is influential on storage requirements and computational efficiency, this API supports
254 measurement of template size. NIST will report statistics on the actual sizes of templates produced by face recognition
255 implementations submitted to FIVE. NIST may report statistics on runtime memory usage. Template sizes were reported
256 in the FRVT 2013 test⁴, IREX III test⁵, and the MBE-STILL 2010 test⁶.

257 **1.16. Reporting computational efficiency**

258 As with other tests, NIST will compute and report recognition accuracy. In addition, NIST will also report timing statistics
259 for all core functions of the submitted SDK implementations. This includes feature extraction and 1:N recognition. For an
260 example of how efficiency can be reported, see the final report of the FRVT 2013 test, IREX III test, and the MBE-STILL
261 2010 test.

262 **1.17. Exploring the accuracy-speed trade-space**

263 NIST will explore the accuracy vs. speed tradeoff for face recognition algorithms running on a fixed platform. NIST will
264 report both accuracy and speed of the implementations tested. While NIST cannot force submission of "fast vs. slow"
265 variants, participants may choose to submit variants on some other axis (e.g. "experimental vs. mature")
266 implementations. NIST encourages "fast-less-accurate vs. slow-more-accurate" with a factor of three between the speed
267 of the fast and slow versions.

268 **1.18. Hardware specification**

269 NIST intends to support high performance by specifying the runtime hardware beforehand. There are several types of
270 computer blades that may be used in the testing. The blades are labeled as Dell M905, M910, M605, and M610. The
271 following list gives some details about the hardware of each blade type:

- 272 • Dell M605 - Dual Intel Xeon E5405 2 GHz CPUs (4 cores each)
- 273 • Dell M905 - Quad AMD Opteron 8376HE 2 GHz CPUs (4 cores each)
- 274 • Dell M610 - Dual Intel Xeon X5680 3.3 GHz CPUs (6 cores each)
- 275 • Dell M910 - Dual Intel Xeon X7560 2.3 GHz CPUs (8 cores each)

276 Each CPU has 512K cache. The bus runs at 667 Mhz. The main memory is 192 GB Memory as 24 8GB modules. We
277 anticipate that 16 processes can be run without time slicing.

278 The minimum instruction set across all processors used in the evaluation is specified here⁷. Dependence on instructions
279 not included in the minimum instruction set is prohibited.

280 NIST is requiring use of 64 bit implementations throughout. This will support large memory allocation to support 1:N
281 identification task with image and video frame counts in the millions. For still images, if all templates were to be held in
282 memory, the 192GB capacity implies a limit of ~19KB per template, for a 10 million image enrollment. For video, given
283 the data expectations and the occurrence of faces in the imagery, we anticipate the developers will have sufficient

⁴ See the FRVT 2013 test report: NIST Interagency Report 8009, linked from <http://face.nist.gov/frvt>

⁵ See the IREX III test report: NIST Interagency Report 7836, linked from <http://iris.nist.gov/irex>

⁶ See the MBE-STILL 2010 test report, NIST Interagency Report 7709, linked from <http://face.nist.gov/mbc>

⁷ `cat /proc/cpuinfo` returns `fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pat pse36 clflush mmx fxsr sse sse2 ht syscall nx mmxext fxsr_opt pdpe1gb rdtscp lm 3wext 3dnow constant_tsc nonstop_tsc pni cx16 popcnt lahf_lm cmp_legacy svm extapic cr8_legacy altmovcr8 abm sse4a misalignsse 3dnowprefetch osvw`

284 memory for video templates. Note that while the API allows read access of the disk during the 1:N search, the disk is, of
 285 course, relatively slow.

286 Some of the section 3 API calls allow the implementation to write persistent data to hard disk. The amount of data shall
 287 not exceed 200 kilobytes per enrolled image. NIST will respond to prospective participants' questions on the hardware,
 288 by amending this section.

289 **1.19. Operating system, compilation, and linking environment**

290 The operating system that the submitted implementations shall run on will be released as a downloadable file accessible
 291 from <http://nigos.nist.gov:8080/evaluations/>, which is the 64-bit version of **CentOS 7.0 running Linux kernel 3.10.0.**

292 For this test, Windows machines will not be used. Windows-compiled libraries are not permitted. All software must run
 293 under Linux.

294 NIST will link the provided library file(s) to our C++ language test drivers. Participants are required to provide their library
 295 in a format that is linkable using **g++ version 4.8.2.** The standard libraries are:

```
296 /usr/lib64/libstdc++.so.6.0.19 lib64/libc.so.6 -> libc-2.17.so lib64/libm.so.6 -> libm-2.17.so
```

297 A typical link line might be

```
298 g++ -I. -Wall -m64 -o fivetest fivetest.cpp -L. -lfive_Enron_A_07
```

299 The Standard C++ library should be used for development of the SDKs. The prototypes from the API of this document will
 300 be written to a file "five.h" which will be included via

```
#include <five.h>
```

301 The header files will be made available to implementers at <http://nigos.nist.gov:8080/five>.

302 NIST will handle all input of images via the JPEG and PNG libraries, sourced, respectively from <http://www.iij.org/> and see
 303 <http://libpng.org>.

304 All compilation and testing will be performed on x86 platforms. Thus, participants are strongly advised to verify library-
 305 level compatibility with g++ (on an equivalent platform) prior to submitting their software to NIST to avoid linkage
 306 problems later on (e.g. symbol name and calling convention mismatches, incorrect binary file formats, etc.).

307 Dependencies on external dynamic/shared libraries such as compiler-specific development environment libraries are
 308 discouraged. If absolutely necessary, external libraries must be provided to NIST upon prior approval by the Test Liaison.

309 **1.20. Software and documentation**

310 **1.20.1. SDK Library and platform requirements**

311 Participants shall provide NIST with binary code only (i.e. no source code). Header files (".h") are allowed, but these shall
 312 not contain intellectual property of the company nor any material that is otherwise proprietary. It is preferred that the
 313 SDK be submitted in the form of a single static library file. However, dynamically linked shared library files are permitted.

314 The core library shall be named according to Table 5. Additional shared object library files may be submitted that support
 315 this "core" library file (i.e. the "core" library file may have dependencies implemented in these other libraries).

316 Intel Integrated Performance Primitives (IPP) libraries are permitted if they are delivered as a part of the developer-
 317 supplied library package. It is the provider's responsibility to establish proper licensing of all libraries. The use of IPP
 318 libraries shall not inhibit the SDK's ability to run on CPUs that do not support IPP. Please take note that some IPP
 319 functions are multithreaded and threaded implementations may complicate comparative timing.

320 Access to any GPUs is not permitted.

321 **Table 5 – Implementation library filename convention**

Form	libFIVE_provider_sequence.ending			
Underscore delimited parts of	libFIVE	provider	sequence	ending

FIVE

the filename				
Description	First part of the name, required to be this.	Single word name of the main provider EXAMPLE: Acme	A two digit decimal identifier to start at 00 and increment by 1 every time any SDK is sent to NIST. EXAMPLE: 07	Either .so or .a
Example	libFIVE_Acme_C_07.a			

322

323 NIST will report the size of the supplied libraries.

324 **1.20.2. Configuration and developer-defined data**

325 The implementation under test may be supplied with configuration files and supporting data files. The total size of the
326 SDK, that is all libraries, include files, data files and initialization files shall be less than or equal to 1 073 741 824 bytes =
327 1024³ bytes.

328 NIST will report the size of the supplied configuration files.

329 **1.20.3. Installation and Usage**

330 The SDK must install easily (i.e. one installation step with no participant interaction required) to be tested, and shall be
331 executable on any number of machines without requiring additional machine-specific license control procedures or
332 activation.

333 The SDK shall be installable using simple file copy methods. It shall not require the use of a separate installation program.

334 The SDK shall neither implement nor enforce any usage controls or limits based on licenses, number of executions,
335 presence of temporary files, etc. The submitted implementations shall remain operable with no expiration date.

336 Hardware (e.g. USB) activation dongles are not acceptable.

337 **1.20.4. Hard disk space**

338 FIVE participants should inform NIST if their implementations require more than 100K of persistent storage, per enrolled
339 image on average.

340 **1.20.5. Documentation**

341 Participants shall provide complete documentation of the SDK and detail any additional functionality or behavior beyond
342 that specified here. The documentation must define all (non-zero) developer-defined error or warning return codes.

343 **1.20.6. Modes of operation**

344 Individual SDKs provided shall not include multiple “modes” of operation, or algorithm variations. No switches or options
345 will be tolerated within one library. For example, the use of two different “coders” by a feature extractor must be split
346 across two separate SDK libraries, and two separate submissions.

347 **1.21. Runtime behavior**

348 **1.21.1. Interactive behavior**

349 The SDK will be tested in non-interactive “batch” mode (i.e. without terminal support). Thus, the submitted library shall
350 not use any interactive functions such as graphical user interface (GUI) calls, or any other calls which require terminal
351 interaction e.g. reads from “standard input”.

352 **1.21.2. Error codes and status messages**

353 The SDK will be tested in non-interactive “batch” mode, without terminal support. Thus, the submitted library shall run
 354 quietly, i.e. it should not write messages to "standard error" and shall not write to “standard output”. An SDK may write
 355 debugging messages to a log file - the name of the file must be declared in documentation.

356 **1.21.3. Exception Handling**

357 The application should include error/exception handling so that in the case of a fatal error, the return code is still
 358 provided to the calling application.

359 **1.21.4. External communication**

360 Processes running on NIST hosts shall not side-effect the runtime environment in any manner, except for memory
 361 allocation and release. Implementations shall not write any data to external resource (e.g. server, file, connection, or
 362 other process), nor read from such. If detected, NIST will take appropriate steps, including but not limited to, cessation of
 363 evaluation of all implementations from the supplier, notification to the provider, and documentation of the activity in
 364 published reports.

365 **1.21.5. Stateless behavior**

366 All components in this test shall be stateless, except as noted. This applies to face detection, feature extraction and
 367 matching. Thus, all functions should give identical output, for a given input, independent of the runtime history. NIST
 368 will institute appropriate tests to detect stateful behavior. If detected, NIST will take appropriate steps, including but not
 369 limited to, cessation of evaluation of all implementations from the supplier, notification to the provider, and
 370 documentation of the activity in published reports.

371 **1.22. Threaded computations**

372 Table 6 shows the limits on the numbers of threads a face recognition implementation may use. **Threading is prohibited**
 373 **for feature extraction and search**, because NIST will parallelize the test by dividing the workload across many cores and
 374 many machines. For the finalization function, if threading is used, NIST requires the provider to disclose the maximum
 375 number of threads to us.

376 **Table 6 – Number of threads allowed for each function**

Function	Video
Feature extraction	1
Finalize enrollment (before 1:N)	1 ≤ T ≤ 16
Identification	1

377 To expedite testing NIST will run up to P >> 1 processes concurrently. NIST's calling applications are single-threaded.

378 **1.23. Time limits**

379 The elemental functions of the implementations shall execute under the time constraints of Table 7. These time limits
 380 apply to the function call invocations defined in section 3. Assuming the times are random variables, NIST cannot regulate
 381 the maximum value, so the time limits are 90-th percentiles. This means that 90% of all operations should take less than
 382 the identified duration.

383 The time limits apply per image or video frame. When K images of a person are present or K frames are in a video clip,
 384 the time limits shall be increased by a factor K.

385 **Table 7 – Processing time limits in milliseconds**

Function	Video-to-Video	Video-to-Still	Still-to-Video
Feature extraction enrollment	5 * 1500 per video frame (1 core)	1500 per image (1 core)	5 * 1500 per video frame (1 core)
Feature extraction for identification	5 * 1500 per video frame (1 core)	1500 per image (1 core)	5 * 1500 per video frame (1 core)

386 For video: the multiple of K=5 is a notional average of the number of persons expected in any given frame. This figure is
 387 proportionally unreliable for any given sample.

388 While there is no time limit for the enrollment finalization procedure, NIST will report the execution duration.

389 **1.24. Test datasets**

390 This section is under development. The data has, in some cases, been estimated from initial small partitions. The
 391 completion of this section depends on further work. The information is subject to change. We intend to update this
 392 section as fully as possible.

393 NIST is likely to use other datasets, in addition.

394 **Table 8 – Main video corpora (others will be used)**

	Dataset P	Dataset T	Dataset B	Other datasets - Undisclosed
Collection, environment	Indoor public space with individuals walking mostly toward cameras as could occur on a transit terminal		Television footage, indoor and outdoor	
Number of individuals in field of view	Multiple, usually below 20 many not fully visible but usually more than 1.		Few, most often 1, occasionally others in background	
View angle	Various pitch due to different heights of camera installation, some yaw also due to subject behavior		Pitch variation present, but yaw angles vary more due to subject behavior	
Video frame size	1920 x 1080	Various	Various	
Eye to eye distance	10-100 pixels	10-150 pixels	10-120	
	The above values are guidelines; exceptions will inevitably occur in large datasets.			
Camera properties	Consumer-grade video	Professional-grade video	Professional-grade video cameras	
Camera motion	Fixed geometry, fixed optics		Usually camera is still or slowly panning or zooming	
Frames per second	24	Up to 30	Up to 30	
Similar composition to	Compare to the iLids data but with higher spatial resolution on the face.		Similar to YouTubeFaces in that typically one subject is present and in the foreground	
Accompanying stills	Yes, for video-to-still and still-to-video searches, high-resolution stills approximating ISO/IEC 19794-5 are available. In addition, off-angle images exist with many combinations of pitch and yaw. In addition, less formal “social-media” like stills are available also. Various galleries will be formed from these images. Images for which interocular distance exceeds 240 pixels will be downsized.		Stills usually resemble frames from the video. ISO/IEC 19794-5 images are not usually available.	

395
 396 NIST does not know the minimum and maximum numbers of persons appearing in video sequences. Moreover, NIST will
 397 apply the algorithms to other databases. The maximum number of frames in a video sequence will be limited by the
 398 duration of the sequence. NIST expects to use sequences whose duration extends from a few seconds to a few minutes

399 Some notes regarding the video data:

- 400 – NIST does not anticipate using interlaced video.
- 401 – The videos are contiguous in time, without interruptions.
- 402 – Some sequences exist at much higher frame rates. NIST will examine whether this offers benefit.
- 403 – Some of the datasets were collected using consumer-grade cameras capturing video in standard formats while others
- 404 were collected using professional-grade cameras captured in modern proprietary video codecs.

405 In some videos, the scenes capture people walking towards the camera. Occasionally, there are people walking in various
 406 transverse directions including people walking away from the camera. The cameras have varying pitch angles ranging
 407 from 0 degrees (frontal) to higher values. The depth of scene varies between the cameras such that the sizes of the faces
 408 vary, with the following:

- 409 — Eye-to-eye distances range from approximately 10 pixels to 120 pixels
- 410 — Amount of time a face is fully visible in a scene can vary from approximately 0 to 30 seconds
- 411 — Some of the captures include non-uniform lighting due to light coming through adjacent windows

412
 413 Please note that the properties stated above may not hold for all datasets that might be employed in FIVE.

414 **1.25. Ground truth integrity**

415 Some of the test databases will be derived from operational systems. They may contain ground truth errors in which

- 416 — a single person is present under two different identifiers, or
- 417 — two persons are present under one identifier, or
- 418 — in which a face is not present in the image.

419 If these errors are detected, they will be removed. NIST will use aberrant scores (high impostor scores, low genuine
 420 scores) to detect such errors. This process will be imperfect, and residual errors are likely. For comparative testing,
 421 identical datasets will be used and the presence of errors should give an additive increment to all error rates. For very
 422 accurate implementations this will dominate the error rate. NIST intends to attach appropriate caveats to the accuracy
 423 results. For prediction of operational performance, the presence of errors gives incorrect estimates of performance.

424 **2. Data structures supporting the API**

425 **2.1. Overview**

426 This section describes the API for the face recognition in video applications described in section 1.8. All SDK's submitted
 427 to FIVE shall implement the functions required in Section 3.

428 **2.2. Requirement**

429 FIVE participants shall submit an SDK which implements the relevant C++ prototyped interfaces of clause 3. C++ was
 430 chosen in order to make use of some object-oriented features.

431 **2.3. File formats and data structures**

432 **2.3.1. Overview**

433 In this test, an individual is represented by $K \geq 1$ two-dimensional facial images, and by subject and image-specific
 434 metadata.

435 **2.3.2. Dictionary of terms describing images and MULTIFACES**

436 Images will be accompanied by one of the labels given in Table 9. Face recognition implementations submitted to FIVE
 437 should tolerate images of any category.

438 **Table 9 – Labels describing types of images**

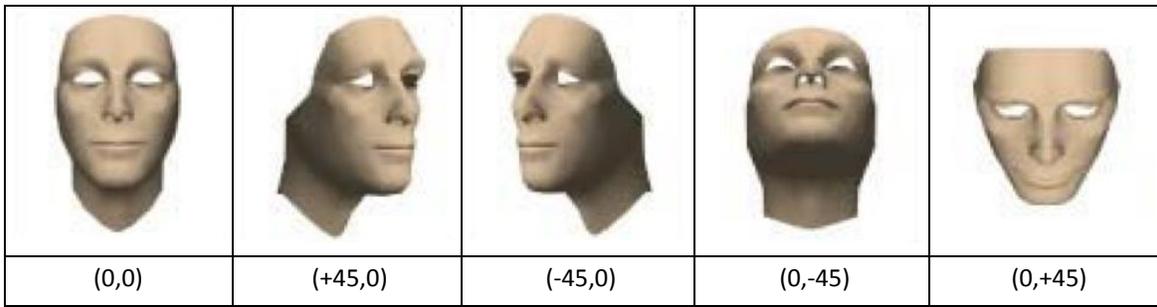
	Label as C++ string	Meaning	Yaw (degrees)	Pitch (degrees)
1.	"unknown"	Either the label is unknown or unassigned.		
2.	"uncontrolled"	Any illumination, pose is unknown and could be frontal		
3.	"FF"	Full frontal	0	0
4.	"FD"	Face down	0	10 to 40

FIVE

5.	"FU"	Face up	0	-10 to -40
6.	"QL"	Quarter left	-10 to -45	0
7.	"QR"	Quarter right	10 to 45	0
8.	"HL"	Half left	-46 to -80	0
9.	"HR"	Half right	46 to 80	0
10.	"PL"	Profile left	-90	0
11.	"PR"	Profile right	90	0
12.	"QLU"	Quarter left up	-10 to -45	-10 to -40
13.	"QRU"	Quarter right up	10 to 45	-10 to -40
14.	"HLU"	Half left up	-46 to -80	-10 to -40
15.	"HRU"	Half right up	46 to 80	-10 to -40
16.	"HLD"	Half left down	-46 to -80	10 to 40
17.	"HRD"	Half right down	46 to 80	10 to 40

439 **Figure 2** provides examples of pose angles and their encoding (yaw, pitch) as specified in the ISO/IEC 19794-5 [ISO], with
 440 yaw angle defined as the rotation in degrees about the y-axis (vertical axis) and pitch angle defined as the rotation in
 441 degrees about the x-axis (horizontal axis).

442



443

Figure 2 – Examples of pose angles and their encodings (yaw, pitch)

444

NOTE 1: We do not intend to deliberately include non-face images in this test.

445

NOTE 2: **MULTIFACES** will contain face images of only one person.

446

447

A **MULTIFACE** (see Table 12) will be accompanied by one of the labels given in Table 10. Face recognition implementations submitted to FIVE should tolerate **MULTIFACES** of any category.

449

Table 10 – Labels describing types of MULTIFACES

	Label as C++ string	Meaning
1.	"FRONTAL"	All ONEFACE s contain nominally frontal images and are labeled "FF".
2.	"MULTIPOSE"	Each ONEFACE is labeled with one of the following: "FF", "FD", "FU", "QL", "QR", "HL", "HR", "PL", "PR", "QLU", "QRU", "HLU", "HRU", "HLD", "HRD".
3.	"INFORMAL"	All ONEFACE s contain informal images that are labeled "uncontrolled".
4.	"UNKNOWN"	Each ONEFACE is labeled with one of the labels from Table 9, including possibly "unknown" or "uncontrolled".

450

451

2.3.3. Data structures for encapsulating multiple still images

452

The standardized formats for facial images are the ISO/IEC 19794-5:2005 and the ANSI/NIST ITL 1-2007 type 10 record.

453

The ISO record can store multiple images of an individual in a standalone binary file. In the ANSI/NIST realm, K images of an individual are usually represented as the concatenation of one Type 1 record + K Type 10 records. The result is usually stored as an EFT file.

455

456 An alternative method of representing K images of an individual is to define a structure containing an image filename and
 457 metadata fields. Each file contains a standardized image format, e.g. PNG (lossless) or JPEG (lossy).

458 **2.3.4. Class for encapsulating a single face image**

459 **Table 11 – ONEFACE class**

	C++ code fragment	Remarks
1.	<code>class ONEFACE</code>	
2.	<code>{</code> <code>private:</code>	
3.	<code>uint16_t imageWidth;</code>	Number of pixels horizontally
4.	<code>uint16_t imageHeight;</code>	Number of pixels vertically
5.	<code>uint16_t imageDepth;</code>	Number of bits per pixel. Legal values are 8 and 24.
6.	<code>uint8_t format;</code>	Flag indicating native format of the image as supplied to NIST 0x01 = JPEG (i.e. compressed data) 0x02 = PNG (i.e. never compressed data)
7.	<code>uint8_t *data;</code>	Pointer to raster scanned data. Either RGB color or intensity. If image_depth == 24 this points to 3WH bytes RGBRGBRGB... If image_depth == 8 this points to WH bytes I I I I I I I I
8.	<code>std::string description;</code>	Single description of the image. The allowed values for this string are given in Table 9.
9.	<code>public:</code> <code>//getter/setter methods</code>	
10.	<code>};</code>	

460 **2.3.5. Class for encapsulating a set of face images from a single person**

461 **Table 12 – MULTIFACE class**

	C++ code fragment	Remarks
1.	<code>class MULTIFACE</code> <code>{</code> <code>private:</code>	Vector containing F pre-allocated face images of the same person. The number of items stored in the vector is accessible via the vector::size() function.
2.	<code>std::vector<ONEFACE> faces;</code> <code>std::string description;</code>	Single description of the vector of ONEFACES. The allowed values for this string are given in Table 10.
3.	<code>public:</code> <code>//getter/setter methods</code>	
4.	<code>};</code>	

462 **2.3.6. Dictionary of terms describing ONEVIDEOS**

463 A ONEVIDEO will be accompanied by one of the labels given in Table 13, describing the density of people in the video
 464 frames. Face recognition implementations submitted to FIVE should tolerate ONEVIDEOS of any category.

465 **Table 13 – Labels describing the density of people in the video frames**

	Label as C++ string	Meaning
1.	"SINGLE"	All of the video frames contain one and only one person. Such video might arise from a TV interview or speech. An algorithm should produce one template from the ONEVIDEO.
2.	"UNKNOWN"	Video frames can contain zero or more people in each frame. Such video might arise in a surveillance clip. The number of templates to return would be a random variable.

466 **2.3.7. Class for encapsulating a video sequence**

467 **Table 14 – ONEVIDEO Class**

	C++ code fragment	Remarks
--	-------------------	---------

1.	class ONEVIDEO	
2.	{	
3.	private:	
4.	uint16_t frameWidth;	Number of pixels horizontally of all frames
5.	uint16_t frameHeight;	Number of pixels vertically of all frames
6.	uint8_t frameDepth;	Number of bits per pixel for all frames. Legal values are 8 and 24.
7.	uint16_t framesPerSec;	The frame rate of the video sequence. If this value is 0, the frames are sampled irregularly and perhaps infrequently from the parent video clip (e.g. manually selected frames, or just the I-frames).
8.	std::string peopleDensity;	Single description of the density of people in the video frames. The allowed values for this string are given in Table 13.
9.	std::vector<const uint8_t*> data;	Vector of pointers to data from each frame in the video sequence. The number of frames (i.e. size of the vector) can be obtained by calling vector::size(). The i-th entry in data (ie. data[i]) points to frame_width x frame_height pixels of data for the i-th frame.
10.	public:	
11.	//getter/setter methods	
12.	};	

468 **2.3.8. Class representing a pair of eye coordinates**

469 The data structure for reporting person locations in video appears in Table 15. The coordinates may be useful to NIST for
 470 relating spatial location to recognition success during our analysis.

471 **Table 15 – EYEPAIR Class**

	C++ code fragment	Remarks
1.	class EYEPAIR	
2.	{	
3.	private:	
4.	bool isSet;	If the eye coordinates have been computed and assigned successfully, this value should be set to true, otherwise it should be set to false.
5.	int16_t xLeft; int16_t yLeft;	X and Y coordinate of the center of the subject's left eye. Out-of-range values (e.g. x < 0 or x >= width) indicate the implementation believes the eye center is outside the image.
6.	int16_t xRight; int16_t yRight;	X and Y coordinate of the center of the subject's right eye. Out-of-range values (e.g. x < 0 or x >= width) indicate the implementation believes the eye center is outside the image.
7.	uint16_t frameNum	For video: the frame number that corresponds to the video frame from which the eye coordinates were generated. (ie. the i-th frame from the video sequence). This field should not be set for eye coordinates for a single still image.
8.	public:	
9.	//getter/setter methods	
10.	};	

472 **2.3.9. Data type for representing a person’s trajectory via eye coordinates from a video sequence**

473 **Table 16 – PersonTrajectory typedef**

	C++ code fragment	Remarks
1.	typedef std::vector< EYEPAIR > PersonTrajectory;	Vector of EYEPAIR (see 2.3.8) objects for video frames where eyes were detected. This data structure should store eye coordinates for each video frame where eyes were detected for a particular person. For video frames where the person’s eyes were not detected, the SDK shall not add an EYEPAIR to this data structure. If a face can be detected, but not the eyes, the implementation should nevertheless fill this data structure with (x,y) _{LEFT} == (x,y) _{RIGHT} representing some point on the center of the face.

474 **2.3.10. Class for representing a person from a video sequence or an image**

475 **Table 17 – PERSONREP Class**

	C++ code fragment	Remarks
1.	class PERSONREP	
2.	{	
	private:	
3.	PersonTrajectory eyeCoordinates;	Data structure for capturing eye coordinates
4.	PersonTemplate proprietaryTemplate;	PersonTemplate is a wrapper to a uint8_t* for capturing proprietary template data representing a person from a video sequence or an image.
5.	public:	
6.	PERSONREP(const uint64_t inSize);	The constructor takes a size parameter and allocates memory of <i>inSize</i> . getPersonTemplatePtr() should be called to access the newly allocated memory for SDK manipulation. Please note that this class will take care of all memory allocation and de-allocation of its own memory. The SDK shall not de-allocate memory created by this class.
7.	void pushBackEyeCoord(const EYEPAIR &eyes);	This function should be used to add EYEPAIRS for the video frames or images where eye coordinates were detected.
8.	uint8_t* getPersonTemplatePtr();	This function returns a uint8_t* to the template data.
9.	uint64_t getPersonTemplateSize() const;	This function returns the size of the template data.
10.	//... getter methods, copy constructor,	
	//... assignment operator	
11.	};	

476 **2.3.11. Class for result of an identification search**

477 All identification searches shall return a candidate list of a NIST-specified length. The list shall be sorted with the most
478 similar matching entries list first with lowest rank.

479 **Table 18 – CANDIDATE Class**

	C++ code fragment	Remarks
1.	class CANDIDATE	
2.	{	
	private:	
3.	bool isSet	If the candidate is valid, this should be set to true. If the candidate computation failed, this should be set to false.
4.	uint32_t templateId;	The Template ID integer from the enrollment database manifest defined in clause 2.3.6.
5.	double similarityScore;	Measure of similarity between the identification template and the enrolled candidate. Higher scores mean more likelihood that the samples are of the same person. An algorithm is free to assign any value to a candidate. The distribution of values will have an impact on the appearance of a plot of false-negative and false-positive identification rates.
6.	public:	
	//getter/setter methods	
7.	};	

480 **2.3.12. Data type for representing a list of results of an identification search**

481 **Table 19 – CANDIDATELIST typedef**

	C++ code fragment	Remarks
1.	typedef std::vector<CANDIDATE> CANDIDATELIST ;	A vector containing objects of CANDIDATES . The CANDIDATE class is defined in section 2.3.11.

482

483 **2.3.13. Class representing return code values**

484 **Table 20 – ReturnCode class**

	C++ code fragment	Remarks
	class ReturnCode { public:	
1.	typedef enum	
2.	{	
3.	Success=0,	Success
4.	MissingConfig=1,	The configuration data is missing or unreadable
5.	EnrollDirFailed=2,	An operation on the enrollment directory failed
6.	InitNumData=3,	The SDK can't support the number of images or videos
7.	InitBadDesc=4,	The image descriptions are unexpected or unusable
8.	RefuseInput=5,	Elective refusal to process this kind of input (ONEVIDEO or MULTIFACE)
9.	FailExtract=6,	Involuntary failure to extract features
10.	FailTempl=7,	Elective refusal to produce a template
11.	FailParse=8,	Cannot parse input data
12.	FinInputData=9,	Cannot locate input data
13.	FinTemplFormat=10,	One or more template files are in an incorrect format
14.	IdBadTempl=11,	The input template was defective
15.	ImgSizeNotSupported=12,	Size of input image/frame not supported
16.	Vendor=13	Vendor-defined failure
17.	} Status;	
18.	ReturnCode(const Status inStatus);	Constructor that takes an input parameter of a Status enum value. All of the functions that need to be implemented for the Video API return an instantiation of a ReturnCode object with a valid status value passed in as a parameter.
19.	Status getStatus() const;	Getter method to return status value
20.	private:	
21.	Status status;	Member variable for storing status
22.	};	

485 **2.4. File structures for enrolled template collection**

486 For still image enrollment, an SDK converts a **MULTIFACE** into a template using the
 487 ImageEnrollment::generateEnrollmentTemplate() function of section 3.3.5.2. For video enrollment, an SDK converts a
 488 **ONEVIDEO** into one or more templates, using the VideoEnrollment::generateEnrollmentTemplate() of section 3.3.1.2. To
 489 support the identification functions, NIST will concatenate enrollment templates into a single large file. This file is called
 490 the EDB (for enrollment database). The EDB is a simple binary concatenation of proprietary templates. There is no
 491 header. There are no delimiters. The EDB may extend to hundreds of gigabytes in length.

492 This file will be accompanied by a manifest; this is an ASCII text file documenting the contents of the EDB. The manifest
 493 has the format shown as an example in Table 21. If the EDB contains N templates, the manifest will contain N lines. The
 494 fields are space (ASCII decimal 32) delimited. There are three fields, all containing numeric integers. Strictly speaking, the
 495 third column is redundant.

496 **Table 21 – Enrollment dataset template manifest**

Field name	Template ID	Template Length	Position of first byte in EDB
Datatype required	Unsigned decimal integer	Unsigned decimal integer	Unsigned decimal integer
Datatype length required	4 bytes	4 bytes	8 bytes
Example lines of a manifest file appear to the right. Lines 1, 2, 3 and N appear.	90201744	1024	0
	163232021	1536	1024
	7456433	512	2560
	...		

	183838	1024	307200000
--	--------	------	-----------

497
498 The EDB scheme avoids the file system overhead associated with storing millions of individual files.

499 **3. API Specification**

500 **3.1.1. Definitions**

501 As shown in Table 22, the video API supports 1:N identification of video-to-video, video-to-still image, and still image-to-
502 video. The following hold:

- 503 – A still image is a picture of one and only one person. One or more such images are presented to the implementation
504 using a **MULTIFACE** data structure.
- 505 – A video is a sequence of $F \geq 1$ frames containing $P \geq 0$ persons.
- 506 – A frame is 2D still image containing $P \geq 0$ persons.
- 507 – Any person might be present in $0 \leq f \leq F$ frames, and their presence may be non-contiguous (e.g. due to occlusion).
- 508 – Different videos contain different numbers of frames and people.
- 509 – A **ONEVIDEO** container is used to represent a video. It contains a small header and pointers to F frames.
- 510 – Any person found in a video is represented by proprietary template (feature) data contained with a **PERSONREP** data
511 structure. A proprietary template contains information from one or more frames. Internally, it might embed multiple
512 traditional still-image templates, or it might integrate feature data by tracking a person across multiple frames.
- 513 – A **PERSONREP** structure additionally contains a trajectory indicating the location of the person in each frame.

514
515 All of the code for the classes needed to implement the video API will be provided to implementers at
516 <http://nigos.nist.gov:8080/five>. A single sample video has been made available at the same link. The sample video is
517 only approximately representative of the scene and is not an extraction from the actual video data that will be used in the
518 evaluation. It is only intended to illustrate similarities in terms of camera placement relative to the subject and people
519 behavior. It is not intended to represent the optical properties of the actual imaging systems, particularly the spatial
520 sampling rate, nor the compression characteristics.

521

522 **Table 22 – API implementation requirements for FIVE**

Function	Video-to-video	Still-to-video	Video-to-still
Enroll	Videos	Videos	Stills
Enrollment input datatype	ONEVIDEO	ONEVIDEO	MULTIFACE
Enrollment datatype	PERSONREP	PERSONREP	PERSONREP
Search	Video	Still	Video
Search input datatype	ONEVIDEO	MULTIFACE	ONEVIDEO
Search datatype	PERSONREP	PERSONREP	PERSONREP
Search result	CANDIDATELIST	CANDIDATELIST	CANDIDATELIST
API requirements	3.3.1 + 3.3.2 + 3.3.3 + 3.3.4	3.3.1 + 3.3.2 + 3.3.7 + 3.3.4	3.3.5 + 3.3.6 + 3.3.3 + 3.3.8

523 **3.1.1.1. Video-to-video**

524 Video-to-video identification is the process of enrolling N videos and then searching the enrollment database with a
525 search video. During identification, the SDK shall return a set of indices of candidate videos that contain people who
526 appear in the search video.

- 527 – N templates will be generated from M enrollment videos. If no people appear in the videos, N will be 0. If many
528 people appear in each video, we'd expect $N > M$.
- 529 – The N templates will be concatenated and finalized into a proprietary enrollment data structure.
- 530 – A **ONEVIDEO** will be converted to $S \geq 0$ identification template(s) based on the number of people detected in the
531 video.

- 532 – Each identification template generated will be searched against the enrollment database of templates generated
- 533 from the M input videos.
- 534 – We anticipate that the same person may appear in more than one enrolled video.

535 **3.1.1.2. Still image-to-video**

536 Still image-to-video identification is the process of enrolling N videos and then searching the enrollment database with a
 537 template produced from a **MULTIFACE** as follows:

- 538 – N templates will be generated from $1 < M \leq N$ enrollment videos.
- 539 – The N templates will be concatenated and finalized into a proprietary enrollment data structure.
- 540 – A **MULTIFACE** (still image) will be converted to an identification template.
- 541 – The identification template will be searched against the enrollment database of N templates.
- 542 – We anticipate that the same person may appear in more than one enrolled video.

543 **3.1.1.3. Video-to-still image**

544 Video-to-still image identification is the process of enrolling N **MULTIFACES** (see Table 12) and then searching the
 545 enrollment database with templates from persons found in a video as follows

- 546 – N templates will be generated from N still-image **MULTIFACES**.
- 547 – The N templates will be concatenated and finalized into a proprietary enrollment data structure.
- 548 – A **ONEVIDEO** will be converted to $S \geq 0$ identification template(s) based on the number of people detected in the
 549 video.
- 550 – Each of the S identification templates will be searched separately against the enrollment database of N templates.

551 **3.2. 1:N Identification**

552 **3.2.1. Overview**

553 The 1:N application proceeds in two phases, enrollment and identification. The identification phase includes separate
 554 pre-search feature extraction stage, and a search stage.

555 The design reflects the following *testing* objectives for 1:N implementations.

- support distributed enrollment on multiple machines, with multiple processes running in parallel
- allow recovery after a fatal exception, and measure the number of occurrences
- allow NIST to copy enrollment data onto many machines to support parallel testing
- respect the black-box nature of biometric templates
- extend complete freedom to the provider to use arbitrary algorithms
- support measurement of duration of core function calls
- support measurement of template size

556 **Table 23 – Procedural overview of the identification test**

Phase	#	Name	Description	Performance Metrics to be reported by NIST

FIVE

Enrollment	E1	Initialization	<p>For still image enrollment, give the implementation advance notice of the number of individuals and images that will be enrolled.</p> <p>Give the implementation the name of a directory where any provider-supplied configuration data will have been placed by NIST. This location will otherwise be empty.</p> <p>The implementation is permitted read-write-delete access to the enrollment directory during this phase. The implementation is permitted read-only access to the configuration directory.</p> <p>After enrollment, NIST may rename and relocate the enrollment directory - the implementation should not depend on the name of the enrollment directory.</p>	
	E2	Parallel Enrollment	<p>For still image enrollment, for each of N individuals, pass multiple images to the implementation for conversion to a combined template. For video enrollment, for each of M video clips, pass multiple video frames to the implementation for generation of N templates, based on the number of people detected in the videos. The implementation will return a template to the calling application.</p> <p>The implementation is permitted read-only access to the enrollment directory during this phase. NIST's calling application will be responsible for storing all templates as binary files. These will not be available to the implementation during this enrollment phase.</p> <p>Multiple instances of the calling application may run simultaneously or sequentially. These may be executing on different computers. For still image enrollment, the same person will not be enrolled twice.</p>	<p>Statistics of the times needed to enroll an individual or video clip.</p> <p>Statistics of the sizes of created templates.</p> <p>The incidence of failed template creations.</p>
	E3	Finalization	<p>Permanently finalize the enrollment directory. This supports, for example, dis-interleaving of internal feature representations, writing of a manifest, indexing, tree construction, computation of statistical information over the enrollment dataset, and adaptation of the representation.</p> <p>The implementation is permitted read-write-delete access to the enrollment directory during this phase.</p>	<p>For still image enrollment, size of the enrollment database as a function of population size N and the number of images.</p> <p>Duration of this operation. The time needed to execute this function shall be reported with the preceding enrollment times.</p>
Pre-search	S1	Initialization	<p>Tell the implementation the location of an enrollment directory. The implementation could look at the enrollment data.</p> <p>The implementation is permitted read-only access to the enrollment directory during this phase.</p>	<p>Statistics of the time needed for this operation.</p>
	S2	Template preparation	<p>For each probe, create a template from a set of input images or one or more templates from a set of video clips. This operation will generally be conducted in a separate process invocation to step S2.</p> <p>The implementation is permitted no access to the enrollment directory during this phase.</p> <p>The result of this step is a search template.</p>	<p>Statistics of the time needed for this operation.</p> <p>Statistics of the size of the search template(s).</p>
Search	S3	Initialization	<p>Tell the implementation the location of an enrollment directory. The implementation should read all or some of the enrolled data into main memory, so that searches can commence.</p> <p>The implementation is permitted read-only access to the enrollment directory during this phase.</p>	<p>Statistics of the time needed for this operation.</p>
	S4	Search	<p>A template or multiple templates is searched against the enrollment database.</p> <p>The implementation is permitted read-only access to the enrollment directory during this phase.</p>	<p>Statistics of the time needed for this operation.</p> <p>Accuracy metrics - Type I + II error rates.</p>

				Failure rates.
--	--	--	--	----------------

557 **3.3. Interfaces**

558 **3.3.1. The VideoEnrollment Interface**

559 The abstract class VideoEnrollment must be implemented by the SDK developer in a class named exactly
 560 SdkVideoEnrollment. The processing that takes place during each phase of the test is done via calls to the methods
 561 declared in the interface as pure virtual, and therefore is to be implemented by the SDK. The test driver will call these
 562 methods, handling all return values.

	C++ code fragment	Remarks
1.	class VideoEnrollment	
2.	{ public:	
3.	virtual ReturnCode initialize(const string &configDir, const string &enrollDir, const uint32_t numVideos) = 0 ;	Initialize the enrollment session.
4.	virtual ReturnCode generateEnrollmentTemplate(const ONEVIDEO &inputVideo, vector< PERSONREP > &enrollTemplates) = 0;	Generate enrollment template(s) for the persons detected in the input video. This function takes a ONEVIDEO (see 2.3.6) as input and populates a vector of PERSONREP (see 2.3.10) with the number of persons detected from the video sequence. The implementation could call vector::push_back to insert into the vector.
5.	// Destructor	
6.	};	

563 **3.3.1.1. Initialization of the video enrollment session**

564 Before any enrollment feature extraction calls are made, the NIST test harness will call the initialization below for video-
 565 to-video and still image-to-video.

566 **Table 24 – VideoEnrollment::initialize**

Prototype	ReturnCode initialize(const string &configDir, const string &enrollDir, const uint32_t numVideos);	
		Input
		Input
		Input
Description	This function initializes the SDK under test and sets all needed parameters. This function will be called N=1 times by the NIST application immediately before any $M \geq 1$ calls to generateEnrollmentTemplate. Caution: The implementation should tolerate execution of $P > 1$ processes on the one or more machines each of which may be reading and writing to this same enrollment directory in parallel. File locking or process-specific temporary filenames would be needed to safely write content in the enrollDir.	
Input Parameters	configDir	A read-only directory containing any developer-supplied configuration parameters or run-time data files.
	enrollDir	The directory will be initially empty, but may have been initialized and populated by separate invocations of the enrollment process. When this function is called, the SDK may populate this folder in any manner it sees fit. Permissions will be read-write-delete.
	numVideos	The total number of videos that will be passed to the SDK for enrollment.
Output Parameters	none	
ReturnCode	Success	Success
	MissingConfig	The configuration data is missing, unreadable, or in an unexpected format.
	EnrollDirFailed	An operation on the enrollment directory failed (e.g. permission, space).
	InitNumData	The SDK cannot support the number of videos.
	Vendor	Vendor-defined failure

567 **3.3.1.2. Video enrollment**

568 A **ONEVIDEO** is converted to enrollment template(s) for each person detected in the **ONEVIDEO** using the function below.

569 **Table 25 – VideoEnrollment::generateEnrollmentTemplate**

Prototypes	ReturnCode generateEnrollmentTemplate(const ONEVIDEO &inputVideo, std::vector< PERSONREP > &enrollTemplates);	
	Input	Output
Description	<p>This function takes a ONEVIDEO, and outputs a vector of PERSONREP objects. If the function executes correctly (i.e. returns a ReturnCode::Success exit status), the NIST calling application will store the template. The NIST application will concatenate the templates and pass the result to the enrollment finalization function. For a video in which no persons appear, a valid output is an empty vector (i.e. size() == 0).</p> <p>If the function gives a non-zero exit status:</p> <ul style="list-style-type: none"> – If the exit status is ReturnCode::FailParse, NIST will debug, otherwise – the test driver will ignore the output template (the template may have any size including zero) – the event will be counted as a failure to enroll. Such an event means that this person can never be identified correctly. <p>IMPORTANT. NIST's application writes the template to disk. The implementation must not attempt writes to the enrollment directory (nor to other resources). Any data needed during subsequent searches should be included in the template, or created from the templates during the enrollment finalization function.</p>	
Input Parameters	inputVideo	An instance of a Table 14 class.
Output Parameters	enrollTemplates	For each person detected in the ONEVIDEO , the function shall identify the person's estimated eye centers for each video frame where the person's eye coordinates can be calculated. The eye coordinates shall be captured in the PERSONREP .eyeCoordinates variable, which is a vector of EYEPAIR objects. The frame number from the video of where the eye coordinates were detected shall be captured in the EYEPAIR .frameNum variable for each pair of eye coordinates. In the event the eye centers cannot be calculated (ie. the person becomes out of sight for a few frames in the video), the SDK shall not store an EYEPAIR for those frames.
ReturnCode	Success	Success
	RefuseInput	Elective refusal to process this kind of ONEVIDEO
	FailExtract	Involuntary failure to extract features (e.g. could not find face in the input-image)
	FailTempl	Elective refusal to produce a template (e.g. insufficient pixels between the eyes)
	FailParse	Cannot parse input data (i.e. assertion that input record is non-conformant)
	ImgSizeNotSupported	Input image/frame size too small or large
	Vendor	Vendor-defined failure. Failure codes must be documented and communicated to NIST with the submission of the implementation under test.

570 **3.3.2. The VideoFinalize Interface**

571 The abstract class VideoFinalize must be implemented by the SDK developer in a class named exactly SdkVideoFinalize.
 572 The finalize function in this class takes the name of the top-level directory where enrollment database (EDB) and its
 573 manifest have been stored. These are described in section 2.3.6. The enrollment directory permissions will be read +
 574 write.

	C++ code fragment	Remarks
1.	class VideoFinalize	
2.	{ public:	
3.	virtual ReturnCode finalize(const string &enrollDir, const string &edbName, const string &edbManifest) = 0;	This function supports post-enrollment developer-optional book-keeping operations and statistical processing. The function will generally be called in a separate process after all the enrollment processes are complete.
4.	// Destructor	

```
5. };
```

575 **3.3.2.1. Finalize video enrollment**

576 After all templates have been created, the function of Table 26 will be called. This freezes the enrollment data. After this
 577 call the enrollment dataset will be forever read-only. This API does not support interleaved enrollment and search
 578 phases.

579 The function allows the implementation to conduct, for example, statistical processing of the feature data, indexing and
 580 data re-organization. The function may alter the file structure. It may increase or decrease the size of the stored data.
 581 No output is expected from this function, except a return code.

582 **Table 26 – VideoFinalize::finalize**

Prototypes	ReturnCode finalize (
	const string &enrollDir,	Input
	const string &edbName,	Input
	const string &edbManifest);	Input
Description	<p>This function takes the name of the top-level directory where enrollment database (EDB) and its manifest have been stored. These are described in section 2.3.6. The enrollment directory permissions will be read + write.</p> <p>The function supports post-enrollment developer-optional book-keeping operations and statistical processing. The function will generally be called in a separate process after all the enrollment processes are complete.</p> <p>This function should be tolerant of being called two or more times. Second and third invocations should probably do nothing.</p>	
Input Parameters	enrollDir	The top-level directory in which enrollment data was placed. This variable allows an implementation to locate any private initialization data it elected to place in the directory.
	edbName	The name of a single file containing concatenated templates, i.e. the EDB of section 2.3.6. While the file will have read-write-delete permission, the SDK should only alter the file if it preserves the necessary content, in other files for example. The file may be opened directly. It is not necessary to prepend a directory name.
	edbManifest	The name of a single file containing the EDB manifest of section 2.3.6. The file may be opened directly. It is not necessary to prepend a directory name.
Output Parameters	None	
ReturnCode	Success	Success
	FinInputData	Cannot locate the input data - the input files or names seem incorrect.
	EnrollDirFailed	An operation on the enrollment directory failed (e.g. permission, space).
	FinTemplFormat	One or more template files are in an incorrect format.
	Vendor	Vendor-defined failure. Failure codes must be documented and communicated to NIST with the submission of the implementation under test.

583 **3.3.3. The VideoFeatureExtraction Interface**

584 The abstract class VideoFeatureExtraction must be implemented by the SDK developer in a class named exactly
 585 SdkVideoFeatureExtraction.

	C++ code fragment	Remarks
1.	class VideoFeatureExtraction	
2.	{ public:	
3.	virtual ReturnCode initialize(const string &configDir, const string &enrollDir) = 0;	Initialize the feature extraction session.

4.	<pre>virtual ReturnCode generateIdTemplate(const ONEVIDEO &inputVideo, vector<PERSONREP> &idTemplates) = 0;</pre>	<p>Generate identification template(s) for the persons detected in the input video. This function takes a ONEVIDEO (see 2.3.6) as input and populates a vector of PERSONREP (see 2.3.10) with the number of persons detected from the video sequence. The implementation could call <code>vector::push_back</code> to insert into the vector.</p>
5.	<pre>// Destructor</pre>	
6.	<pre>};</pre>	

586 **3.3.3.1. Video feature extraction initialization**

587 Before one or more **ONEVIDEOS** are sent to the identification feature extraction function, the test harness will call the
 588 initialization function below.

589 **Table 27 – VideoFeatureExtraction::initialize**

Prototype	ReturnCode initialize(
	const string &configDir,		Input
	const string &enrollDir);		Input
Description	<p>This function initializes the SDK under test and sets all needed parameters. This function will be called once by the NIST application immediately before any $M \geq 1$ calls to <code>generateIdTemplate</code>.</p> <p>The implementation has read-only access to <code>enrollDir</code> (containing prior enrollment data) and to <code>configDir</code>.</p>		
Input Parameters	configDir	A read-only directory containing any developer-supplied configuration parameters or run-time data files.	
	enrollDir	The read-only top-level directory in which enrollment data was placed and then finalized by the implementation. The implementation can parameterize subsequent template production on the basis of the enrolled dataset.	
Output Parameters	none		
ReturnCode	Success	Success	
	MissingConfig	The configuration data is missing, unreadable, or in an unexpected format.	
	EnrollDirFailed	An operation on the enrollment directory failed (e.g. permission).	
	Vendor	Vendor-defined failure	

590 **3.3.3.2. Video feature extraction**

591 A **ONEVIDEO** is converted to one or more identification templates using the function below. The result may be stored by
 592 NIST, or used immediately. The SDK shall not attempt to store any data.

593 **Table 28 – VideoFeatureExtraction::generateIdTemplate**

Prototypes	ReturnCode generateIdTemplate(
	const ONEVIDEO &inputVideo,		Input
	std::vector<PERSONREP> &idTemplates);		Output
Description	<p>This function takes a ONEVIDEO (see 2.3.6) as input and populates a vector of PERSONREP (see 2.3.10) with the number of persons detected from the video sequence. The implementation could call <code>vector::push_back</code> to insert into the vector.</p> <p>If the function executes correctly, it returns a zero exit status. The NIST calling application may commit the template to permanent storage, or may keep it only in memory (the implementation does not need to know). If the function returns a non-zero exit status, the output template will be not be used in subsequent search operations.</p> <p>The function shall not have access to the enrollment data, nor shall it attempt access.</p>		
Input Parameters	InputVideo	An instance of a section 2.3.6 class. Implementations must alter their behavior according to the people detected in the video sequence.	
Output Parameters	IdTemplates	For each person detected in the video, the function shall create a PERSONREP (see section 2.3.10) object, populate it with a template and eye coordinates for each frame where eyes	

		were detected, and add it to the vector.
ReturnCode	Success	Success
	RefuseInput	Elective refusal to process this kind of ONEVIDEO
	FailExtract	Involuntary failure to extract features (e.g. could not find face in the input-image)
	FailTempl	Elective refusal to produce a template (e.g. insufficient pixels between the eyes)
	FailParse	Cannot parse input data (i.e. assertion that input record is non-conformant)
	ImgSizeNotSupported	Input image/frame size too small or large
	Vendor	Vendor-defined failure. Failure codes must be documented and communicated to NIST with the submission of the implementation under test.

594 **3.3.4. The VideoSearch Interface**

595 The abstract class VideoSearch must be implemented by the SDK developer in a class named exactly SdkVideoSearch.

	C++ code fragment	Remarks
1.	class VideoSearch	
2.	{ public:	
3.	virtual ReturnCode initialize(const string &configDir, const string &enrollDir) = 0;	Initialize the search session.
4.	virtual ReturnCode identifyVideo(const PERSONREP &idVideoTemplate, const uint32_t candListLength, CANDIDATELIST &candList) = 0;	For video-to-video identification This function searches a template generated from a ONEVIDEO against the enrollment set, and outputs a vector containing candListLength objects of Candidates (see section 2.3.12).
5.	virtual ReturnCode identifyImage(const PERSONREP &idImageTemplate, const uint32_t candListLength, CANDIDATELIST &candList) = 0;	For still-to-video identification This function searches a template generated from a MULTIFACE against the enrollment set, and outputs a vector containing candListLength objects of Candidates.
6.	// Destructor	
7.	};	

596 **3.3.4.1. Video identification initialization**

597 The function below will be called once prior to one or more calls of the searching function of Table 30. The function might
598 set static internal variables so that the enrollment database is available to the subsequent identification searches.

599 **Table 29 – VideoSearch::initialize**

Prototype	ReturnCode initialize(const string &configDir, const string &enrollDir);	Input
		Input
Description	This function reads whatever content is present in the enrollment_directory, for example a manifest placed there by the VideoFinalize::finalize function.	
Input Parameters	configDir	A read-only directory containing any developer-supplied configuration parameters or run-time data files.
	enrollDir	The read-only top-level directory in which enrollment data was placed.
ReturnCode	Success	Success
	MissingConfig	The configuration data is missing, unreadable, or in an unexpected format.
	EnrollDirFailed	An operation on the enrollment directory failed (e.g. permission).
	Vendor	Vendor-defined failure

600 **3.3.4.2. Video identification search**

601 The function below compares a proprietary identification template against the enrollment data and returns a candidate
 602 list.

603 **Table 30 – VideoSearch::identifyVideo and VideoSearch::identifyImage**

Prototype	ReturnCode identifyVideo(const PERSONREP &idVideoTemplate, const uint32_t candListLength, CANDIDATELIST &candList);	Searches a template generated from a ONEVIDEO against the enrollment set (video-to-video)
		Input
		Output
	ReturnCode identifyImage(const PERSONREP &idImageTemplate, const uint32_t candListLength, CANDIDATELIST &candList);	Searches a template generated from a MULTIFACE against the enrollment set (still-to-video)
		Input
		Output
Description	This function searches an identification template against the enrollment set, and outputs a vector containing candListLength Candidates (see section 2.3.12). Each candidate shall be populated by the implementation and added to candList. Note that candList will be an empty vector when passed into this function. The candidates shall appear in descending order of similarity score - i.e. most similar entries appear first.	
Input Parameters	idTemplate	A template from generateIdTemplate() - If the value returned by that function was non-zero the contents of idTemplate will not be used and this function (i.e. identifyVideo) will not be called.
	candListLength	The number of candidates the search should return
Output Parameters	candList	A vector containing candListLength objects of Candidates. The datatype is defined in section 2.3.12. Each candidate shall be populated by the implementation and added to this vector. The candidates shall appear in descending order of similarity score - i.e. most similar entries appear first.
ReturnCode	Success	Success
	IdBadTempl	The input template was defective.
	Vendor	Vendor-defined failure

604 **3.3.5. The ImageEnrollment Interface**

605 The abstract class ImageEnrollment must be implemented by the SDK developer in a class named exactly
 606 SdkImageEnrollment.

	C++ code fragment	Remarks
1.	class ImageEnrollment	
2.	{ public:	
3.	virtual ReturnCode initialize(const string &configDir, const string &enrollDir, const uint32_t numPersons, const uint32_t numImages, const vector<string> &descriptions) = 0 ;	Initialize the enrollment session.
4.	virtual ReturnCode generateEnrollmentTemplate(const MULTIFACE &inputFaces, PERSONREP &outputTemplate) = 0;	This function takes a MULTIFACE (see 2.3.3) as input and outputs a proprietary template represented by a PERSONREP (see 2.3.10). For each input image in the MULTIFACE , the function shall return the estimated eye centers by setting PERSONREP .eyeCoordinates.
5.	// Destructor	
6.	};	

607 **3.3.5.1. Initialization of the image enrollment session**

608 Before any enrollment feature extraction calls are made, the NIST test harness will call the initialization below for video-
609 to-still.

610 **Table 31 – ImageEnrollment::initialize**

Prototype	ReturnCode initialize(const string &configDir, const string &enrollDir, const uint32_t numPersons, const uint32_t numImages, const std::vector<string> &descriptions);	
		Input
Description	This function initializes the SDK under test and sets all needed parameters. This function will be called N=1 times by the NIST application immediately before any $M \geq 1$ calls to generateEnrollmentTemplate. Caution: The implementation should tolerate execution of $P > 1$ processes on the one or more machines each of which may be reading and writing to this same enrollment directory in parallel. File locking or process-specific temporary filenames would be needed to safely write content in the enrollDir.	
Input Parameters	configDir	A read-only directory containing any developer-supplied configuration parameters or run-time data files.
	enrollDir	The directory will be initially empty, but may have been initialized and populated by separate invocations of the enrollment process. When this function is called, the SDK may populate this folder in any manner it sees fit. Permissions will be read-write-delete.
	numPersons	The number of persons who will be enrolled.
	numImages	The total number of images that will be enrolled, summed over all identities.
	descriptions	A lexicon of labels one of which will be assigned to each enrollment image. See Table 9 for valid values. NOTE: The identification search images may or may not be labeled. An identification image may carry a label not in this set of labels. The number of items stored in the vector is accessible via the vector::size() function.
Output Parameters	none	
ReturnCode	Success	Success
	MissingConfig	The configuration data is missing, unreadable, or in an unexpected format.
	EnrollDirFailed	An operation on the enrollment directory failed (e.g. permission, space).
	InitNumData	The SDK cannot support the number of videos.
	InitBadDesc	The descriptions are unexpected, or unusable.
	Vendor	Vendor-defined failure

611 **3.3.5.2. Image enrollment**

612 A **MULTIFACE** (see Table 12) is converted to a single enrollment template using the function below.

613 **Table 32 – ImageEnrollment::generateEnrollmentTemplate**

Prototypes	ReturnCode generateEnrollmentTemplate(const MULTIFACE &inputFaces, PERSONREP &outputTemplate);	
		Input Output
Description	<p>This function takes a MULTIFACE, and outputs a proprietary template in the form of a PERSONREP object. If the function executes correctly (i.e. returns a ReturnCode::Success exit status), the NIST calling application will store the template. The NIST application will concatenate the templates and pass the result to the enrollment finalization function.</p> <p>If the function gives a non-zero exit status:</p> <ul style="list-style-type: none"> – If the exit status is ReturnCode::FailParse, NIST will debug, otherwise – the test driver will ignore the output template (the template may have any size including zero) – the event will be counted as a failure to enroll. Such an event means that this person can never be identified 	

	correctly. IMPORTANT. NIST's application writes the template to disk. The implementation must not attempt writes to the enrollment directory (nor to other resources). Any data needed during subsequent searches should be included in the template, or created from the templates during the enrollment finalization function.	
Input Parameters	inputFaces	An instance of a Table 12 structure.
Output Parameters	outputTemplate	An instance of a section 2.3.10 class, which stores proprietary template data and eye coordinates. The function shall identify the person's estimated eye centers for each image in the MULTIFACE . The eye coordinates shall be captured in the PERSONREP.eyeCoordinates variable, which is a vector of EYEPAIR objects. In the event the eye centers cannot be calculated, the SDK shall store an EYEPAIR and set EYEPAIR.isSet to false to indicate there was a failure in generating eye coordinates. In other words, for N images in the MULTIFACE .
ReturnCode	Success	Success
	RefuseInput	Elective refusal to process this kind of MULTIFACE
	FailExtract	Involuntary failure to extract features (e.g. could not find face in the input-image)
	FailTempl	Elective refusal to produce a template (e.g. insufficient pixels between the eyes)
	FailParse	Cannot parse input data (i.e. assertion that input record is non-conformant)
	ImgSizeNotSupported	Input image/frame size too small or large
	Vendor	Vendor-defined failure. Failure codes must be documented and communicated to NIST with the submission of the implementation under test.

614 **3.3.6. The ImageFinalize Interface**

615 The abstract class ImageFinalize must be implemented by the SDK developer in a class named exactly SdkImageFinalize.
 616 The finalize function in this class takes the name of the top-level directory where enrollment database (EDB) and its
 617 manifest have been stored. These are described in section 2.3.6. The enrollment directory permissions will be read +
 618 write.

	C++ code fragment	Remarks
1.	class ImageFinalize	
2.	{ public:	
3.	virtual ReturnCode finalize(const string &enrollDir, const string &edbName, const string &edbManifest) = 0;	This function supports post-enrollment developer-optional book-keeping operations and statistical processing. The function will generally be called in a separate process after all the enrollment processes are complete.
4.	// Destructor	
5.	};	

619 **3.3.6.1. Finalize image enrollment**

620 After all templates have been created, the function of Table 33 will be called. This freezes the enrollment data. After this
 621 call the enrollment dataset will be forever read-only. This API does not support interleaved enrollment and search
 622 phases.

623 The function allows the implementation to conduct, for example, statistical processing of the feature data, indexing and
 624 data re-organization. The function may alter the file structure. It may increase or decrease the size of the stored data.
 625 No output is expected from this function, except a return code.

626 **Table 33 – ImageFinalize::finalize**

Prototypes	ReturnCode finalize(const string &enrollDir, const string &edbName, const string &edbManifest);	
		Input
		Input
		Input
Description	This function takes the name of the top-level directory where enrollment database (EDB) and its manifest have	

	<p>been stored. These are described in section 2.3.6. The enrollment directory permissions will be read + write.</p> <p>The function supports post-enrollment developer-optional book-keeping operations and statistical processing. The function will generally be called in a separate process after all the enrollment processes are complete.</p> <p>This function should be tolerant of being called two or more times. Second and third invocations should probably do nothing.</p>	
Input Parameters	enrollDir	The top-level directory in which enrollment data was placed. This variable allows an implementation to locate any private initialization data it elected to place in the directory.
	edbName	The name of a single file containing concatenated templates, i.e. the EDB of section 2.3.6. While the file will have read-write-delete permission, the SDK should only alter the file if it preserves the necessary content, in other files for example. The file may be opened directly. It is not necessary to prepend a directory name.
	edbManifest	The name of a single file containing the EDB manifest of section 2.3.6. The file may be opened directly. It is not necessary to prepend a directory name.
Output Parameters	None	
ReturnCode	Success	Success
	FinInputData	Cannot locate the input data - the input files or names seem incorrect.
	EnrollDirFailed	An operation on the enrollment directory failed (e.g. permission, space).
	FinTemplFormat	One or more template files are in an incorrect format.
	Vendor	Vendor-defined failure. Failure codes must be documented and communicated to NIST with the submission of the implementation under test.

627 **3.3.7. The ImageFeatureExtraction Interface**

628 The abstract class ImageFeatureExtraction must be implemented by the SDK developer in a class named exactly
629 SdkImageFeatureExtraction.

	C++ code fragment	Remarks
1.	class ImageFeatureExtraction	
2.	{	
	public:	
3.	virtual ReturnCode initialize(const string &configDir, const string &enrollDir) = 0;	Initialize the feature extraction session.
4.	virtual ReturnCode generateIdTemplate(const MULTIFACE &inputFaces, PERSONREP &outputTemplate) = 0;	This function takes a MULTIFACE (see 2.3.3) as input and outputs a proprietary template represented by a PERSONREP (see 2.3.10). For each input image in the MULTIFACE , the function shall return the estimated eye centers by setting PERSONREP .eyeCoordinates.
5.	// Destructor	
6.	};	

630 **3.3.7.1. Image feature extraction initialization**

631 Before one or more **MULTIFACE**s are sent to the identification feature extraction function, the test harness will call the
632 initialization function below.

633 **Table 34 – ImageFeatureExtraction::initialize**

Prototype	ReturnCode initialize(const string &configDir, const string &enrollDir);	Input
		Input
Description	This function initializes the SDK under test and sets all needed parameters. This function will be called once by	

	the NIST application immediately before $M \geq 1$ calls generateIdTemplate. The implementation has read-only access to enrollDir (containing prior enrollment data) and to configDir.	
Input Parameters	configDir	A read-only directory containing any developer-supplied configuration parameters or run-time data files.
	enrollDir	The read-only top-level directory in which enrollment data was placed and then finalized by the implementation. The implementation can parameterize subsequent template production on the basis of the enrolled dataset.
Output Parameters	none	
ReturnCode	Success	Success
	MissingConfig	The configuration data is missing, unreadable, or in an unexpected format.
	EnrollDirFailed	An operation on the enrollment directory failed (e.g. permission).
	Vendor	Vendor-defined failure

634 **3.3.7.2. Image feature extraction**

635 A **MULTIFACE** is converted to one identification template using the function below. The result may be stored by NIST, or
 636 used immediately. The SDK shall not attempt to store any data.

637 **Table 35 – ImageFeatureExtraction::generateIdTemplate**

Prototypes	ReturnCode generateIdTemplate(const MULTIFACE &inputFaces, PERSONREP &outputTemplate);	Input Output
Description	<p>This function takes a MULTIFACE (see 2.3.3) as input and populates a PERSONREP (see 2.3.10) with a proprietary template and eye coordinates.</p> <p>If the function executes correctly, it returns a zero exit status. The NIST calling application may commit the template to permanent storage, or may keep it only in memory (the developer implementation does not need to know). If the function returns a non-zero exit status, the output template will be not be used in subsequent search operations.</p> <p>The function shall not have access to the enrollment data, nor shall it attempt access.</p>	
Input Parameters	inputFaces	An instance of a Table 12 structure.
Output Parameters	outputTemplate	An instance of a section 2.3.10 class, which stores proprietary template data and eye coordinates. The function shall identify the person’s estimated eye centers for each image in the MULTIFACE . The eye coordinates shall be captured in the PERSONREP.eyeCoordinates variable, which is a vector of EYEPAIR objects. In the event the eye centers cannot be calculated, the SDK shall store an EYEPAIR and set EYEPAIR.isSet to false to indicate there was a failure in generating eye coordinates. In other words, for N images in the MULTIFACE .
ReturnCode	Success	Success
	RefuseInput	Elective refusal to process this kind of MULTIFACE
	FailExtract	Involuntary failure to extract features (e.g. could not find face in the input-image)
	FailTempl	Elective refusal to produce a template (e.g. insufficient pixels between the eyes)
	FailParse	Cannot parse input data (i.e. assertion that input record is non-conformant)
	ImgSizeNotSupported	Input image/frame size too small or large
	Vendor	Vendor-defined failure. Failure codes must be documented and communicated to NIST with the submission of the implementation under test.

638 **3.3.8. The ImageSearch Interface**

639 The abstract class ImageSearch must be implemented by the SDK developer in a class named exactly SdkImageSearch.

	C++ code fragment	Remarks
1.	class VideoFeatureExtraction	
2.	{ public:	

3.	<code>virtual ReturnCode initialize(const string &configDir, const string &enrollDir) = 0;</code>	Initialize the search session.
4.	<code>virtual ReturnCode identifyVideo(const PERSONREP &idTemplate, const uint32_t candListLength, CANDIDATELIST &candList) = 0;</code>	For video-to-still identification This function searches a template generated from a ONEVIDEO against the enrollment set, and outputs a vector containing <code>candListLength</code> objects of Candidates (see section 2.3.12). Each candidate shall be populated by the implementation and added to <code>candList</code> . The candidates shall appear in descending order of similarity score - i.e. most similar entries appear first.
5.	<code>// Destructor</code>	
6.	<code>};</code>	

640 **3.3.8.1. Image identification initialization**

641 The function below will be called once prior to one or more calls of the searching function of Table 37. The function might
642 set static internal variables so that the enrollment database is available to the subsequent identification searches.

643 **Table 36 – ImageSearch::initialize**

Prototype	ReturnCode initialize(const string &configDir, const string &enrollDir);	
		Input Input
Description	This function reads whatever content is present in the <code>enrollment_directory</code> , for example a manifest placed there by the <code>ImageFinalize::finalize</code> function.	
Input Parameters	<code>configDir</code>	A read-only directory containing any developer-supplied configuration parameters or run-time data files.
	<code>enrollDir</code>	The read-only top-level directory in which enrollment data was placed.
ReturnCode	Success	Success
	MissingConfig	The configuration data is missing, unreadable, or in an unexpected format.
	EnrollDirFailed	An operation on the enrollment directory failed (e.g. permission).
	Vendor	Vendor-defined failure

644 **3.3.8.2. Image identification search**

645 The function below performs a video-to-still identification and compares a proprietary identification template generated
646 from a video against the enrollment data and returns a candidate list.

647 **Table 37 – ImageSearch::identifyVideo**

Prototype	ReturnCode identifyVideo(const PERSONREP &idVideoTemplate, const uint32_t candListLength, CANDIDATELIST &candList);	
		Searches a template generated from a ONEVIDEO against the enrollment set (video-to-still)
		Input
		Input Output
Description	This function searches an identification template against the enrollment set, and outputs a vector containing <code>candListLength</code> objects of Candidates (see section 2.3.12). Each candidate shall be populated by the implementation and added to <code>candList</code> . Note that <code>candList</code> will be an empty vector when passed into this function. The candidates shall appear in descending order of similarity score - i.e. most similar entries appear first.	
Input Parameters	<code>idTemplate</code>	A template from <code>VideoFeatureExtraction::generatedIdTemplate()</code> - If the value returned by that function was non-zero the contents of <code>idTemplate</code> will not be used and this function (i.e. <code>identifyVideo</code>) will not be called.
	<code>candListLength</code>	The number of candidates the search should return
Output Parameters	<code>candList</code>	A vector containing <code>candListLength</code> objects of Candidates. The datatype is defined in section 2.3.12. Each candidate shall be populated by the implementation and

FIVE

		added to this vector. The candidates shall appear in descending order of similarity score - i.e. most similar entries appear first.
ReturnCode	Success	Success
	IdBadTempl	The input template was defective.
	Vendor	Vendor-defined failure

648 NOTE: Ordinarily the calling application will set the input candidate list length to operationally typical values, say $0 \leq L \leq$
 649 200, and $L \ll N$. However, there is interest in the presence of mates much further down the candidate list. We may
 650 therefore extend the candidate list length such that L approaches N.

651 **4. References**

AN27	NIST Special Publication 500-271: American National Standard for Information Systems — Data Format for the Interchange of Fingerprint, Facial, & Other Biometric Information – Part 1. (ANSI/NIST ITL 1-2007). Approved April 20, 2007.
FRVT 2002	Face Recognition Vendor Test 2002: Evaluation Report, NIST Interagency Report 6965, P. Jonathon Phillips, Patrick Grother, Ross J. Micheals, Duane M. Blackburn, Elham Tabassi, Mike Bone
FRVT 2002b	Face Recognition Vendor Test 2002: Supplemental Report, NIST Interagency Report 7083, Patrick Grother
FRVT 2006	P. Jonathon Phillips, W. Todd Scruggs, Alice J. O’Toole, Patrick J. Flynn, Kevin W. Bowyer, Cathy L. Schott, and Matthew Sharpe. "FRVT 2006 and ICE 2006 Large-Scale Results." NISTIR 7408, March 2007.
FRVT 2013	P. Grother and M. Ngan, Face Recognition Vendor Test (FRVT), Performance of Face Identification Algorithms, NIST Interagency Report 8009, Released May 26, 2014. http://face.nist.gov/frvt
IREX III	P. Grother, G.W. Quinn, J. Matey, M. Ngan, W. Salamon, G. Fiumara, C. Watson, Iris Exchange III, Performance of Iris Identification Algorithms, NIST Interagency Report 7836, Released April 9, 2012. http://iris.nist.gov/irex
ISO STD05	ISO/IEC 19794-5:2005 — Information technology — Biometric data interchange formats — Part 5: Face image data. The standard was published in 2005, and can be purchased from ANSI at http://webstore.ansi.org/ Multipart standard of "Biometric data interchange formats". This standard was published in 2005. It was amended twice to include guidance to photographers, and then to include 3D information. Two corrigenda were published. All these changes and new material is currently being incorporated in revision of the standard. Publication is likely in early 2011. The documentary history is as follows. ISO/IEC 19794-5: Information technology — Biometric data interchange formats — Part 5:Face image data. First edition: 2005-06-15. International Standard ISO/IEC 19794-5:2005 Technical Corrigendum 1: Published 2008-07-01 International Standard ISO/IEC 19794-5:2005 Technical Corrigendum 2: Published 2008-07-01 Information technology — Biometric data interchange formats — Part 5: Face image data AMENDMENT 1: Conditions for taking photographs for face image data. Published 2007-12-15 Information technology — Biometric data interchange formats — Part 5: Face image data AMENDMENT 2: Three dimensional image data. JTC 1/SC37/N3303. FCD text of the second edition. Contact pgrother AT nist DOT gov for more information.
MBE	P. Grother, G .W. Quinn, and P. J. Phillips, Multiple-Biometric Evaluation (MBE) 2010, Report on the Evaluation of 2D Still Image Face Recognition Algorithms, NIST Interagency Report 7709, Released June 22, 2010. Revised August 23, 2010. http://face.nist.gov/mbe
MINEX	P. Grother et al., Performance and Interoperability of the INCITS 378 Template, NIST IR 7296 http://fingerprint.nist.gov/minex04/minex_report.pdf
MOC	P. Grother and W. Salamon, MINEX II - An Assessment of ISO/IEC 7816 Card-Based Match-on-Card Capabilities http://fingerprint.nist.gov/minex/minexII/NIST_MOC_ISO_CC_interop_test_plan_1102.pdf
PERFSTD INTEROP	ISO/IEC 19795-4 — Biometric Performance Testing and Reporting — Part 4: Interoperability Performance Testing. Posted as document 37N2370. The standard was published in 2007. It can be purchased from ANSI at http://webstore.ansi.org/ .

652

Annex A

Submission of Implementations to the FIVE

653
654

655 **A.1 Submission of implementations to NIST**

656 NIST requires that all software, data and configuration files submitted by the participants be signed and encrypted.
657 Signing is done with the participant's private key, and encryption is done with the NIST public key. The detailed
658 commands for signing and encrypting are given here: <http://www.nist.gov/itl/iad/ig/encrypt.cfm>

659 NIST will validate all submitted materials using the participant's public key, and the authenticity of that key will be verified
660 using the key fingerprint. This fingerprint must be submitted to NIST by writing it on the signed participation agreement.

661 By encrypting the submissions, we ensure privacy; by signing the submission, we ensure authenticity (the software
662 actually belongs to the submitter). NIST will reject any submission that is not signed and encrypted. NIST accepts no
663 responsibility for anything that is transmitted to NIST that is not signed and encrypted with the NIST public key.

664 **A.2 How to participate**

665 Those wishing to participate in FIVE testing must do all of the following, on the schedule listed on Page 2.

- 666 — **IMPORTANT:** Follow the instructions for cryptographic protection of your SDK and data here.
667 <http://www.nist.gov/itl/iad/ig/encrypt.cfm>
- 668 — Send a signed and fully completed copy of the *Application to Participate in the Face In Video Evaluation (FIVE)*. This is
669 available at <http://www.nist.gov/itl/iad/ig/five.cfm>. This must identify, and include signatures from, the Responsible
670 Parties as defined in the application. The properly signed FIVE Application to Participate shall be sent to NIST as a
671 PDF.
- 672 — Provide an SDK (Software Development Kit) library which complies with the API (Application Programmer Interface)
673 specified in this document.
- 674 • Encrypted data and SDKs below 20MB can be emailed to NIST at five@nist.gov
 - 675 • Encrypted data and SDKS above 20MB shall be
- 676 EITHER
- 677 ▪ Split into sections AFTER the encryption step. Use the unix "split" commands to make 9MB chunks,
678 and then rename to include the filename extension need for passage through the NIST firewall.
 - 679 ▪ `you% split -a 3 -d -b 9000000 libFIVE_enron_A_02.tgz.gpg`
 - 680 ▪ `you% ls -l x??? | xargs -iQ mv Q libFIVE_enron_A_02_Q.tgz.gpg`
 - 681 ▪ Email each part in a separate email. Upon receipt NIST will
 - 682 ▪ `nist% cat FIVE2012_enron_A02_*.tgz.gpg > libFIVE_enron_A_02.tgz.gpg`
- 683 OR
- 684 ▪ Made available as a file.zip.gpg or file.zip.asc download from a generic http webserver⁸,
- 685 OR
- 686 ▪ Mailed as a file.zip.gpg or file.zip.asc on CD / DVD to NIST at this address:

FIVE Test Liaison (A203) 100 Bureau Drive A203/Tech225/Stop 8940 NIST Gaithersburg, MD 20899-8940 USA	In cases where a courier needs a phone number, please use NIST shipping and handling on: 301 -- 975 -- 6296.
--	---

⁸ NIST will not register, or establish any kind of membership, on the provided website.

687 A.3 Implementation validation

688 Registered Participants will be provided with a small validation dataset and test program available on the website

689 <http://www.nist.gov/itl/iad/ig/five.cfm> shortly after the final evaluation plan is released.

690 The validation test programs shall be compiled by the provider. The output of these programs shall be submitted to NIST.

691 Prior to submission of the SDK and validation data, the Participant must verify that their software executes on the
692 validation images, and produces correct similarity scores and templates.

693 Software submitted shall implement the FIVE API Specification as detailed in the body of this document.

694 Upon receipt of the SDK and validation output, NIST will attempt to reproduce the same output by executing the SDK on
695 the validation imagery, using a NIST computer. In the event of disagreement in the output, or other difficulties, the
696 Participant will be notified.