

1
2
3
4 Face Recognition Vendor Test (FRVT)
5 (Ongoing)
6
7
8
9

10 Still Face 1:1 Verification
11 Concept, Evaluation Plan and API
12 Version 2.0

13 Updates since version 1.0 of this document are highlighted in **green**.
14

15
16 Patrick Grother and Mei Ngan

17 Contact via frvt@nist.gov

Image Group
Information Access Division
Information Technology Laboratory



May 22, 2017

18
19
20

Table of Contents

21

22 1. FRVT 3

23 1.1. Scope 3

24 1.2. Audience 3

25 1.3. Rules for Participation 3

26 1.4. Reporting 3

27 1.5. Hardware specification 3

28 1.6. Operating system, compilation, and linking environment..... 4

29 1.7. Software and Documentation..... 4

30 1.8. Runtime behavior 6

31 1.9. Single-thread Requirement/Parallelization 6

32 1.10. Time limits 6

33 2. Data structures supporting the API..... 7

34 2.1. Requirement 7

35 2.2. File formats and data structures..... 7

36 3. API Specification 10

37 3.1. Namespace 10

38 3.2. Overview..... 10

39 3.3. API..... 11

40

41 **List of Tables**

42 Table 1 – Implementation library filename convention 5

43 Table 2 – Processing time limits in milliseconds, per 640 x 480 image 6

44 Table 3 – Structure for a single image..... 7

45 Table 4 – Labels describing categories of Images..... 7

46 Table 5 – Structure for a set of images from a single person..... 8

47 Table 6 – Structure for a pair of eye coordinates..... 8

48 Table 7 – Labels describing template role..... 8

49 Table 8 – Enumeration of return codes..... 8

50 Table 9 – ReturnStatus structure 9

51 Table 10 – Structure containing subject metadata information 9

52 Table 11 – Structure representing face image and associated attributes..... 10

53 Table 12 – Functional summary of the 1:1 application 10

54 Table 13 – Initialization 12

55 Table 14 – GPU index specification 12

56 Table 15 – Template generation 12

57 Table 16 – Template matching 13

58 Table 17 – Training..... 14

59

60 **List of Figures**

61 Figure 1 – Schematic of 1:1 verification 10

62 Figure 2 – Schematic of training..... 14

63

64

65 1. FRVT

66 1.1. Scope

67 This document establishes a concept of operations and an application programming interface (API) for evaluation of face
68 recognition (FR) implementations submitted to NIST's ongoing Face Recognition Vendor Test. This API is for the 1:1
69 identity verification track. Separate API documents will be published for future additional tracks to FRVT. All images
70 include exactly one face.

71 1.2. Audience

72 Participation in FRVT is open to any organization worldwide. There is no charge for participation. The target audience is
73 researchers and developers of FR algorithms. While NIST intends to evaluate stable technologies that could be readily
74 made operational, the test is also open to experimental, prototype and other technologies. All algorithms **must** be
75 submitted as implementations of the API defined in this document.

76 1.3. Rules for Participation

77 1.3.1. Participation Agreement

78 A participant must properly follow, complete, and submit the FRVT Participation Agreement. This must be done once,
79 either prior or in conjunction with the very first algorithm submission. It is not necessary to do this for each submitted
80 implementation thereafter UNLESS there are major organizational changes to the submitting entity.

81 1.3.2. Number and Schedule of Submissions

82 Participants may send up to two initial submissions that run to completion. After that, participations may send one
83 submission as often as every **120 days three calendar months** from the last submission for evaluation. NIST will evaluate
84 implementations on a first-come-first-served basis, and quickly publish results.

85 1.4. Reporting

86 For all algorithms that complete the evaluations, NIST will post performance results on the NIST FRVT website. NIST will
87 maintain an email list to inform interested parties of updates to the website. Artifacts will include a leaderboard
88 highlighting the top performing submissions in various areas (e.g., accuracy, speed etc.) and individual implementation-
89 specific report cards. NIST will maintain reporting on the two most recent algorithm submissions from any organization.
90 Prior submission results will be archived but remain accessible via a public link.

91
92 **Important:** This is an open test in which NIST will identify the algorithm and the developing organization. Algorithm
93 results will be attributed to the developer. Results will be machine generated (i.e. scripted) and will include timing,
94 accuracy and other performance results. These will be posted alongside results from other implementations. Results will
95 be expanded and modified as additional implementations are tested, and as analyses are implemented. Results may be
96 regenerated on-the-fly, usually whenever additional implementations complete testing, or when new analysis is added.
97

98 NIST may additionally report results in workshops, conferences, conference papers and presentations, journal articles and
99 technical reports.

100 1.5. Hardware specification

101 NIST intends to support high performance by specifying the runtime hardware beforehand. There are several types of
102 computer blades that may be used in the testing. Each CPU has 512K cache. The bus runs at 667 Mhz. The main memory
103 is 192 GB Memory as 24 8GB modules. We anticipate that 16 processes can be run without time slicing, though NIST will
104 handle all multiprocessing work via `fork()`¹. Participant-initiated multiprocessing is not permitted.

¹ <http://man7.org/linux/man-pages/man2/fork.2.html>

105 NIST is requiring use of 64 bit implementations throughout.

106 **1.5.1. Central Processing Unit (CPU)-only platforms**

107 The following list gives some details about the hardware of each CPU-only blade type:

- 108 • Dual Intel Xeon X5680 3.3 GHz CPUs (6 cores each)
- 109 • Dual Intel Xeon X7560 2.3 GHz CPUs (8 cores each)
- 110 • Dual Intel(R) Xeon(R) CPU E5-2630 v4 @ 2.20GHz (10 cores each)

111 **1.5.2. Graphics Processing Units (GPU)-enabled platforms**

112 The following provides some details about the hardware of GPU-enabled machines:

- 113 • Dual Intel Xeon E5-2695 3.3 GHz CPUs (14 cores each; 56 logical CPUs total) with Dual NVIDIA Tesla K40 GPUs,
114 with 12GB of memory per GPU

115 All GPU-enabled machines will be running CUDA version 7.5. cuDNN v5 for CUDA 7.5 will also be installed on these
116 machines. Implementations that use GPUs will only be run on GPU-enabled machines. Please note that GPU-dependent
117 implementations submitted to FRVT will have longer test turnaround times than CPU-only implementations due to
118 resource constraints. Developers submitting GPU implementations are encouraged to submit “CPU-equivalent”
119 implementations of their algorithms for timing comparisons. Algorithms using GPUs will be identified as such in public
120 reports.

121 **1.6. Operating system, compilation, and linking environment**

122 The operating system that the submitted implementations shall run on will be released as a downloadable file accessible
123 from http://nigos.nist.gov:8080/evaluations/CentOS-7-x86_64-Everything-1511.iso, which is the 64-bit version of CentOS
124 7.2 running Linux kernel 3.10.0.

125 For this test, Windows machines will not be used. Windows-compiled libraries are not permitted. All software must run
126 under CentOS 7.2.

127 NIST will link the provided library file(s) to our C++ language test drivers. Participants are required to provide their library
128 in a format that is dynamically-linkable using the C++11 compiler, g++ version 4.8.5.

129 A typical link line might be

```
130 g++ -std=c++11 -I. -Wall -m64 -o frvt11 frvt11.cpp -L. -lfrvt11_acme_07_cpu
```

131 The Standard C++ library should be used for development. The prototypes from this document will be written to a file
132 "frvt11.h" which will be included via

```
#include <frvt11.h>
```

133 The header files will be made available to implementers at <https://github.com/usnistgov/frvt>.

134 All compilation and testing will be performed on x86_64 platforms. Thus, participants are strongly advised to verify
135 library-level compatibility with g++ (on an equivalent platform) prior to submitting their software to NIST to avoid linkage
136 problems later on (e.g. symbol name and calling convention mismatches, incorrect binary file formats, etc.).

137 **1.7. Software and Documentation**

138 **1.7.1. Library and Platform Requirements**

139 Participants shall provide NIST with binary code only (i.e. no source code). The implementation should be submitted in
140 the form of a dynamically-linked library file.

141

142 The core library shall be named according to Table 1. Additional supplemental libraries may be submitted that support
143 this “core” library file (i.e. the “core” library file may have dependencies implemented in these other libraries).

144 Supplemental libraries may have any name, but the “core” library must be dependent on supplemental libraries in order
145 to be linked correctly. The **only** library that will be explicitly linked to the FRVT 1:1 test driver is the “core” library.

146
147 Intel Integrated Performance Primitives (IPP)® libraries are permitted if they are delivered as a part of the developer-
148 supplied library package. It is the provider’s responsibility to establish proper licensing of all libraries. The use of IPP
149 libraries shall not prevent running on CPUs that do not support IPP. Please take note that some IPP functions are
150 multithreaded and threaded implementations are prohibited.

151
152 NIST will report the size of the supplied libraries.

153 **Table 1 – Implementation library filename convention**

Form	libfrvt11_provider_sequence_processor.ending				
Underscore delimited parts of the filename	libfrvt11	provider	sequence	processor	ending
Description	First part of the name, required to be this.	Single word, non-infringing name of the main provider EXAMPLE: Acme	A three digit decimal identifier to start at 000 and incremented by 1 every time a library is sent to NIST. EXAMPLE: 007	“gpu” if implementation uses GPUs; “cpu” otherwise	.so
Example	libfrvt11_acme_007_cpu.so				

154 Important: Public results will be attributed with the provider name and the 3-digit sequence number in the submitted
155 library name.

156 **1.7.2. Configuration and developer-defined data**

157 The implementation under test may be supplied with configuration files and supporting data files. NIST will report the
158 size of the supplied configuration files.

159 **1.7.3. Submission folder hierarchy**

160 Participant submissions shall contain the following folders at the top level

- 161 • lib/ - contains all participant-supplied software libraries
- 162 • config/ - contains all configuration and developer-defined data
- 163 • doc/ - contains any participant-provided documentation regarding the submission
- 164 • validation/ - contains validation output

165 **1.7.4. Installation and Usage**

166 The implementation shall be installable using simple file copy methods. It shall not require the use of a separate
167 installation program and shall be executable on any number of machines without requiring additional machine-specific
168 license control procedures or activation. The implementation shall not use nor enforce any usage controls or limits based
169 on licenses, number of executions, presence of temporary files, etc. The implementation shall remain operable for at
170 least six months from the submission date.

171 **1.7.5. Documentation**

172 Participants shall provide documentation of additional functionality or behavior beyond that specified here. The
173 documentation must define all (non-zero) developer-defined error or warning return codes.

174 **1.7.6. Modes of operation**

175 Implementations shall not require NIST to switch “modes” of operation or algorithm parameters. For example, the use of
176 two different feature extractors must either operate automatically or be split across two separate library submissions.

177 1.8. Runtime behavior

178 1.8.1. Interactive behavior, stdout, logging

179 The implementation will be tested in non-interactive “batch” mode (i.e. without terminal support). Thus, the submitted
180 library shall:

- 181 – Not use any interactive functions such as graphical user interface (GUI) calls, or any other calls which require
182 terminal interaction e.g. reads from “standard input”.
- 183 – Run quietly, i.e. it should not write messages to "standard error" and shall not write to “standard output”.
- 184 – Only if requested by NIST for debugging, include a logging facility in which debugging messages are written to a
185 log file whose name includes the provider and library identifiers and the process PID.

186 1.8.2. Exception Handling

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

189 1.8.3. External communication

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

195 1.8.4. Stateless behavior

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

201 1.9. Single-thread Requirement/Parallelization

202 Implementations must run in single-threaded mode, because NIST will parallelize the test by dividing the workload across
203 many cores and many machines. Implementations must ensure that there are no issues with their software being
204 parallelized via the `fork()` function – this applies to both GPU and CPU implementations submitted to FRVT.

205 For implementations using the GPU: For any given GPU, NIST will run a single implementation process (i.e., `fork()` once per
206 GPU), with 12GB of main memory available for use by the algorithm. NIST machines are equipped with dual GPUs, and
207 the NIST test harness will load balance by telling the implementation which GPU to use via the section 3.3.2.1 `setGPU()`
208 function call. All calls to `setGPU()` will be performed after a call to `fork()`. Implementations using the GPU are encouraged
209 to perform initialization within the `setGPU()` function where 1. which GPU to use is provided to the implementation and 2.
210 to support known limitations of commonly used deep learning frameworks such as Caffe, where initialization must take
211 place in the worker process.

212 1.10. Time limits

213 The elemental functions of the implementations shall execute under the time constraints of Table 2. These time limits
214 apply to the function call invocations defined in section 3. Assuming the times are random variables, NIST cannot regulate
215 the maximum value, so the time limits are 90-th percentiles. This means that 90% of all operations should take less than
216 the identified duration.

217 The time limits apply per image. When K images of a person are present, the time limits shall be increased by a factor K.

218 **Table 2 – Processing time limits in milliseconds, per 640 x 480 image**

Function	1:1 verification
Training	12 hours for an input set of 6000 images
Feature extraction enrollment	1000 (1 core) 640x480 pixels
Feature extraction for verification	1000 (1 core) 640x480 pixels
Matching	5 (1 core)

219

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

221 **2.1. Requirement**

222 FRVT 1:1 participants shall implement the relevant C++ prototyped interfaces of clause 3. C++ was chosen in order to
223 make use of some object-oriented features.

224 **2.2. File formats and data structures**

225 **2.2.1. Overview**

226 In this face recognition test, an individual is represented by $K \geq 1$ two-dimensional facial images. All facial images in the
227 test will contain one and only one face per image.

228 **Table 3 – Structure for a single image**

C++ code fragment	Remarks
typedef struct Image	
{	
uint16_t image_width;	Number of pixels horizontally
uint16_t image_height;	Number of pixels vertically
uint16_t image_depth;	Number of bits per pixel. Legal values are 8 and 24.
std::shared_ptr<uint8_t> data;	Managed 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
Label description;	Single description of the image. The allowed values for this field are specified in the enumeration in Table 4.
} Image;	

229

230 An **Image** will be accompanied by one of the labels given below. Face recognition implementations should tolerate
231 **Images** of any category.

232 **Table 4 – Labels describing categories of Images**

Label as C++ enumeration	Meaning
enum class Label {	
UNKNOWN=0,	Either the label is unknown or unassigned.
ISO=1,	Frontal, intended to be in conformity to ISO/IEC 19794-5:2005.
MUGSHOT=2,	From law enforcement booking processes. Nominally frontal.
PHOTOJOURNALISM=3,	The image might appear in a news source or magazine. The images are typically taken by professional photographer and are well exposed and focused but exhibit pose and illumination variations.
EXPLOITATION=4	The image is taken from a child exploitation database. This imagery has highly unconstrained pose and illumination, expression and resolution.
WILD=5	Unconstrained image, taken by an amateur photographer, exhibiting wide variations in pose, illumination, and resolution.
};	

233

234

Table 5 – Structure for a set of images from a single person

C++ code fragment	Remarks
<code>using Multiface = std::vector<Image>;</code>	Vector of Image objects

235 **2.2.2. Data structure for eye coordinates**

236 Implementations should return eye coordinates of each facial image. This function, while not necessary for a recognition
 237 test, will assist NIST in assuring the correctness of the test database. The primary mode of use will be for NIST to inspect
 238 images for which eye coordinates are not returned, or differ between implementations.

239 The eye coordinates shall follow the placement semantics of the ISO/IEC 19794-5:2005 standard - the geometric
 240 midpoints of the endocanthion and exocanthion (see clause 5.6.4 of the ISO standard).

241 Sense: The label "left" refers to subject's left eye (and similarly for the right eye), such that `xright < xleft`.

242 **Table 6 – Structure for a pair of eye coordinates**

C++ code fragment	Remarks
<code>typedef struct EyePair</code>	
<code>{</code>	
<code> bool isLeftAssigned;</code>	If the subject's left eye coordinates have been computed and assigned successfully, this value should be set to true, otherwise false.
<code> bool isRightAssigned;</code>	If the subject's right eye coordinates have been computed and assigned successfully, this value should be set to true, otherwise false.
<code> uint16_t xleft;</code> <code> uint16_t yleft;</code>	X and Y coordinate of the center of the subject's left eye. If the eye coordinate is out of range (e.g. <code>x < 0</code> or <code>x >= width</code>), <code>isLeftAssigned</code> should be set to false.
<code> uint16_t xright;</code> <code> uint16_t yright;</code>	X and Y coordinate of the center of the subject's right eye. If the eye coordinate is out of range (e.g. <code>x < 0</code> or <code>x >= width</code>), <code>isRightAssigned</code> should be set to false.
<code>} EyePair;</code>	

243 **2.2.3. Template Role**

244 Labels describing the type/role of the template to be generated will be provided as input to template generation.

245 **Table 7 – Labels describing template role**

Label as C++ enumeration	Meaning
<code>enum class TemplateRole {</code>	
<code> Enrollment_11,</code>	Enrollment template for 1:1 matching
<code> Verification_11</code>	Verification template for 1:1 matching
<code>};</code>	

246 **2.2.4. Data type for similarity scores**

247 Identification and verification functions shall return a measure of the similarity between the face data contained in the
 248 two templates. The datatype shall be an eight byte double precision real. The legal range is `[0, DBL_MAX]`, where the
 249 `DBL_MAX` constant is larger than practically needed and defined in the `<climits>` include file. Larger values indicate more
 250 likelihood that the two samples are from the same person.

251 Providers are cautioned that algorithms that natively produce few unique values (e.g. integers on `[0,127]`) will be
 252 disadvantaged by the inability to set a threshold precisely, as might be required to attain a false match rate of exactly
 253 `0.0001`, for example.

254 **2.2.5. Data structure for return value of API function calls**

255 **Table 8 – Enumeration of return codes**

Return code as C++ enumeration	Meaning
enum class ReturnCode {	
Success=0,	Success
ConfigError=1,	Error reading configuration files
RefuseInput=2,	Elective refusal to process the input, e.g. because cannot handle greyscale
ExtractError=3,	Involuntary failure to process the image, e.g. after catching exception
ParseError=4,	Cannot parse the input data
TemplateCreationError=5,	Elective refusal to produce a template (e.g. insufficient pixels between the eyes)
VerifTemplateError=6,	For matching, either or both of the input templates were result of failed feature extraction
NumDataError=7,	The implementation cannot support the number of images
TemplateFormatError=8,	Template file is in an incorrect format or defective
GPUError=9,	There was a problem setting or accessing the GPU
VendorError=10	Vendor-defined failure. Failure codes must be documented and communicated to NIST with the submission of the implementation under test.
};	

256

257

Table 9 – ReturnStatus structure

C++ code fragment	Meaning
struct ReturnStatus {	
ReturnCode code;	Return Code
std::string info;	Optional information string
// constructors	
};	

258

2.2.6. Data structure for encapsulating training data

259

The following structure represents subject attributes that may be available to the implementation during training.

260

Table 10 – Structure containing subject metadata information

C++ code fragment	Meaning
typedef struct Attributes {	
enum class Gender {Unknown, Male, Female};	
enum class Race {Unknown, White, Black, EastAsian, SouthAsian, Hispanic};	
enum class EyeGlasses {Unknown, NotWearing, Wearing};	
enum class FacialHair {Unknown, Moustache, Goatee, Beard};	
enum class SkinTone {Unknown, LightPink, LightYellow, MediumPinkBrown, MediumYellowBrown, MediumDarkBrown, DarkBrown};	
std::string id;	A subject ID that identifies a person. Images of the same person will have the same subject ID.
double age;	Subject age (in years). A negative value indicates age is unknown.
Gender gender;	Subject gender
Race race;	Subject race. This value may be a proxy.
EyeGlasses eyeglasses;	Whether the subject is wearing eyeglasses
FacialHair facialhair;	Facial hair type if applicable
double height;	Subject height (in meters). A negative value indicates height is unknown.
double weight;	Subject weight (in kilograms). A negative value indicates weight is unknown.
SkinTone skintone;	Subject skin tone

```
} Attributes;
```

261

Table 11 – Structure representing face image and associated attributes

C++ code fragment	Remarks
using faceAttributePair = std::pair<Image, Attributes>;	A pair of face image and associated attributes

262

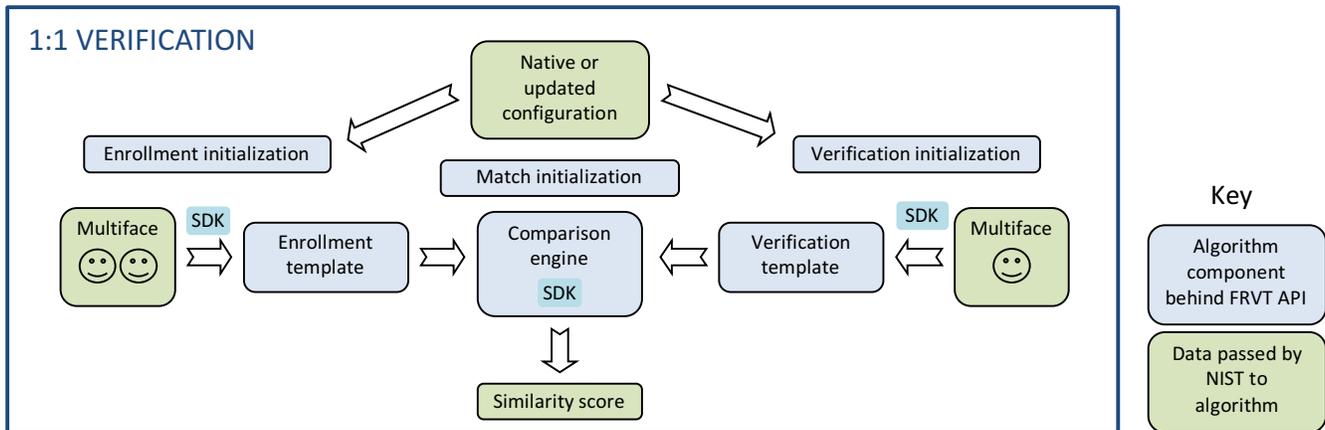
263 3. API Specification

264 3.1. Namespace

265 All data structures and API interfaces/function calls will be declared in the FRVT namespace.

266 3.2. Overview

267



268

269

Figure 1 – Schematic of 1:1 verification

270

271 The 1:1 testing will proceed in the following phases: optional offline training; preparation of enrollment templates;
 272 preparation of verification templates; and matching. Note that training, template creation, and matching may all be
 273 performed as separate processes. These are detailed in Table 12.

274

Table 12 – Functional summary of the 1:1 application

Phase	Description	Performance Metrics to be reported by NIST
Training (Optional)	Given 1) $K \geq 1$ images with associated subject ID and attribute data and 2) the implementation's configuration directory, the implementation may use the provided training data to populate a new "trained" configuration directory. This directory will be used to initialize the algorithm during subsequent template creation and matching processes. Images of the same person will have the same subject ID. Images with different subject IDs indicate they are different people. Attribute data may include the subject's age, gender, race, and other information. Please note that this function may or may not be called prior to creation of templates or matching. The implementation's ability to create or match templates should not be dependent on this function.	
Initialization	Function to read configuration data, if any.	None
Enrollment	Given $K \geq 1$ input images of an individual, the implementation will create a proprietary enrollment template. NIST will	Statistics of the time needed to produce a template. Statistics of template size. Rate of failure to produce a

	manage storage of these templates.	template
Verification	Given $K \geq 1$ input images of an individual, the implementation will create a proprietary verification template. NIST will manage storage of these templates.	Statistics of the time needed to produce a template. Statistics of template size. Rate of failure to produce a template.
Matching (i.e. comparison)	Given a proprietary enrollment and a proprietary verification template, compare them to produce a similarity score.	Statistics of the time taken to compare two templates. Accuracy measures, primarily reported as DETs, including for partitions of the input datasets.

275
 276 NIST requires that these operations may be executed in a loop in a single process invocation, or as a sequence of independent process
 277 invocations, or a mixture of both.

278 **3.3. API**

279 **3.3.1.1. Interface**

280 The software under test must implement the interface `Interface` by subclassing this class and implementing each
 281 method specified therein.

	C++ code fragment	Remarks
1.	<code>class Interface</code>	
2.	<code>{</code> <code>public:</code>	
3.	<code>virtual ReturnStatus initialize(</code> <code>const std::string &configDir) = 0;</code>	
4.	<code>virtual ReturnStatus createTemplate(</code> <code>const Multiface &faces,</code> <code>TemplateRole role,</code> <code>std::vector<uint8_t> &templ,</code> <code>std::vector<EyePair> &eyeCoordinates) = 0;</code>	
5.	<code>virtual ReturnStatus matchTemplates(</code> <code>const std::vector<uint8_t> &verifTemplate,</code> <code>const std::vector<uint8_t> &enrollTemplate,</code> <code>double &similarity) = 0;</code>	
6.	<code>virtual void ReturnStatus setGPU(uint8_t gpuNum) = 0;</code>	
7.	<code>static std::shared_ptr<Interface> getImplementation();</code>	Factory method to return a managed pointer to the <code>Interface</code> object. This function is implemented by the submitted library and must return a managed pointer to the <code>Interface</code> object.
8.	<code>virtual ReturnStatus train(</code> <code>const std::string &configDir,</code> <code>const std::string &trainedConfigDir,</code> <code>const std::vector<faceAttributePair> &faces) = 0;</code>	
9.	<code>};</code>	

282
 283 There is one class (static) method declared in `Interface.getImplementation()` which must also be implemented
 284 by the implementation. This method returns a shared pointer to the object of the interface type, an instantiation of the
 285 implementation class. A typical implementation of this method is also shown below as an example.
 286

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

```
#include "frvt11.h"

using namespace FRVT;

NullImpl:: NullImpl () { }

NullImpl::~ NullImpl () { }

std::shared_ptr<Interface>
Interface::getImplementation()
{
    return std::make_shared<NullImpl>();
}

// Other implemented functions
```

287 **3.3.2. Initialization**

288 The NIST test harness will call the initialization function in Table 13 before calling template generation or matching. **This**
 289 **function will be called BEFORE any calls to fork() are made.**

290 **Table 13 – Initialization**

Prototype	ReturnStatus initialize(const string &configDir);	Input
Description	This function initializes the implementation under test. It will be called by the NIST application before any call to createTemplate() or matchTemplates(). The implementation under test should set all parameters. This function will be called N=1 times by the NIST application, prior to parallelizing M >= 1 calls to createTemplate() via fork().	
Input Parameters	configDir	A read-only directory containing any developer-supplied configuration parameters or run-time data files. The name of this directory is assigned by NIST, not hardwired by the provider. The names of the files in this directory are hardwired in the implementation and are unrestricted.
Output Parameters	none	
Return Value	See Table 8 for all valid return code values.	

291 **3.3.2.1. GPU Index Specification**

292 For implementations using GPUs, the function of Table 14 specifies a sequential index for which GPU device to execute
 293 on. This enables the test software to orchestrate load balancing across multiple GPUs. **This function will be called AFTER**
 294 **a call to fork() is made.**

295 **Table 14 – GPU index specification**

Prototypes	void ReturnStatus setGPU (uint8_t gpuNum);	Input
Description	This function sets the GPU device number to be used by all subsequent implementation function calls. gpuNum is a zero-based sequence value of which GPU device to use. 0 would mean the first detected GPU, 1 would be the second GPU, etc. If the implementation does not use GPUs, then this function call should simply do nothing.	
Input Parameters	gpuNum	Index number representing which GPU to use.
Return Value	See Table 8 for all valid return code values.	

296 **3.3.3. Template generation**

297 The function of Table 15 supports role-specific generation of a template data. Template format is entirely proprietary.

298 **Table 15 – Template generation**

Prototypes	ReturnStatus createTemplate(const Multiface &faces,	Input
------------	---	-------

	TemplateRole role, std::vector<uint8_t> &templ, std::vector<EyePair> &eyeCoordinates);	Input Output Output
Description	Takes a Multiface and outputs a proprietary template and associated eye coordinates. The vectors to store the template and eye coordinates will be initially empty, and it is up to the implementation to populate them with the appropriate data. In all cases, even when unable to extract features, the output shall be a template that may be passed to the matchTemplates() function without error. That is, this routine must internally encode "template creation failed" and the matcher must transparently handle this.	
Input Parameters	faces	Implementations must alter their behavior according to the number of images contained in the structure and the TemplateRole type.
	role	Label describing the type/role of the template to be generated
Output Parameters	templ	The output template. The format is entirely unregulated. This will be an empty vector when passed into the function, and the implementation can resize and populate it with the appropriate data.
	eyeCoordinates	For each input image in the Multiface, the function shall return the estimated eye centers. This will be an empty vector when passed into the function, and the implementation shall populate it with the appropriate number of entries. Values in eyeCoordinates[i] shall correspond to faces[i].
Return Value	See Table 8 for all valid return code values.	

299 **3.3.4. Matching**

300 Matching of one enrollment against one verification template shall be implemented by the function of Table 16.

301 **Table 16 – Template matching**

Prototype	ReturnStatus matchTemplates(const std::vector<uint8_t> &verifTemplate, const std::vector<uint8_t> &enrollTemplate, double &similarity);	Input Input Output
Description	Compare two proprietary templates and output a similarity score, which need not satisfy the metric properties. When either or both of the input templates are the result of a failed template generation (see Table 15), the similarity score shall be -1 and the function return value shall be <code>VerifTemplateError</code> .	
Input Parameters	verifTemplate	A verification template from createTemplate(role=Verification_11). The underlying data can be accessed via verifTemplate.data(). The size, in bytes, of the template could be retrieved as verifTemplate.size().
	enrollTemplate	An enrollment template from createTemplate(role=Enrollment_11). The underlying data can be accessed via enrollTemplate.data(). The size, in bytes, of the template could be retrieved as enrollTemplate.size().
Output Parameters	similarity	A similarity score resulting from comparison of the templates, on the range [0,DBL_MAX]. See section 2.2.4.
Return Value	See Table 8 for all valid return code values.	

302 **3.3.1. Training**

303

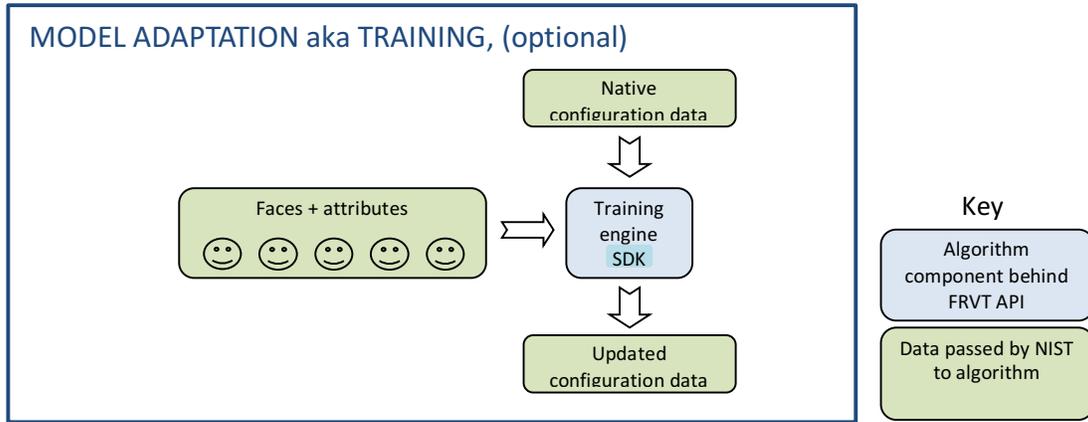


Figure 2 – Schematic of training

The NIST test harness may optionally call the training function in Table 17 as a separate process outside of the template generation and matching process. The implementation will be provided with the read-only configuration directory as supplied to NIST in the original submission, a read-write directory to store output(s) from training, and a set of face images and subject attributes where available.

Table 17 – Training

Prototype	ReturnStatus train(const std::string &configDir, const std::string &trainedConfigDir, const std::vector<faceAttributePair> &faces);	Input Input Input
Description	This function provides the implementation with face images and associated attributes where available. Attributes include a subject ID (this value is always assigned), and where available, subject data such as age, gender, race, and other information. Images of the same person will have the same subject ID. Genuine associations can be created using images with the same subject ID, and imposter associations can be derived using images with different subject IDs. This function may or may not be called prior to creation of templates or matching. The implementation’s ability to create or match templates should not be dependent on this function.	
Input Parameters	configDir	A read-only directory containing any developer-supplied configuration parameters or runtime data files. The name of this directory is assigned by NIST, not hardwired by the provider. The names of the files in this directory are hardwired in the implementation and are unrestricted.
	trainedConfigDir	A directory with read-write permissions where the implementation can store any training output. The name of this directory is assigned by NIST, not hardwired by the provider. The names of the files in this directory are hardwired in the implementation and are unrestricted. Important: This directory is what will subsequently be provided to the implementation’s initialize() function as the input configuration directory if this training function is invoked. Therefore, at a minimum, even if you choose not to implement this function, the necessary data from the original configuration configDir must be copied over into this directory.
	faces	A vector of face image-attribute pairs provided to the implementation for training purposes
Output Parameters	none	
Return Value	See Table 8 for all valid return code values.	

Purpose of training: Broadly NIST is seeking a repeatable and robust mechanism to provide end users with an easy to use, automated, mechanism to get the benefits of training on their own data. The training function is intended to improve some aspect of recognition.

NIST’s first attempt at exploiting the functionality of the train() API function call will be to address this problem: Some recognition algorithms give different imposter distributions for different age groups. So NIST will call train with thousands of images associated with an identity and age labels. An effective training mechanism would yield some configuration

318 data that allowed the recognition components (createTemplate() and matchTemplates()) to improve stability of the
319 impostor distribution across age groups. As a second test, we will then repeat this with race labels.

320

321 That said, developers can use this function for any purpose. You can assume that the tests that use the result of this step
322 will be with images of the same type as that passed to the function. The training and test sets will have disjoint sets of
323 people, reflecting the operational case where a training function should have utility over new users of a system.