

Face Recognition Vendor Test
Ongoing

Performance of Automated
Presentation Attack Detection (PAD) Algorithms
Concept, Evaluation Plan, and API

VERSION 0.1

Mei Ngan
Patrick Grother
Kayee Hanaoka
Austin Hom
Joyce Yang
*Information Access Division
Information Technology Laboratory*

March 23, 2022

Revision History

1

Date	Version	Description
March 23, 2022	0.1	Draft document for public comment

2

3 **Table of Contents**

4 **1. PAD 4**

5 1.1. SCOPE 4

6 1.2. GENERAL FRVT EVALUATION SPECIFICATIONS 5

7 1.3. REPORTING 5

8 1.4. ACCURACY METRICS 5

9 1.5. TIME LIMITS..... 6

10 **2. RULES FOR PARTICIPATION..... 6**

11 2.1. PARTICIPATION AGREEMENT..... 6

12 2.2. VALIDATION 6

13 **3. DATA STRUCTURES SUPPORTING THE API 6**

14 **4. IMPLEMENTATION LIBRARY FILENAME 6**

15 **5. API SPECIFICATION 6**

16 5.1. HEADER FILE..... 7

17 5.2. NAMESPACE 7

18 5.3. API 7

19 5.3.1. *Interface*..... 7

20 5.3.2. *Initialization* 7

21 5.3.3. *Presentation Attack Detection*..... 8

22

23

24 **List of Tables**

25 Table 1 – Software vs. hardware-based PAD 4

26 Table 2 – Processing time limits in milliseconds, per 1280 x 960 image 6

27 Table 3 – Initialization 7

28 Table 4 – Presentation Attack Detection 8

29

30

31 1. PAD

32 1.1. Scope

33 A presentation attack (PA), as defined by the ISO/IEC 30107¹ standard on biometric presentation attack detection, is
 34 “the presentation of an artefact or of human characteristics to a biometric capture subsystem in a fashion intended to
 35 interfere with system policy”. A presentation attack is often launched with the intent of *impersonation* (the user is
 36 trying to authenticate as a target identity) or *evasion* (the user is trying to fool the biometric system into not
 37 recognizing their true identity). The goals of impersonation include trying to gain positive access privilege as someone
 38 else, for example, trying to unlock someone’s cell phone or gain access to a facility. The goals of evasion are typically
 39 to conceal one’s true identity to evade recognition from say a watchlist, or to create a separate enrollment under a
 40 different name. Biometric systems can potentially be attacked by an unknown number of presentation attack
 41 instruments, and the number or type of attack instruments in existence is not well-known. Some examples of known
 42 presentation attack instruments include artificial “gummy” fingers², “replay” attacks where the attacker is holding a
 43 photo or video of someone’s face to the camera³, and iris photo and contact lens attacks⁴.

44 Presentation attack of face recognition systems (and the ability to detect it) is an area of high interest given the
 45 widespread deployment of face recognition systems, particularly in unmanned/unsupervised and remote enrollment
 46 and authentication scenarios. PAD capabilities generally fall into two categories – software-based and hardware-
 47 based PAD. Table 1 summarizes the relevance and applications of software vs. hardware-based PAD. The FRVT PAD
 48 test will provide ongoing independent testing of software-based facial PAD detection technologies. The evaluation is
 49 designed to assess software-based PA detection capability to inform developers and current and prospective end-
 50 users. Software-based PAD solutions operate only on the captured imagery. This document establishes an initial
 51 concept of operations and an application programming interface (API) for evaluation of algorithmic capability to
 52 detect facial presentation attack from still photographs and/or video frames.

53 **Note:** Hardware-based PAD solutions are currently out of scope in *FRVT PAD*. For developers interested in evaluation
 54 of hardware-based PAD capabilities, the [DHS Science and Technology Directorate](#) is planning a future technology
 55 demonstration to include testing of hardware-based PAD capabilities.

56 **Table 1 – Software vs. hardware-based PAD**

	Software-based PAD	Hardware-based PAD
Input	Image/video	Image/video + other non-standardized data or signals sensed by dedicated hardware
Mode of operation	Server-based PAD with non-face-aware capture device; offline PAD in existing/legacy systems	Client or edge-based PAD with dedicated face-aware capture device
Applications	Applications where capture processes and devices are not controlled or cannot be configured to perform PAD	Applications where hardware is controllable/configurable during the capture process to perform PAD

57

¹ [ISO/IEC 30107-1:2016](#) Information technology — Biometric presentation attack detection — Part 1: Framework

² Tsutomu Matsumoto, Hiroyuki Matsumoto, Koji Yamada, and Satoshi Hoshino "Impact of artificial "gummy" fingers on fingerprint systems", Proc. SPIE 4677, Optical Security and Counterfeit Deterrence Techniques IV, (19 April 2002);

<https://doi.org/10.1117/12.462719>

³ <https://mobidev.biz/blog/face-anti-spoofing-prevent-fake-biometric-detection>

⁴ Chaos Computer Club Berlin: Chaos Computer Clubs breaks iris recognition system of the Samsung Galaxy S8 (2017).

<https://www.ccc.de/en/updates/2017/iriden>

58 1.2. General FRVT Evaluation Specifications

59 General and common information shared between all Ongoing FRVT tracks are documented in the FRVT General
60 Evaluation Specifications document – https://pages.nist.gov/frvt/api/FRVT_common_2.0_draft.pdf. This includes
61 rules for participation, hardware and operating system environment, software requirements, reporting, and common
62 data structures that support the APIs.

63 1.3. Reporting

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

67
68 **Important:** NIST will publish the name of the developer’s organization, the algorithm identifiers, and the performance
69 results with attribution to the developer. Results will be machine generated (i.e. scripted) and will include timing,
70 accuracy and other performance results. These will be provided alongside results from other implementations. Results
71 will be expanded and modified as additional implementations are tested, and as analyses are implemented. Results
72 may be regenerated on-the-fly, usually whenever additional implementations complete testing, or when new analyses
73 are added.

74
75 **Note:** Due to data sensitivities, NIST does not intend on disclosing and describing the presentation attack instruments
76 used in our evaluation. We will report PA detection metrics for each PAI but without the name or description of the
77 PAI. This is intended to encourage broad PAD effectiveness across unknown PAs, and to discourage tuning to specific
78 attacks. This reflects the operational reality that attackers don’t advertise their methods.

79

80 1.4. Accuracy metrics

81 This test will evaluate algorithmic ability to detect whether an image or a video contains a presentation attack or not.
82 Per established metrics⁵ for assessment of presentation attacks, NIST will compute and report:

- 83 • Attack Presentation Classification Error Rate (APCER) – the proportion of presentation attack samples
84 incorrectly classified as bona fide presentation
- 85 • Bona Fide Presentation Classification Error Rate (BPCER) – the proportion of bona fide samples incorrectly
86 classified as presentation attack samples
- 87 • Attack Presentation Non-Response Rate (APNRR⁶) and Bona Fide Presentation Non-Response Rate (BPNRR) –
88 the proportion of presentation attack and bona fide samples, respectively, that do not generate a response
89 and fail to be processed by the algorithm software, whether it’s elective refusal to process the imagery or an
90 involuntary error. Failure to process events will be logged when the algorithm software returns a non-
91 successful return code from the PAD function, indicating that something went wrong while processing the
92 imagery for PAD.

93 We intend on reporting the above quantities for various presentation attack species.

94 We intend on incorporating failure to process events into the calculation of BPCER and APCER. All occurrences of
95 failure to process by an algorithm will be treated as if a presentation attack is detected with the confidence score set
96 to +1.

97 We will also publish error tradeoff plots (BPCER vs. APCER, parametric on threshold) and other analyses as
98 appropriate.

⁵ International Organization for Standardization: Information Technology – Biometric presentation attack detection – Part 3: Testing and reporting. [ISO/IEC FDIS 30107-3:2017](https://www.iso.org/standard/54441.html), JTC 1/SC 37, Geneva, Switzerland, 2017

⁶ ISO/IEC 30107-3 includes APNRR “proportion of attack presentations using the same PAI species that cause no response at the PAD subsystem or data capture subsystem” to quantify outcomes where the sensor is not even triggered by the presented PA sample. We use APNRR to quantify that for the PAD-subsystem, which here is the algorithm under test.

99 1.5. Time limits

100 The elemental functions of the implementations shall execute under the time constraints of Table 2. These time limits
 101 apply to the function call invocations defined in Section 5. Assuming the times are random variables, NIST cannot
 102 regulate the maximum value, so the time limits are median values. This means that the median of all operations
 103 should take less than the identified duration. NIST will publish duration statistics.

104 The time limits apply per image.

105 **Table 2 – Processing time limits in milliseconds, per 1280 x 960 image**

Function	
detectPA ()	5000 (1 core)

106

107 2. Rules for participation

108 2.1. Participation agreement

109 A participant must properly follow, complete, and submit the [FRVT Participation Agreement](#). This must be done once,
 110 either prior or in conjunction with the very first algorithm submission. It is not necessary to do this for each submitted
 111 implementation thereafter. **Note:** Organizations that have already submitted a participation agreement for FRVT
 112 Ongoing 1:1 do not need to send in a new participation agreement unless the organization updates their
 113 cryptographic signing key.

114 2.2. Validation

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

119 3. Data structures supporting the API

120 The data structures supporting this API are documented in the FRVT - General Evaluation Specifications document
 121 available at – https://pages.nist.gov/frvt/api/FRVT_common_2.0_draft.pdf with corresponding header file named
 122 *frvt_structs.h* published at https://github.com/usnistgov/frvt/blob/pad/common/src/include/frvt_structs.h.

123 4. Implementation Library Filename

124 The core library shall be named as `libfrvt_pad_<provider>_<sequence>.so`, with

- 125 • provider: single word, non-infringing name of the main provider. Example: acme
- 126 • sequence: a three digit decimal identifier to start at 000 and incremented by 1 every time a library is sent to
 127 NIST. Example: 007

128

129 Example core library names: *libfrvt_pad_acme_000.so*, *libfrvt_pad_mycompany_006.so*.

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

132 5. API specification

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

136 **5.1. Header File**

137 The prototypes from this document will be written to a file named `frvt_pad.h` and are be available to implementers at
 138 https://github.com/usnistgov/frvt/blob/pad/pad/src/include/frvt_pad.h.

139 **5.2. Namespace**

140 All supporting data structures will be declared in the `FRVT` namespace. All API interfaces/function calls for this track
 141 will be declared in the `FRVT_PAD` namespace.

142 **5.3. API**

143 **5.3.1. Interface**

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

	C++ code fragment	Remarks
1.	<code>Class PADInterface</code>	
2.	<code>{</code> <code>public:</code>	
3.	<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.
4.	<code>// Other functions to implement</code>	
5.	<code>};</code>	

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

	C++ code fragment	Remarks
	<code>#include <frvt_pad.h></code> <code>using namespace FRVT_PAD;</code> <code>NullImpl:: NullImpl () { }</code> <code>NullImpl::~ NullImpl () { }</code> <code>std::shared_ptr<Interface></code> <code>Interface::getImplementation()</code> <code>{</code> <code>return std::make_shared<NullImpl>();</code> <code>}</code> <code>// Other implemented functions</code>	

149 **5.3.2. Initialization**

150 Before any presentation attack detection calls are made, the NIST test harness will call the initialization function of
 151 Table 3. This function will be called BEFORE any calls to `fork()`⁷ are made. This function **must** be implemented.

152 **Table 3 – Initialization**

Prototype	ReturnStatus initialize(const std::string &configDir);	Input
Description	This function initializes the implementation under test. It will be called by the NIST application before any calls to the presentation attack detection functions of this API. The implementation under test should set all parameters.	

⁷ <http://man7.org/linux/man-pages/man2/fork.2.html>

	This function will be called N=1 times by the NIST application, prior to parallelizing M >= 1 calls to any other functions via <code>fork()</code> . This function will be called from a single process/thread.	
Input Parameters	<code>configDir</code>	A read-only directory containing any developer-supplied configuration parameters or run-time data files.
Output Parameters	None	
Return Value	See General Evaluation Specifications document for all valid return code values. This function <u>must</u> be implemented.	

153

154 **5.3.3. Presentation Attack Detection**

155 The function of Table 4 evaluates presentation attack detection on still photos and/or sequential video frames. A
 156 single image or a sequence of video frames is provided to the function for detection of a presentation attack. Both PA
 157 imagery and non-PA (bona fide) imagery will be used, which will support measurement of attack presentation
 158 classification error rate (APCER) with a bona fide classification error rate (BPCER). This function must be implemented.

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

161

Table 4 – Presentation Attack Detection

Prototypes	ReturnStatus detectPA(const Media &suspectedPA, bool &isPA, double &score);	
		Input
		Output
		Output
Description	This function takes a piece of input media containing an image or sequence of video frames and outputs a binary decision on whether the media represents a PA and a "padiness" score on [-1, 1] representing how confident the algorithm is that the piece of media contains a PA. A value of -1 means certainty that the media does not contain a PA, and +1 represents certainty that the media does contain a PA. A value near 0 will indicate uncertainty. Question to developers: Should this function call include an additional input parameter specifying PA intent? PAD implementations are often fielded in applications where the classes of risk are known. For example, in authentication, a primary concern is impersonation. In background-checks, the concern is of evasion, concealment. We could add a parameter that communicates "PA intent", for example: A value of 1 indicates the sample is an evasion PA or bona fide; a value of 2 indicates the sample is an impersonation PA or bona fide; a value of 3 would mean the software is not provided with that information.	
Input Parameters	<code>suspectedPA</code>	Input media of a single still image, or a sequence of video frames
Output Parameters	<code>isPA</code>	True if media contains a PA; False otherwise
	<code>score</code>	A real-valued score on [-1, 1] Developers are cautioned that if software only ever reports a few discrete values, the resulting error-tradeoff characteristic will have steps, such that end-users will not be able to set thresholds that finely target some objective (e.g. BPCER = 0.01). This limitation will not occur if the algorithm emits scores with a continuous distribution.
Return Value	See General Evaluation Specifications document for all valid return code values. This function <u>must</u> be implemented.	

162