

Face Recognition Vendor Test  
MORPH

Performance of Automated Facial Morph Detection and  
Morph Resistant Face Recognition Algorithms  
Concept, Evaluation Plan and API

VERSION 3.0

Mei Ngan  
Patrick Grother  
Kayee Hanaoka  
*Information Access Division  
Information Technology Laboratory*

May 19, 2022

**NIST**  
**National Institute of  
Standards and Technology**  
U.S. Department of Commerce

**Revision History**

1

<b>Date</b>	<b>Version</b>	<b>Description</b>
July 12, 2019	2.0	Initial document
September 9, 2020	2.0.1	Update link to General Evaluation Specifications document
July 7, 2021	2.1	Add optional <code>ageDeltaInDays</code> input argument to function <code>detectMorphDifferentially</code> (see Section 5.3.5)
May 19, 2022	3.0	<ul style="list-style-type: none"><li>- Remove optional <code>ageDeltaInDays</code> input argument to differential morph detection function in Section 5.3.5</li><li>- Add new function to support differential morph detection with additional subject metadata in Section 5.3.6</li></ul>

2

3 **Table of Contents**

4 **1. MORPH** ..... **4**

5 1.1. SCOPE ..... 4

6 1.2. GENERAL FRVT EVALUATION SPECIFICATIONS ..... 4

7 1.3. REPORTING ..... 4

8 1.4. ACCURACY METRICS ..... 4

9 **2. RULES FOR PARTICIPATION**..... **5**

10 2.1. IMPLEMENTATION REQUIREMENTS..... 5

11 2.2. PARTICIPATION AGREEMENT..... 5

12 2.3. NUMBER AND SCHEDULE OF SUBMISSIONS..... 5

13 2.4. VALIDATION ..... 5

14 **3. DATA STRUCTURES SUPPORTING THE API** ..... **5**

15 3.1. SUBJECT METADATA ..... 5

16 3.2. REQUIREMENT ..... 6

17 **4. IMPLEMENTATION LIBRARY FILENAME** ..... **6**

18 4.1. FILE FORMATS AND DATA STRUCTURES..... 6

19 4.1.1. *ImageLabel describing the format of an image* ..... 6

20 **5. API SPECIFICATION** ..... **7**

21 5.1. HEADER FILE..... 7

22 5.2. NAMESPACE ..... 7

23 5.3. API ..... 7

24 5.3.1. *Implementation Requirements*..... 7

25 5.3.2. *Interface*..... 7

26 5.3.3. *Initialization* ..... 8

27 5.3.4. *Single-image Morph Detection*..... 8

28 5.3.5. *Two-image Differential Morph Detection* ..... 9

29 5.3.6. *Two-image differential morph detection with subject metadata*..... 10

30 5.3.7. *1:1 Comparison*..... 11

31  
32

33 **List of Tables**

34 Table 1 – Structure for a single image ..... 6

35 Table 2 –Labels for subject sex ..... 6

36 Table 1 – Enumeration of image label ..... 6

37 Table 2 – API Functions ..... 7

38 Table 3 – Initialization ..... 8

39 Table 4 – Single-image Morph Detection ..... 9

40 Table 5 – Two-image Differential Morph Detection ..... 10

41 Table 6 – Two-image Differential Morph Detection ..... 11

42 Table 7 – 1:1 Comparison ..... 11

43  
44

## 45 1. MORPH

### 46 1.1. Scope

47 Facial morphing (and the ability to detect it) is an area of high interest to a number of photo-credential issuance  
 48 agencies and those employing face recognition for identity verification. The FRVT MORPH test will provide ongoing  
 49 independent testing of prototype facial morph detection technologies. The evaluation is designed to obtain an  
 50 assessment on morph detection capability to inform developers and current and prospective end-users. This  
 51 document establishes a concept of operations and an application programming interface (API) for evaluation of two  
 52 separate tasks:

- 53 1. Algorithmic capability to detect facial morphing (morphed/blended faces) in still photographs
  - 54 a. Single-image morph detection of non-scanned photos, printed-and-scanned photos, and images of  
 55 unknown photo format/origin
  - 56 b. Two-image differential morph detection of non-scanned photos, printed-and-scanned photos, and  
 57 images of unknown photo format/origin
- 58 2. Face recognition algorithm resistance against morphing

### 59 1.2. General FRVT Evaluation Specifications

60 General and common information shared between all Ongoing FRVT tracks are documented in the FRVT General  
 61 Evaluation Specifications document - [https://pages.nist.gov/frvt/api/FRVT\\_common.pdf](https://pages.nist.gov/frvt/api/FRVT_common.pdf). This includes rules for  
 62 participation, hardware and operating system environment, software requirements, reporting, and common data  
 63 structures that support the APIs.

### 64 1.3. Reporting

65 For all algorithms that complete the evaluation, NIST will provide performance results back to the participating  
 66 organizations. NIST may additionally report and share results with partner government agencies and interested  
 67 parties, and in workshops, conferences, conference papers, presentations and technical reports.

68 **Important:** This is a test in which NIST will identify the algorithm and the developing organization. Algorithm results  
 69 will be attributed to the developer. Results will be machine generated (i.e. scripted) and will include timing, accuracy  
 70 and other performance results. These will be provided alongside results from other implementations. Results will be  
 71 expanded and modified as additional implementations are tested, and as analyses are implemented. Results may be  
 72 regenerated on-the-fly, usually whenever additional implementations complete testing, or when new analyses are  
 73 added.  
 74

### 75 1.4. Accuracy metrics

76 This test will evaluate algorithmic ability to detect whether an image is a morphed/blended image of two or more  
 77 faces and/or to correctly reject 1:1 comparisons of morphed images against other images of the subjects used to  
 78 create the morph (but similarly, correctly authenticate legitimate non-morphed, mated pairs and correctly reject non-  
 79 morphed, non-mated pairs). Per established metrics<sup>1,2</sup> for assessment of morphing attacks, NIST will compute and  
 80 report:

---

<sup>1</sup> International Organization for Standardization: Information Technology – Biometric presentation attack detection – Part 3: Testing and reporting. ISO/IEC FDIS 30107-3:2017, JTC 1/SC 37, Geneva, Switzerland, 2017

<sup>2</sup> U. Scherhag, A. Nautsch, C. Rathgeb, M. Gomez-Barrero, R. Veldhuis, L. Spreeuwiers, M. Schils, D. Maltoni, P. Grother, S. Marcel, R. Breithaupt, R. Raghavendra, C. Busch: "Biometric Systems under Morphing Attacks: Assessment of Morphing Techniques and Vulnerability Reporting", in Proceedings of the IEEE 16th International Conference of the Biometrics Special Interest Group (BIOSIG), Darmstadt, September 20-22, (2017)

- 81 • Attack Presentation Classification Error Rate (APCER) – the proportion of morph attack samples incorrectly  
82 classified as bona fide presentation
- 83 • Bona Fide Presentation Classification Error Rate (BPCER) – the proportion of bona fide samples incorrectly  
84 classified as morphed samples
- 85 • Mated Morph Presentation Match Rate (MMPMR) - the proportion of comparisons where the morphed  
86 image successfully authenticates against all constituents
- 87 • True Acceptance Rate (TAR) – the proportion of non-morphed, mated comparisons that correctly  
88 authenticate
- 89 • False Match Rate (FMR) – the proportion of non-morphed, non-mated comparisons that incorrectly  
90 authenticate

91

92 We will report the above quantities as a function of alpha (the fraction of each subject that contributed to the morph),  
93 image compression ratio, image resolution, image size, and others.

94 We will also report error tradeoff plots (BPCER vs. APCER, MMPMR vs. FMR, parametric on threshold).

## 95 2. Rules for participation

### 96 2.1. Implementation Requirements

97 Developers are not required to implement all functions specified in this API. Developers may choose to implement  
98 one or more functions of this API – please refer to Section 5.3.1 for detailed information regarding implementation  
99 requirements.

### 100 2.2. Participation agreement

101 A participant must properly follow, complete, and submit the [FRVT MORPH Participation Agreement](#). This must be  
102 done once, either prior or in conjunction with the very first algorithm submission. It is not necessary to do this for  
103 each submitted implementation thereafter.

### 104 2.3. Number and Schedule of Submissions

105 Currently, the number and schedule of submissions is not regulated, so participants can send submissions at any time.  
106 NIST reserves the right to amend this section with submission volume and frequency limits. NIST will evaluate  
107 implementations on a first-come-first-served basis and provide results back to the participants as soon as possible.

### 108 2.4. Validation

109 All participants must run their software through the provided FRVT MORPH validation package prior to submission.  
110 The validation package will be made available at <https://github.com/usnistgov/frvt>. The purpose of validation is to  
111 ensure consistent algorithm output between the participant’s execution and NIST’s execution. Our validation set is  
112 not intended to provide training or test data.

## 113 3. Data structures supporting the API

114 The data structures supporting this API are documented in this section and in the FRVT - General Evaluation  
115 Specifications document available at – [https://pages.nist.gov/frvt/api/FRVT\\_common.pdf](https://pages.nist.gov/frvt/api/FRVT_common.pdf) with corresponding header  
116 file named *frvt\_structs.h* published at <https://github.com/usnistgov/frvt>.

### 117 3.1. Subject Metadata

118 Data structure representing information about a subject.

119

**Table 1 – Structure for a single image**

C++ code fragment	Remarks
typedef struct SubjectMetadata	
{	
Sex sex;	Sex of the subject
int16_t ageInMonths;	Age of subject (in months) in probe image; -1 indicates an unassigned value
int16_t ageDeltaInMonths;	Age/time difference (in months) between probe and reference image; -1 indicates an unassigned value
} SubjectMetadata;	

120

121

**Table 2 – Labels for subject sex**

Label as C++ enumeration	Meaning
enum class Sex {	
Unknown=0,	Either the label is unknown or unassigned
Female,	
Male,	
};	

122

**123 3.2. Requirement**

124 FRVT MORPH participants should implement the relevant C++ prototyped interfaces of section 5. C++ was chosen in  
 125 order to make use of some object-oriented features. Any functions that are not implemented should return  
 126 ReturnCode::NotImplemented.

**127 4. Implementation Library Filename**

128 The core library shall be named as libfrvt\_morph\_<provider>\_<sequence>.so, with  
 129 • provider: single word, non-infringing name of the main provider. Example: acme  
 130 • sequence: a three digit decimal identifier to start at 000 and incremented by 1 every time a library is sent to  
 131 NIST. Example: 007  
 132

133 Example core library names: *libfrvt\_morph\_acme\_000.so, libfrvt\_morph\_mycompany\_006.so.*  
 134 Important: Public results will be attributed with the provider name and the 3-digit sequence number in the submitted  
 135 library name.

**136 4.1. File formats and data structures**

**137 4.1.1. ImageLabel describing the format of an image**

138

**Table 3 – Enumeration of image label**

Return code as C++ enumeration	Meaning
enum class ImageLabel {	
Unknown=0,	Image origin is unknown or unassigned
NonScanned=1	Non-scanned photo
Scanned=2,	Printed-and-scanned photo
};	

139

140 **5. API specification**

141 Please note that included with the FRVT MORPH validation package (available at <https://github.com/usnistgov/frvt>) is  
 142 a “null” implementation of this API. The null implementation has no real functionality but demonstrates mechanically  
 143 how one could go about implementing this API.

144 **5.1. Header File**

145 The prototypes from this document will be written to a file named `frvt_morph.h` and will be available to implementers  
 146 at <https://github.com/usnistgov/frvt>.

147 **5.2. Namespace**

148 All supporting data structures will be declared in the `FRVT` namespace. All API interfaces/function calls for this track  
 149 will be declared in the `FRVT_MORPH` namespace.

150 **5.3. API**

151 **5.3.1. Implementation Requirements**

152 Developers are not required to implement all functions specified in this API. Developers may choose to implement  
 153 one or more functions of Table 4, but at a minimum, developers must submit a library that implements

- 154 1. `Interface` of Section 5.3.2,
- 155 2. `initialize()` of Section 5.3.3, and
- 156 3. AT LEAST one of the functions from Table 4. For any other function that is not implemented, the function  
 157 shall return `ReturnCode::NotImplemented`.

158 **Table 4 – API Functions**

Function	Section
<code>detectMorph()</code> – single image morph detection of <ul style="list-style-type: none"> <li>• Non-scanned photo</li> <li>• Printed-and-scanned photo</li> <li>• Image of unknown format</li> </ul>	5.3.4
<code>detectMorphDifferentially()</code> – two image differential morph detection of <ul style="list-style-type: none"> <li>• Non-scanned photo</li> <li>• Printed-and-scanned photo</li> <li>• Image of unknown format</li> </ul>	5.3.5
<code>compareImages()</code> – 1:1 comparison	5.3.6

159

160 **5.3.2. Interface**

161 The software under test must implement the interface `Interface` by subclassing this class and implementing AT  
 162 LEAST ONE of the methods specified therein.

	C++ code fragment	Remarks
1.	<code>Class MorphInterface</code>	
2.	<code>{</code> <code>public:</code>	

## FRVT MORPH

3.	<code>static std::shared_ptr&lt;Interface&gt; 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.
4.	<code>// Other functions to implement</code>	
5.	<code>};</code>	

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

C++ code fragment	Remarks
<pre>#include "frvt_morph.h"  using namespace FRVT_MORPH;  NullImpl:: NullImpl () { }  NullImpl::~ NullImpl () { }  std::shared_ptr&lt;Interface&gt; Interface::getImplementation() {     return std::make_shared&lt;NullImpl&gt;(); } // Other implemented functions</pre>	

### 166 5.3.3. Initialization

167 Before any morph detection or matching calls are made, the NIST test harness will call the initialization function of  
 168 Table 5. This function will be called BEFORE any calls to `fork()` are made. This function must be implemented.

169 **Table 5 – Initialization**

Prototype	ReturnStatus initialize( const std::string &configDir, const std::string& configValue);	
		Input
		Input
Description	This function initializes the implementation under test and sets all needed parameters in preparation for template creation. This function will be called N=1 times by the NIST application, prior to parallelizing M >= 1 calls to any morph detection or matching functions via <code>fork()</code> .  This function will be called from a single process/thread.	
Input Parameters	configDir	A read-only directory containing any developer-supplied configuration parameters or run-time data files.
	configValue	An optional string value encoding algorithm-specific configuration parameters. Developers may provide documentation for such configuration parameter(s) in their submission to NIST. Otherwise, the default value for this parameter will be an empty string.
Output Parameters	None	
Return Value	See General Evaluation Specifications document for all valid return code values. This function <u>must</u> be implemented.	

170

### 171 5.3.4. Single-image Morph Detection

172 The function of Table 6 evaluates morph detection on non-scanned photos, scanned photos, and photos of unknown  
 173 formats. A single image along with an associated image label describing the image format/origin is provided to the  
 174 function for detection of morphing. Both morphed images and non-morphed images will be used, which will support



175 measurement of a morph attack presentation classification error rate (APCER) with a bona fide presentation  
 176 classification error rate (BPCER).

177 **Non-scanned photos**

178 Non-scanned photos are digital images known to not have been printed and scanned back in. There are a number of  
 179 operational use-cases for morph detection on such digital images.

180 **Scanned photos**

181 While there are existing techniques to detect manipulation of a digital image, once the image has been printed and  
 182 scanned back in, it leaves virtually no traces of the original image ever being manipulated. So the ability to detect  
 183 whether a printed-and-scanned image contains a morph warrants investigation.

184 **Photos of unknown format**

185 In some cases, the format and/or origin of the image in question is not known, so images with “unknown” labels will  
 186 also be tested.

187

188 Multiple instances of the calling application may run simultaneously or sequentially. These may be executing on  
 189 different computers.

190

**Table 6 – Single-image Morph Detection**

Prototypes	ReturnStatus detectMorph( const Image &suspectedMorph, const ImageLabel &label, bool &isMorph, double &score);	
		Input
		Input
		Output
		Output
Description	This function takes an input image and associated image label describing the image format/origin, and outputs a binary decision on whether the image is a morph and a "morphiness" score on [0, 1] indicating how confident the algorithm thinks the image is a morph, with 0 meaning confidence that the image is not a morph and 1 representing absolute confidence that it is a morph.	
Input Parameters	suspectedMorph	Input Image
	label	ImageLabel (Section 4.1.1) describing the format of the input image <ul style="list-style-type: none"> <li>• NonScanned = non-scanned digital photo</li> <li>• Scanned = a photo that is printed, then scanned</li> <li>• Unknown = unknown photo format/origin</li> </ul>
Output Parameters	isMorph	True if image contains a morph; False otherwise
	score	A score on [0, 1] representing how confident the algorithm is that the image contains a morph. 0 means certainty that image does not contain a morph and 1 represents certainty that image contains a morph.
Return Value	See General Evaluation Specifications document for all valid return code values.  If this function is not implemented, the return code should be set to <code>ReturnCode::NotImplemented</code> .  If this function is not implemented for a certain type of image, for example, the function supports non-scanned photos but not scanned photos, then the function should return <code>ReturnCode::NotImplemented</code> when the function is called with the particular unsupported image type.	

191 **5.3.5. Two-image Differential Morph Detection**

192 Two face samples are provided to the function of Table 7 as input, the first being a suspected morphed facial image  
 193 and the second image representing a known, non-morphed face image of one of the subjects contributing to the  
 194 morph (e.g., live capture image from an eGate). This procedure supports measurement of whether algorithms can

195 detect morphed images when additional information (provided as the second supporting known subject image) is  
 196 provided.

197 Similar to single-image morph detection, the function of Table 7 will support non-scanned, scanned, and photos of  
 198 unknown format/origin. The input image type will be specified by the associated ImageLabel input parameter.

199 Multiple instances of the calling application may run simultaneously or sequentially. These may be executing on  
 200 different computers.

201 **Table 7 – Two-image Differential Morph Detection**

Prototypes	ReturnStatus detectMorphDifferentially( const Image &suspectedMorph, const ImageLabel &label, const Image &probeFace, bool &isMorph, double &score);	
		Input
		Input
		Input
		Output
		Output
Description	This function takes two input images - a known unaltered/not morphed image of the subject ( <code>probeFace</code> ) and an image of the same subject that's in question (may or may not be a morph) ( <code>suspectedMorph</code> ) with an associated image label describing the image format/origin. This function outputs a binary decision on whether <code>suspectedMorph</code> is a morph (given <code>probeFace</code> as a prior) and a "morphiness" score on [0, 1] indicating how confident the algorithm thinks the <code>suspectedMorph</code> is a morph, with 0 meaning confidence that the <code>suspectedMorph</code> is not a morph and 1 representing absolute confidence that it is a morph.	
Input Parameters	<code>suspectedMorph</code>	Input Image
	<code>label</code>	ImageLabel (Section 4.1.1) describing the format of the suspected morph image <ul style="list-style-type: none"> <li>• NonScanned = non-scanned digital photo</li> <li>• Scanned = a photo that is printed, then scanned</li> <li>• Unknown = unknown photo format/origin</li> </ul>
	<code>probeFace</code>	An image of the subject known not to be a morph (e.g., live capture image)
Output Parameters	<code>isMorph</code>	True if image contains a morph; False otherwise
	<code>score</code>	A score on [0, 1] representing how confident the algorithm is that the image contains a morph. 0 means certainty that image does not contain a morph and 1 represents certainty that image contains a morph.
Return Value	See General Evaluation Specifications document for all valid return code values.  If this function is not implemented, the return code should be set to <code>ReturnCode::NotImplemented</code> .  If this function is not implemented for a certain type of image, for example, the function supports non-scanned photos but not scanned photos, then the function should return <code>ReturnCode::NotImplemented</code> when the function is called with the particular unsupported image type.	

202 **5.3.6. Two-image Differential Morph Detection with Subject Metadata**

203 Two face samples are provided to the function of Table 8 as input, the first being a suspected morphed facial image  
 204 and the second image representing a known, non-morphed face image of one of the subjects contributing to the  
 205 morph (e.g., live capture image from an eGate). **In addition**, subject metadata is provided as input to the algorithm,  
 206 which includes sex, age of the subject (in months) at the time the probe image is taken, and the age/time difference  
 207 (in months) between the suspected morph and the live probe image. Operationally, this information might be derived  
 208 from data read from the machine readable zone of a passport for example. This procedure supports measurement of  
 209 whether algorithms can detect morphed images when additional subject metadata is provided.

210

211 Multiple instances of the calling application may run simultaneously or sequentially. These may be executing on  
 212 different computers.

213

**Table 8 – Two-image Differential Morph Detection with Subject Metadata**

Prototypes	ReturnStatus detectMorphDifferentially( const Image &suspectedMorph, const ImageLabel &label, const Image &probeFace, const SubjectMetadata &subjectMetadata, bool &isMorph, double &score);	
		Input
		Input
		Input
		Input
		Output
Description	This function takes two input images - a known unaltered/not morphed image of the subject ( <code>probeFace</code> ) and an image of the same subject that's in question (may or may not be a morph) ( <code>suspectedMorph</code> ) with an associated image label describing the image format/origin. Additionally, subject metadata is provided as input to the algorithm, which include sex, age of the subject (in months) at the time the probe image is taken, and the age/time difference (in months) between the suspected morph and the live probe image. This function outputs a binary decision on whether <code>suspectedMorph</code> is a morph (given <code>probeFace</code> as a prior) and a "morphiness" score on [0, 1] indicating how confident the algorithm thinks the <code>suspectedMorph</code> is a morph, with 0 meaning confidence that the <code>suspectedMorph</code> is not a morph and 1 representing absolute confidence that it is a morph.	
Input Parameters	suspectedMorph	Input Image
	label	ImageLabel (Section 4.1.1) describing the format of the suspected morph image <ul style="list-style-type: none"> <li>• NonScanned = non-scanned digital photo</li> <li>• Scanned = a photo that is printed, then scanned</li> <li>• Unknown = unknown photo format/origin</li> </ul>
	probeFace	An image of the subject known not to be a morph (e.g., live capture image)
	subjectMetadata	SubjectMetadata (Section 3.1) with information about the subject
Output Parameters	isMorph	True if image contains a morph; False otherwise
	score	A score on [0, 1] representing how confident the algorithm is that the image contains a morph. 0 means certainty that image does not contain a morph and 1 represents certainty that image contains a morph.
Return Value	See General Evaluation Specifications document for all valid return code values.  If this function is not implemented, the return code should be set to <code>ReturnCode::NotImplemented</code> .  If this function is not implemented for a certain type of image, for example, the function supports non-scanned photos but not scanned photos, then the function should return <code>ReturnCode::NotImplemented</code> when the function is called with the particular unsupported image type.	

214

215 **5.3.7. 1:1 Comparison**

216 Two face samples are provided to the function of Table 9 for one-to-one comparison of whether the two images are of  
 217 the same subject. The expected behavior from the algorithm is to be able to correctly reject comparisons of morphed  
 218 images against constituents that contributed to the morph. The goal is to show algorithm robustness against  
 219 morphing alterations when morphed images are compared against other images of the subjects used for morphing.  
 220 Comparisons of morphed images against constituents should return a low similarity score, indicating rejection of  
 221 match. Comparisons of unaltered/non-morphed images of the same subject should return a high similarity score,  
 222 indicating acceptance of match.

223

224 Multiple instances of the calling application may run simultaneously or sequentially. These may be executing on  
 225 different computers.

226 **Table 9 – 1:1 Comparison**

Prototypes	ReturnStatus compareImages(	
------------	-----------------------------	--

FRVT MORPH

	const Image &enrollImage, const Image &verifImage, double &similarity);	Input Input Output
Description	This function compares two images and outputs a similarity score. In the event the algorithm cannot perform the comparison operation, the similarity score shall be set to -1.0 and the function return code value shall be set appropriately.	
Input Parameters	enrollImage	The enrollment image
	verifImage	The verification image
Output Parameters	similarity	A similarity score resulting from comparison of the two images, on the range [0,DBL_MAX].
Return Value	See General Evaluation Specifications document for all valid return code values.  If this function is not implemented, the return code should be set to <code>ReturnCode : :NotImplemented</code> .	