

**NIST Interagency Report  
NISTIR 8584**

**Face Analysis Technology Evaluation (FATE)  
MORPH**

*Part 4B: Considerations for Implementing Morph Detection in  
Operations*

Mei Ngan  
Patrick Grother

This publication is available free of charge from:  
<https://doi.org/10.6028/NIST.IR.8584>

**NIST Interagency Report  
NISTIR 8584**

**Face Analysis Technology Evaluation (FATE)  
MORPH**

*Part 4B: Considerations for Implementing Morph Detection in  
Operations*

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

This publication is available free of charge from:  
<https://doi.org/10.6028/NIST.IR.8584>

August 2025



U.S. Department of Commerce  
Howard Lutnick, Secretary

National Institute of Standards and Technology  
Craig Burkhardt, Acting Under Secretary of Commerce for Standards and Technology and Acting NIST Director

Certain equipment, instruments, software, or materials, commercial or non-commercial, are identified in this paper in order to specify the experimental procedure adequately. Such identification does not imply recommendation or endorsement of any product or service by NIST, nor does it imply that the materials or equipment identified are necessarily the best available for the purpose.

**NIST Technical Series Policies**

[Copyright, Use, and Licensing Statements](#)

[NIST Technical Series Publication Identifier Syntax](#)

**Publication History**

Approved by the NIST Editorial Review Board on 2025-08-14

**How to Cite this NIST Technical Series Publication**

Mei Ngan and Patrick Grother (2025) FATE MORPH Part 4B: Considerations for Implementing Morph Detection in Operations. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Interagency Report (NISTIR) 8584.

<https://doi.org/10.6028/NIST.IR.8584>

**Contact Information**

[frvt@nist.gov](mailto:frvt@nist.gov)

## Executive Summary

Morphing of face photographs presents a threat to identity processes; for example, two people being able to use one passport. The threat exists in applications that allow users to submit their own photos where digital history is unknown. The threat can be reduced significantly if photo capture is trusted. When that cannot be implemented, software-based morph detection algorithms can be deployed. Automated morph detection accuracy has improved [1] to the point that agencies considering deployment can weigh capability alongside prior risks that morphs are present in their operational flows, costs of not detecting those, costs of deployment, and costs of false detections. This document guides organizations in detecting, then investigating morphed face photographs in their operational settings.

## Scope

This document is intended to build awareness of the existence of morphs. It envisages applications that involve automated processing of photographs submitted when applying for an ID document such as a passport, and others that involve human subjects presenting photographs in real-time, for example a travel document to a border control authority. Such agencies can employ morph attack detection tools, face recognition engines, and staff trained to inspect, compare, and authenticate documents and human faces.

This document is intended to guide organizations toward effective deployment of tools and practices in situations where morphed photographs are a concern in operational workflows. It includes guidelines for what agencies might consider doing after a morph detector generates a positive indication or a suspicious photo is detected through human review.

## Keywords

face morphing , face analysis technology evaluation, FATE, morph attack, morph attack detection, identity proofing and identification.

## Table of Contents

<b>1. Introduction.....</b>	<b>1</b>
<b>2. Background.....</b>	<b>1</b>
2.1. Morphs, threats, consequences .....	1
2.1.1. Morph creation.....	2
2.1.2. The nature of morphs .....	3
<b>3. Morph detection .....</b>	<b>6</b>
3.1. Automated morph detection tools .....	6
3.2. State-of-the-art.....	6
<b>4. Configuring a morph detector.....</b>	<b>7</b>
<b>5. Investigatory techniques .....</b>	<b>9</b>
5.1. Visual inspection of image(s).....	9
5.1.1. Inspection for presence of artifacts.....	9
5.1.2. Inspection for absence or difference of features .....	10
5.1.3. Inspection of image file metadata.....	11
5.2. Use of a different biometric characteristic for identity verification .....	11
5.3. Use of face recognition .....	12
5.3.1. Applying 1:1 face recognition to a pair of images .....	12
5.3.2. Applying 1:N face recognition to search a database .....	13
<b>6. Procedures for investigation of candidate morphs.....</b>	<b>14</b>
6.1. Procedures after an S-MAD detection .....	14
6.2. Procedures after a D-MAD detection.....	15
<b>7. Preventing morphs from entering operational systems .....</b>	<b>16</b>
7.1. Attended live enrollment .....	16
7.2. Unattended trusted capture .....	16
7.3. Apply morph detection .....	16
<b>8. Additional training .....</b>	<b>17</b>
<b>9. References .....</b>	<b>18</b>

## List of Figures

<b>Figure 1 – Morph that is readily recognized as such because we are familiar with both persons involved .....</b>	<b>1</b>
<b>Figure 2 – Morph that is difficult to recognize because most of us are not familiar with the persons involved ..</b>	<b>2</b>
<b>Figure 3 – Morph examples .....</b>	<b>3</b>
<b>Figure 4 – Examples of morphing artifacts.....</b>	<b>4</b>
<b>Figure 5 – Morph after post-processing.....</b>	<b>4</b>
<b>Figure 6 – Morph of demographically different subjects .....</b>	<b>5</b>
<b>Figure 7 – Morph of demographically similar subjects .....</b>	<b>5</b>
<b>Figure 8 – Comparing S-MAD vs. D-MAD algorithm performance [21] .....</b>	<b>7</b>
<b>Figure 9 – Absence of features – scars.....</b>	<b>10</b>
<b>Figure 10 – Absence of features – moles .....</b>	<b>10</b>
<b>Figure 11 – Difference of features – ears.....</b>	<b>11</b>
<b>Figure 12 – Illustrative face recognition similarity score distributions .....</b>	<b>12</b>
<b>Figure 13 – Morph detection using 1:N face recognition .....</b>	<b>13</b>

## **Acknowledgements**

The authors would like to thank the Department of Homeland Security's Science and Technology Directorate (S&T) Biometric and Identity Technology Center and Office and Biometric Identity Management (OBIM) and the Federal Bureau of Investigation (FBI) for their support of this activity.

The authors are also grateful to the staff in the NIST Biometrics Research Laboratory for infrastructure supporting rapid evaluation and results analysis.

And last but never the least, the authors would like to thank Frøy Løvåsdal, Richard Vorder Bruegge, Ryan Galluzzo, Jason Prince, Adam Forman, and Louis-Guillaume Rigaud for their time and candor in reviewing the contents of this report and for providing many meaningful suggestions to help make this a better publication.

## 1. Introduction

Face morphing is an image manipulation technique where two or more human faces are blended or merged into a single photograph. Since morphing was first described in the academic literature in “The Magic Passport” [2], there has been a global, concerted effort to develop effective countermeasures. The U.S. Department of Homeland Security’s Science and Technology Directorate has funded work on both morph databases and software-based morph detectors [3, 4]. Concurrently, the U.S. Department of State and Australian Department of Foreign Affairs and Trade funded additional datasets which have supported evaluation efforts run by the National Institute of Standards and Technology (NIST) to establish a competitive benchmark [5] for tracking detection technologies. The European Community also tackled the problem in its iMars [6] project to “provide image morphing and manipulation attack detection solutions for the evaluation of ID document authenticity”. Collectively, these initiatives have heightened awareness of morphing threats, produced effective detection tools, and established rigorous assessment facilities.

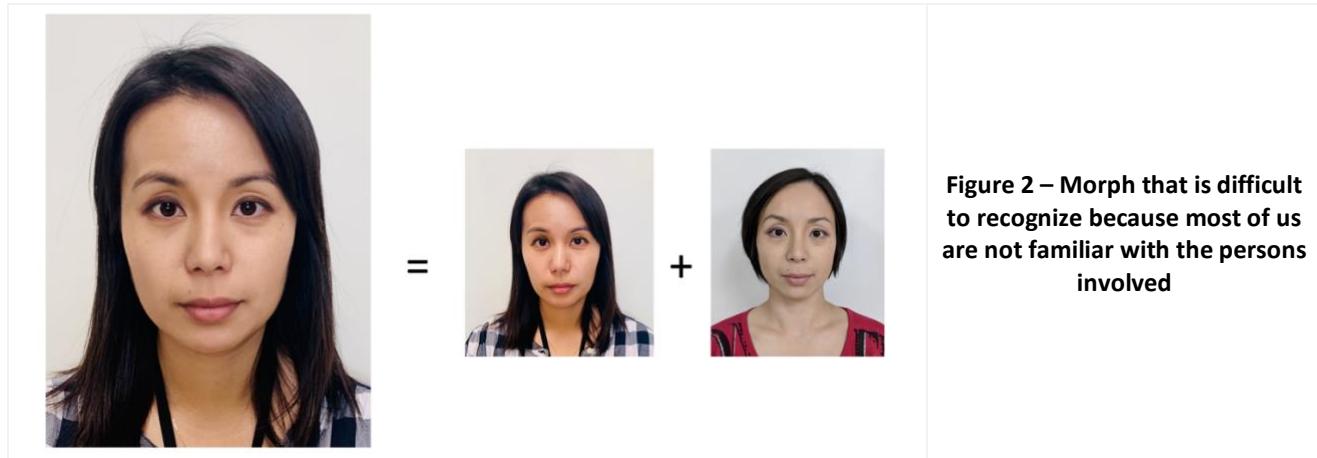
A primary driver for these research and development efforts has been the concern that country passport issuance processes are vulnerable to morphs, which can result in the issuance of a “fraudulently obtained genuine” ID document that can be used by multiple people at borders. A secondary driver is that domestic issuance or renewal processes are vulnerable. A third driver is remote enrollment and ID issuance processes which allow applicants to use their own mobile devices to submit potentially manipulated photos. This document is focused on the first two drivers. The prevalence of morphs is unknown and difficult to quantify without technical means for detection. Those cases that have been detected [7, 8, 9] were revealed through traditional border security and immigration processes.

## 2. Background

### 2.1. Morphs, threats, consequences



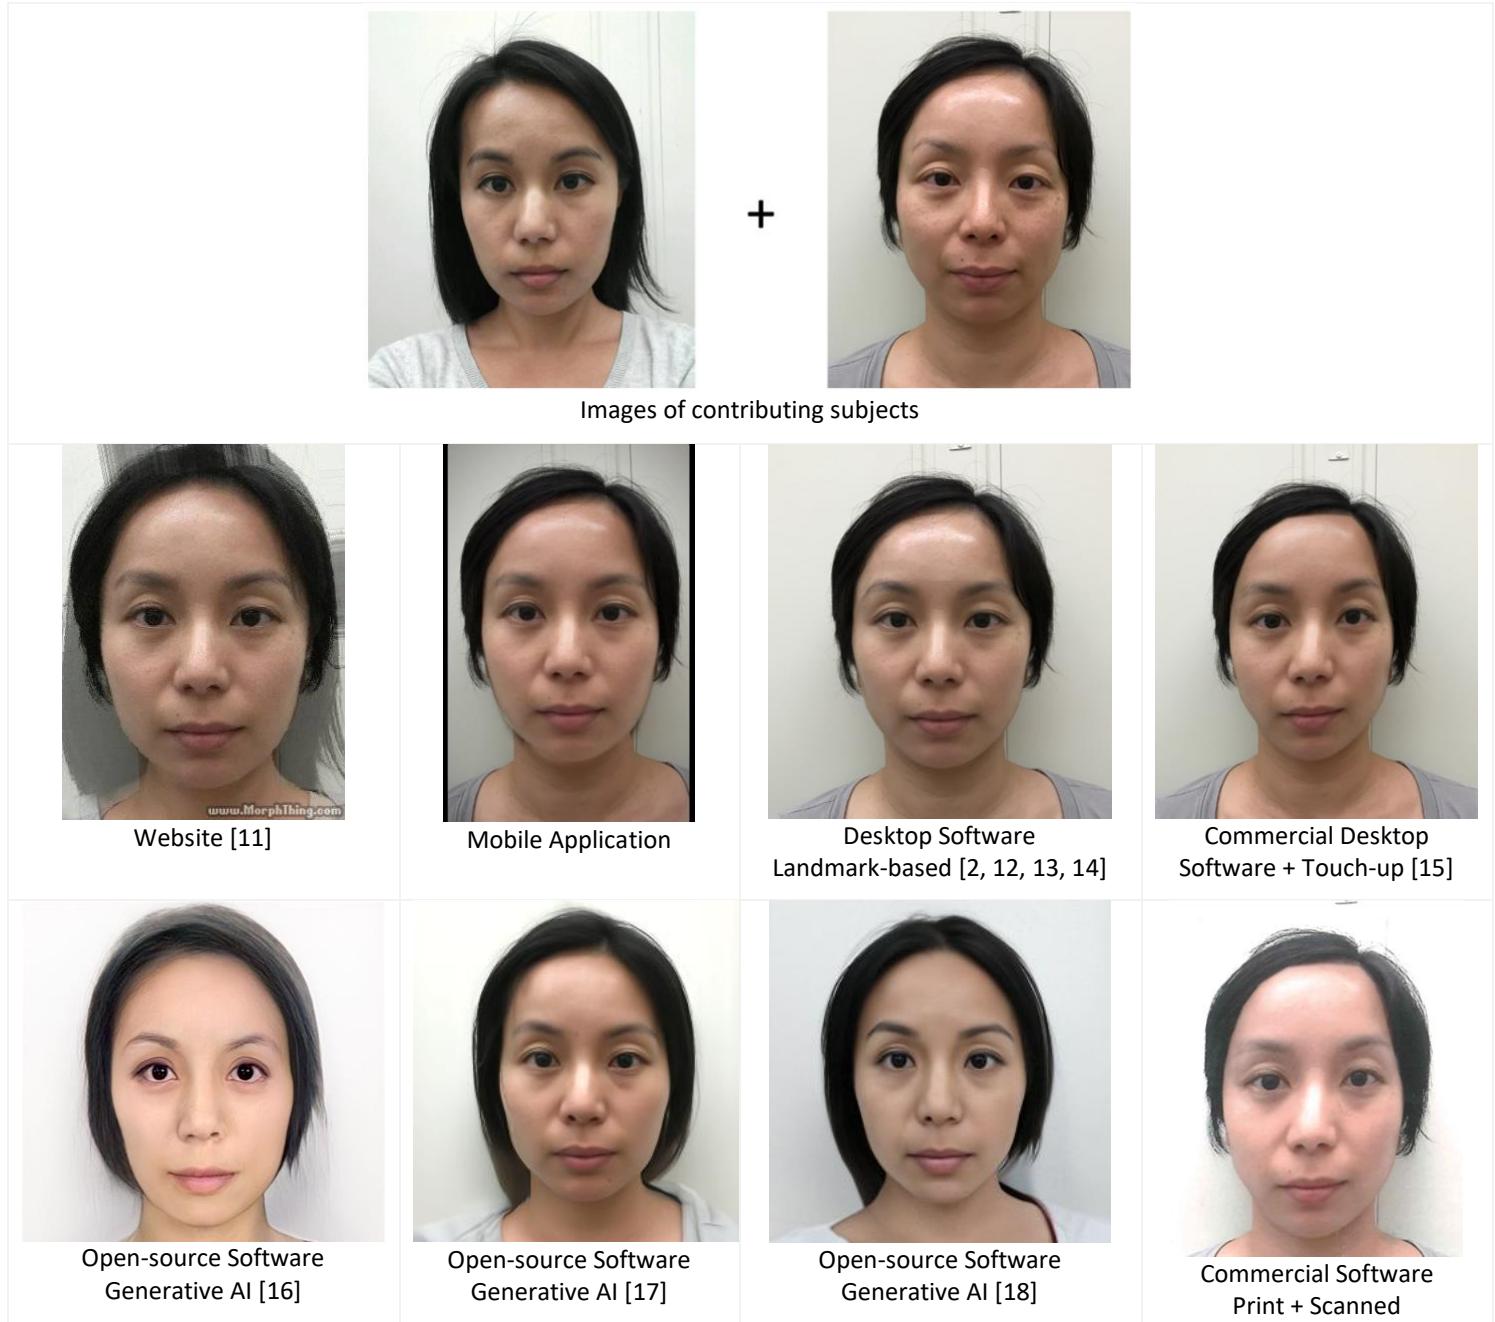
A morphed photograph is a combination of two or more human faces into one photograph. Morphs can be produced using mobile-phone apps, desktop graphics packages, or machine learning-based tools. In the example of Figure 1, human observers will note the resemblance of the morph to the constituent people, because they know the individuals involved. However, in the case of Figure 2 the constituent people are unknown to the observer, so the morph is perceived to be a photograph of just another unknown person.



The main threat associated with morphs is that if the submitted photograph is put on an identity document such as a passport, then both individuals will be able to use that document. Specifically, a human observer, such as a border guard or immigration officer will note the resemblance of the morph to the presenting traveler and likely allow passage. A face recognition system similarly will compute a high enough similarity between the morph and a live authentication photo that it too allows passage. The result is that **two or more people can use one ID document**. In an immigration context, if the passport can be mailed internationally, then both travelers can enter various countries. In a remote identity verification process, morphs can be added to a fake document to defeat the biometric matching process.

### 2.1.1. Morph creation

There are many software packages that will merge two photos into a morph. Some are free, some cost money. Some packages execute on general purpose computers; others can be downloaded from app stores for execution on mobile phones or tablets. Some methods for morphing are based on graphical techniques (e.g., landmark-based approaches); others are based on machine learning and/or neural networks (e.g., generative AI). Such morphing tools are widely accessible to the public and do not require high levels of technical expertise to use. Once produced, morphed photos may be touched-up using photo-editing software packages and/or post-processed (e.g., re-digitized by printing and scanning). Different face recognition systems will have different sensitivities to morphing depending on how the morph was created, and standards are being developed to help quantify the resistance of biometric systems to morphing attacks [10].



**Figure 3 – Morph examples**

### 2.1.2. The nature of morphs

Software packages vary in the quality and type of the morphs they produce. The term quality here mostly relates to how many artifacts or “tells” are present in the morphed photo. This can include artifacts around the iris, nostrils, lips, eyebrows, inconsistent skin texture and color, and/or anatomically impossible features – see Figure 4. Artifacts may be absent, or not visible to a human

observer. If they are present, they may be partially or completely removed in post-processing steps including touch-up, and printing on paper and scanning from paper – see Figure 5.

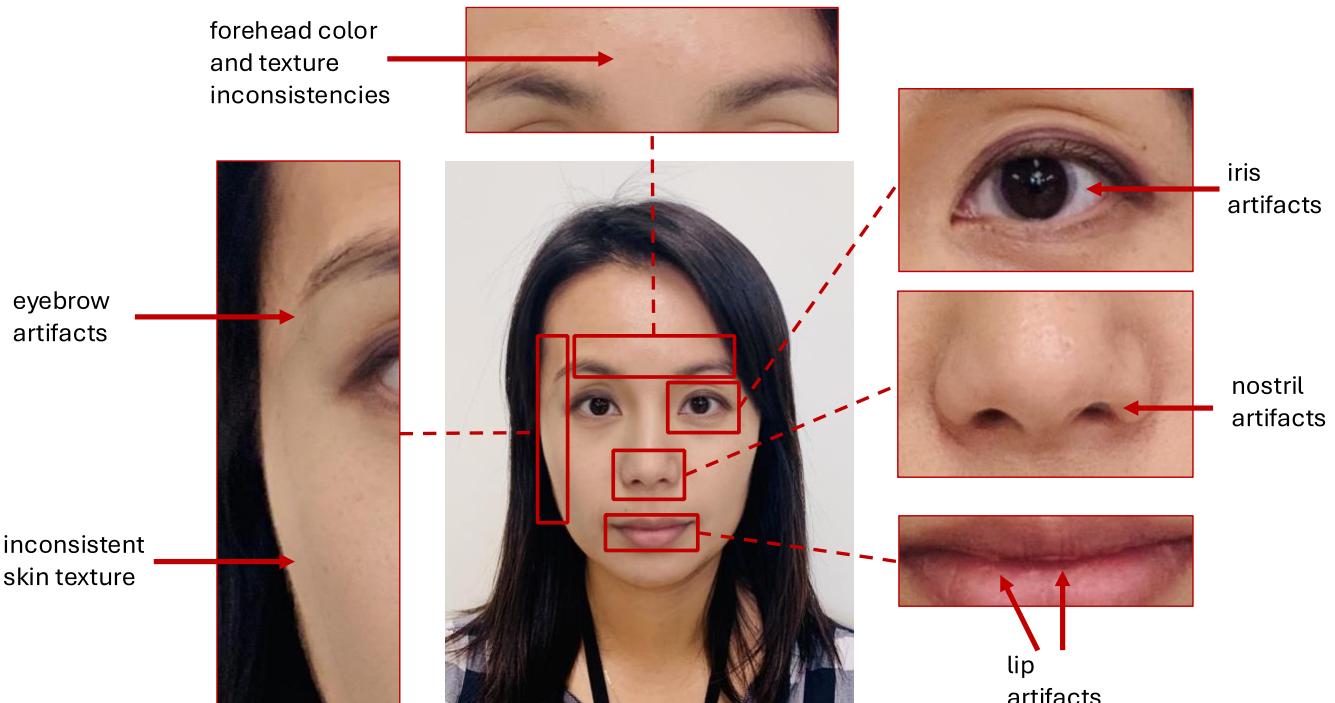


Figure 4 – Examples of morphing artifacts

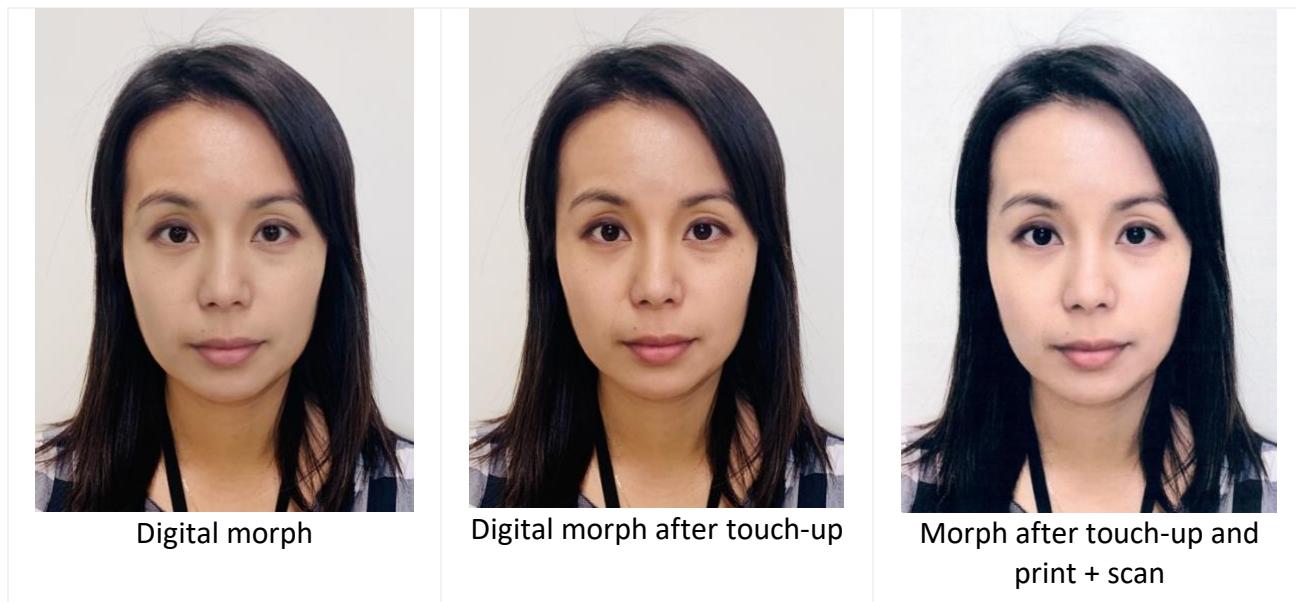
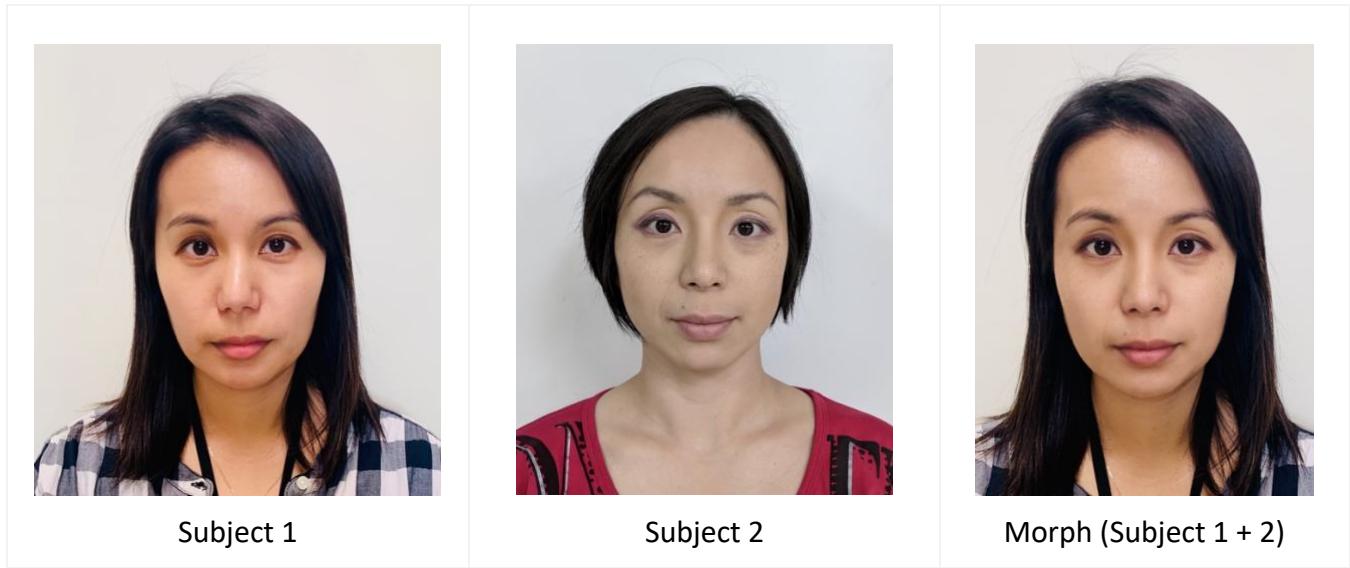


Figure 5 – Morph after post-processing

Morphing methods combine photos of two or more individuals. Software packages will produce morphs for any pair of individuals regardless of their appearance or similarity, and some tools will morph any number of input faces together. However, the most relevant case is the morphing of two individuals.



**Figure 6 – Morph of demographically different subjects**



**Figure 7 – Morph of demographically similar subjects**

A morph will be most effective if both morphed individuals, when challenged, can plausibly state that the photo is a legitimate photo of them alone. For example, if a morph is produced from photos of a

man and woman (e.g., Figure 6), the facial similarity will usually be low, such that an alert human official would notice. If, however, a morph is produced from photos of people who share the same demographics (e.g., same sex, similar age and race – see Figure 7) then human detection of the morph will be considerably more challenging. Similarly, detection of morphs of same-sex siblings or twins would be even more difficult because of natural facial similarity.

### 3. Morph detection

#### 3.1. Automated morph detection tools

There are two classes of automated morph attack detectors (MAD), and they work in different ways.

- **S-MAD** *Single-image morph attack detectors* operate solely on a suspected morph photograph, such as that accompanying a passport application. They apply various image processing and pattern classification techniques to detect artifacts produced by common morphing methods.
- **D-MAD** *Differential morph attack detectors* operate on a suspected morph and a constituent photo that is known not to be a morph, such as a photo collected in an automated border control (ABC) gate or photo collected at primary immigration for authentication to a passport photo.

#### 3.2. State-of-the-art

A number of S-MAD and D-MAD approaches [19] have been developed over the past decade. These have been tested and benchmarked in public evaluations including NIST’s FATE MORPH [5] and UNIBO’s BOEP [20]. NIST has been publishing public results on the performance of morph detection algorithms tested on sequestered data since 2019 on an ongoing basis. The use of sequestered data that developers have never seen or trained on before is standard practice in many of NIST’s biometric evaluations, allowing for truly independent testing of technological capabilities, their generalizability, and limitations. The FATE MORPH benchmark includes morph datasets of escalating difficulty, ranging from rudimentary morphs with clearly visible morphing artifacts to high quality morphs where little to no traces of manipulation are visible to the human eye. The benchmark includes data generated with morphing methods that are available in the public domain, as well as sequestered morphing methods.

Automated morph detection was primarily spearheaded by academic research and in recent years advanced further by emerging commercial development and capabilities. While accuracy has improved, the current [state-of-the-art shows strengths and weaknesses](#) [21] in the different morph detection approaches – see Figure 8.

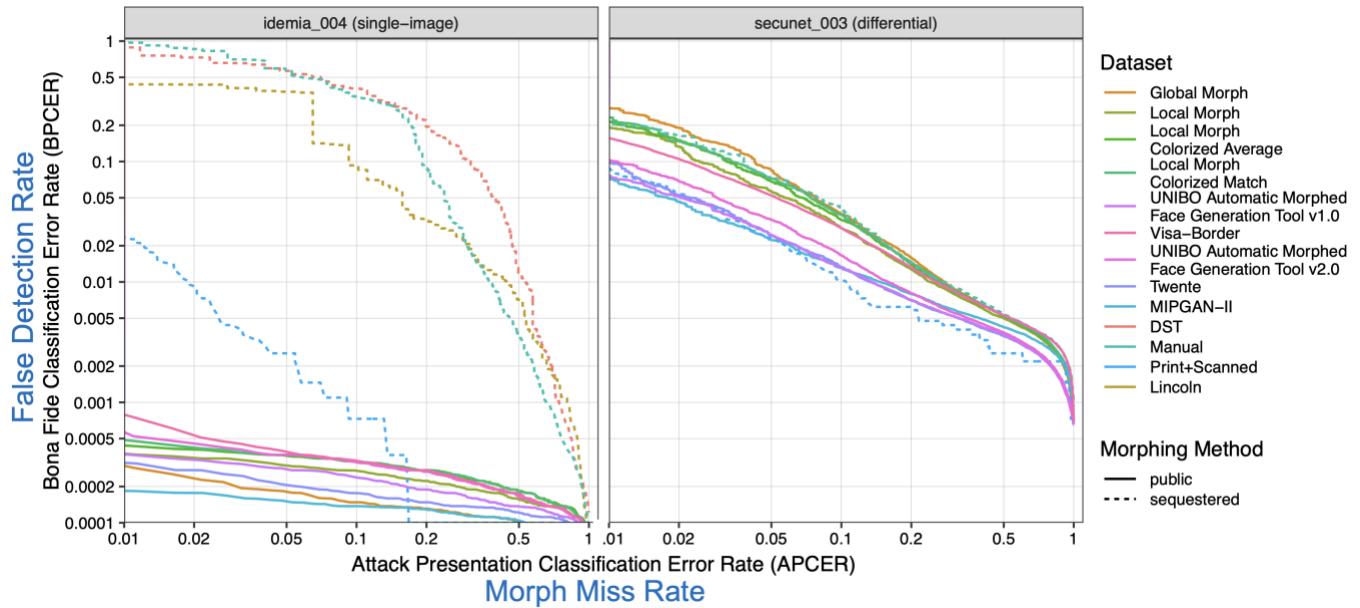


Figure 8 – Comparing S-MAD vs. D-MAD algorithm performance [21]

- S-MAD algorithms have been found [1] to be effective on morphs from a particular morph generator, if they have been trained on examples from that generator. For not-previouslly-seen morph types, morph detection accuracy is often very poor. This makes S-MAD currently not as generalizable across unseen morphing methods.
- D-MAD algorithms have been found [1] to be more generalizable in that accuracy is more consistent across all morphing methods tested, both public domain and sequestered methods. This is because they are usually built on technology that is used for face recognition, often leveraging identity information between the photo in question and the live probe image rather than specific image artifacts or noise patterns from a single image. While D-MAD can be used to test any two images against one another, D-MAD algorithms can be used most effectively when an additional live capture constituent photo is available.

#### 4. Configuring a morph detector

The configuration of a morph detector can depend on a number of variables, such as the expected level of threat or risk of an attack, priority of customer convenience vs. inconvenience, resource availability, and alignment with existing subsystems/components such as a face recognition system. Organizations should adopt a risk-based approach that is aligned with their specific operational context and defined risk tolerance levels. In this section, we provide illustrative examples of how a morph detector might be configured in different scenarios; these scenarios are intended solely for illustrative purposes, not as comprehensive guidance.

## High Threshold

In operations where an organization has conducted risk analysis and has estimated that the likelihood of morph fraud is low, the organization might consider configuring the morph detection software to minimize false positives. This assists with aligning the number of investigations triggered with the available resources to carry them out. A low false positive rate can be achieved by setting a **high threshold** for the MAD output score. However, increasing the threshold also increases the risk of missing actual morphs. This tradeoff between false positives and missed detections can be analyzed through performance tests of MAD implementations [1].

### Example Scenario #1: Document issuance:

Consider a passport agency that has estimated the prevalence of morphs within its operations and uses a morph detection algorithm whose speed and accuracy have been validated in prior testing. Given the agency's daily application volume and the limited capacity of its review team to investigate flagged cases, the algorithm might be configured with a threshold that keeps the false detection rate within manageable limits. A risk-based approach could involve setting this threshold in a way that balances security objectives with the operational impact of false detections. External evaluations (e.g., NIST MORPH [1], UNIBO BOEP [20]) can provide initial insight into how different threshold values may meet this condition. Operational testing conducted by the organization can inform further adjustments (e.g., threshold configuration changes, additional resources to support investigation load, etc.). To match detection performance to operational capacity, the agency might set the threshold high enough to limit the number of false positives to a level that the team can reasonably handle each day (i.e., the number of false detections per day is kept within the maximum daily review capacity).

### Example Scenario #2: Automated border control (ABC) gate:

Consider an airport operation using ABC gates that has estimated the prevalence of morphs in its processes and determined, through operational testing, that most transactions involve legitimate travelers using their own passports. ABC gates rely on automated face recognition software, and in high-volume, low-fraud environments, these systems are generally configured to minimize false rejections. Adding a morph detection capability may increase the overall rejection rate. A risk-based approach could involve estimating the potential impact of this addition using external evaluation data (e.g., NIST MORPH [1], UNIBO BOEP [20]) and then refining those estimates through operational testing. In some cases, the combined effect of both systems may lead to noticeably more travelers being incorrectly rejected. To address this while still mitigating morph-based fraud, the agency might adjust the face recognition system to be slightly less strict, reducing its contribution to false rejections while accepting a possible trade-off in increased impostor acceptance. Another approach could be if the ABC gate's organizational risk assessment processes determine the risk levels for different countries, they may want to take that into account when processing passports from different countries. For instance,

if the passport was issued by a country known to use live, attended enrollment processes (where photos are securely collected and less likely to be manipulated - see sections 7.1 and 7.2), the system could consider skipping morph detection for those travelers, reducing false alarms. Other strategies may also be possible.

## Low Threshold

Setting a high threshold reduces the number of false positives and minimizes operational burden, which may be appropriate in low-risk environments where morphing attacks are unlikely. However, this also increases the risk of failing to detect actual morphs. In higher-risk operational contexts, such as when the likelihood of morphing fraud is elevated, it may be appropriate to **lower the threshold** to increase detection sensitivity. While this increases the chance of identifying morph attacks, it also results in more false positives, requiring greater investigative capacity and operational resources. Threshold selection should therefore be guided by an entity's fraud risk profile, resource availability and capability, tolerance for false positives, and other factors.

**Example Scenario #3: Remote enrollment:** In remote enrollment processes, there is risk that an applicant could present, for example, a printed morph image to the camera. If the morph detection software flags such a transaction as suspicious, the applicant could then be redirected to complete their enrollment in person at a trusted facility with staff oversight. This approach assumes that in-person enrollment is widely accessible and available. Configuring the morph detector to prioritize security will likely result in more legitimate applications being incorrectly flagged and unnecessarily redirected to in-person enrollment.

## 5. Investigatory techniques

If a morph detector flags a photo as being a morph, additional steps may need to be taken by a human officer or examiner to validate the decision of the morph detector and to confirm the traveler/applicant's true identity. This process may be non-trivial and may require escalation to specialists with additional system accesses and training.

### 5.1. Visual inspection of image(s)

#### 5.1.1. Inspection for presence of artifacts

Morphing software vary in the quality of the morphs they produce. The term quality here mostly relates to how many artifacts are present in the morphed photo. This can include artifacts around the iris, nostrils, lips, eyebrows, inconsistent skin texture and color, and/or anatomically impossible features – see Figure 4 for examples of artifacts commonly seen in morphs.

### 5.1.2. Inspection for absence or difference of features

Facial features such as scars, marks/moles, and tattoos have some degree of permanence. If a photo captured today does not include a feature evident in a prior photo, then it could indicate a morph. This, of course, is not conclusive if the prior photo preceded the injury that caused the scar, or the date of getting the tattoo. Likewise, absence of a feature may be inconclusive if the subject claimed to have had surgery remove it, or it may be concealed by cosmetics, etc. If the ears are visible in both photos, check for differences between ear patterns and structure as this could also be evidence of morphing.

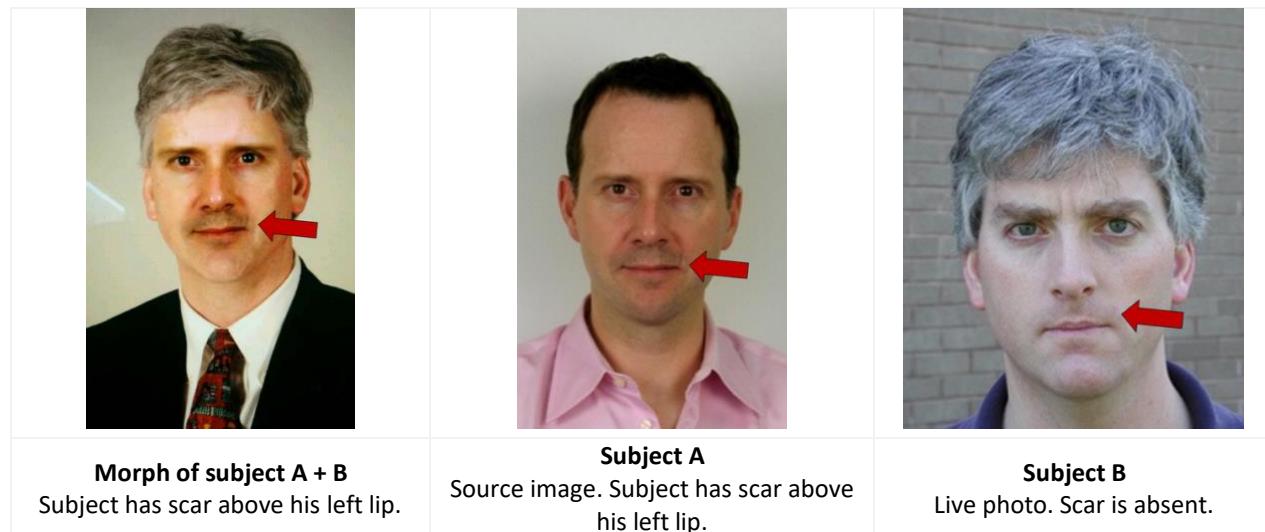


Figure 9 – Absence of features – scars

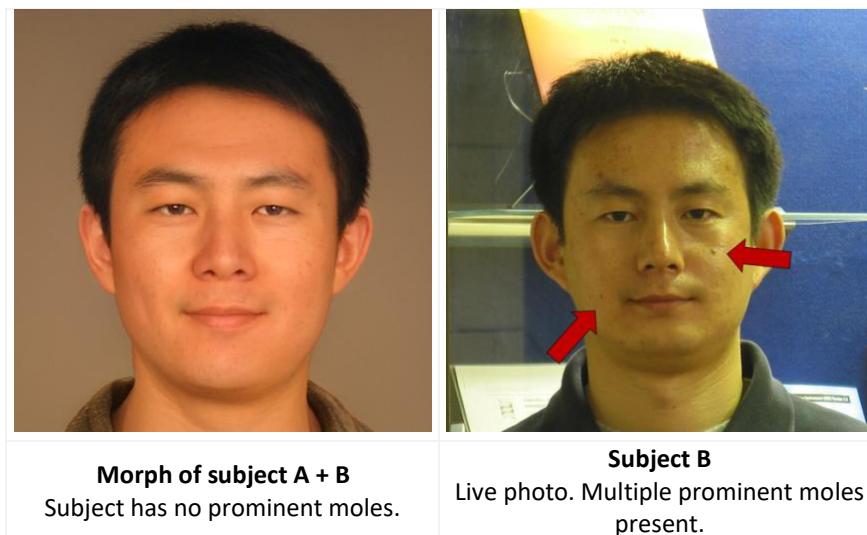
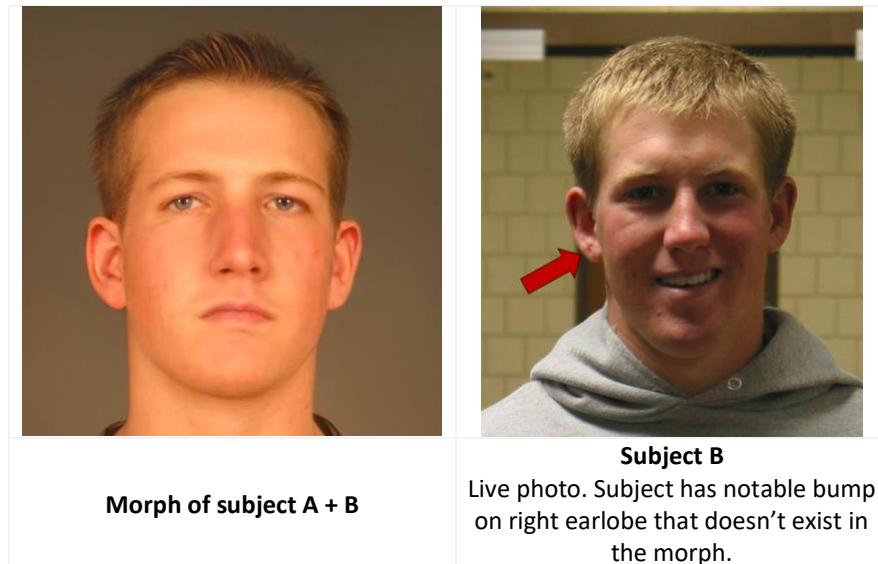


Figure 10 – Absence of features – moles



**Figure 11 – Difference of features – ears**

### 5.1.3. Inspection of image file metadata

JPEG image files are often accompanied by extensive metadata in an EXIF (Exchangeable Image File Format) header. This information can be inspected in most photo editing packages and with dedicated tools such as *jhead*, *ExifTool*, or *ffinfo*.

The EXIF header is sometimes absent but, when present, can include date, location, camera, shutter speed, and software image handling information. This data can have investigational value. For example, the capture or modification dates may not be consistent with information from the applicant or traveler; the geo-location may not be consistent with claimed residence; or an attacker may leave definitive information such as the name of a known morph-preparation tool. A high-effort attacker would be expected to edit the EXIF header to cover their tracks.

## 5.2. Use of a different biometric characteristic for identity verification

In addition to a face photograph, certain countries will collect and store fingerprint and/or iris imagery on the passport. Because fingerprint and iris imagery are often collected in person, the chances of those images being manipulated are much lower, leading to more confident identity verification outcomes. If a different biometric modality is available, agencies should consider collecting the biometric characteristic from the traveler and validate that it matches what is stored in the passport.

## 5.3. Use of face recognition

### 5.3.1. Applying 1:1 face recognition to a pair of images

Face recognition algorithms compare faces and return a measure of similarity. High scores are consistent with the pair being of the same person. Low scores are produced from photos of different people or, importantly, from poor quality photos. In morph investigations, a score between those two possibilities may be an indicator that a morph is present – see Figure 12.

An investigator will need quantitative data on what range of scores are expected from comparison of photos of the same person and different people. Such data is best produced through operational experience with that algorithm applied to photos of the operational population photographed in the normal operational environments. This is critical as different operations may encounter photos of varied quality characteristics or a non-homogenous population. Score distribution data may also be available from a test lab, or from the supplier.

If the score is reduced and looks anomalous with respect to that range, then this could be an indication that the pair includes a morph. However, the investigator must realize that a reduced score could also arise due to a number of factors, including natural ageing of the subject's face – for example when the pair of images are collected a decade or more apart – or some other change in facial appearance, such as injury or plastic surgery. Image quality deficiencies for one or both images will also lower the match score, which may arise, for example, when one image is from a passport and the other is a selfie from the user's mobile device.

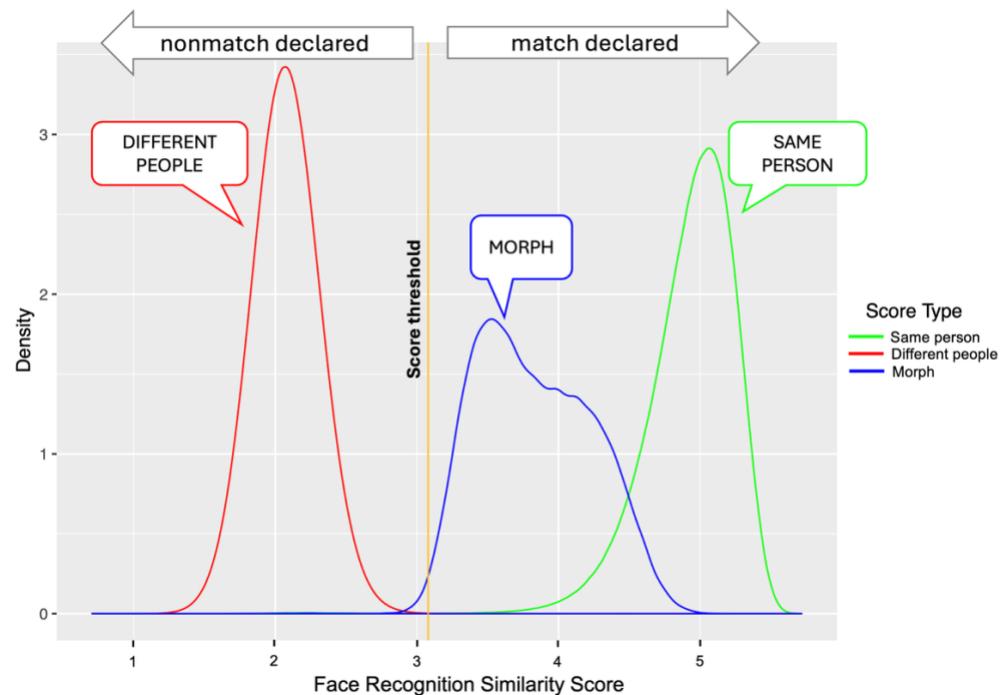


Figure 12 – Illustrative face recognition similarity score distributions

### 5.3.2. Applying 1:N face recognition to search a database

In the context of a potentially morphed image being used to either establish a new identity or to renew an existing identity credential, if the issuing organization maintains or has access to a centralized database of past applicant facial photos, there is an opportunity to search the application photo against the database with the goal of detecting morphs. Based on the outcomes reported in [22], 1:N face recognition systems may have utility in detection of morphed photos in operational pipelines, with particularly promising results under an ID renewal scenario. One potential advantage of using such a 1:N approach is that many ID issuance agencies (e.g., passport offices) will already have a 1:N face recognition system within their operational pipeline, so there is opportunity to reuse existing infrastructure in lieu of procuring a dedicated morph detection capability.

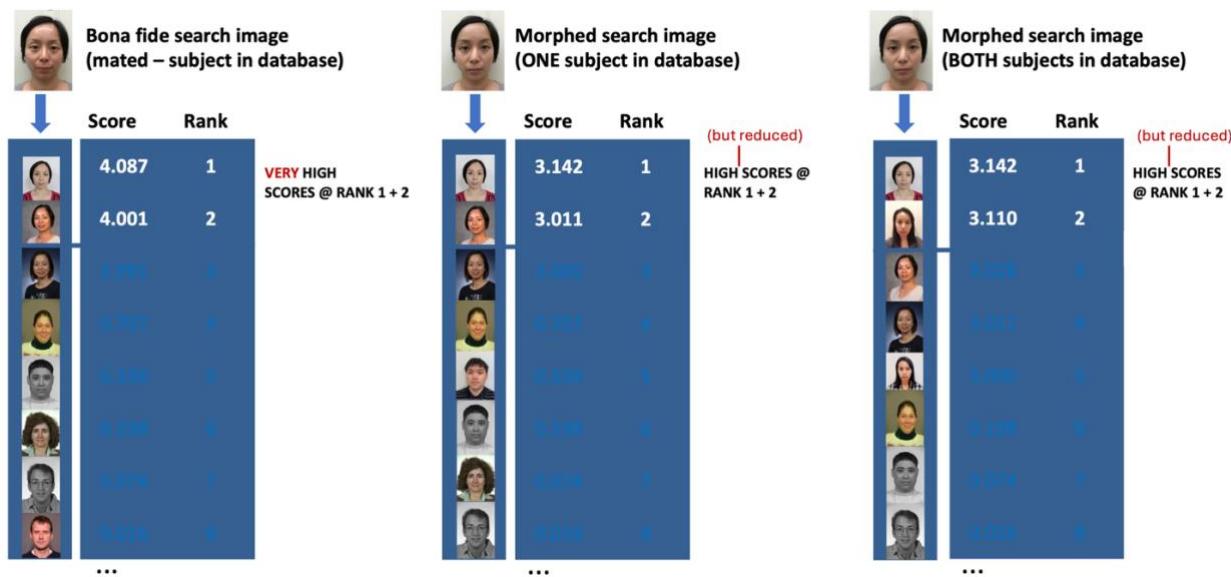


Figure 13 – Morph detection using 1:N face recognition

When a morph (subject A + subject B) or bona fide probe is searched against a gallery, if images of the subject(s) exist in the gallery, possible outcomes for what gets returned on the candidate list include:

Under a **renewal** scenario (see Figure 13)

- For a bona fide (non-morph) search, prior photos of the subject would be returned at top rank with very high similarity score(s).
- For a morph search
  - If only one of the contributing subjects exists in the gallery, prior photos of that subject is expected to be returned at top rank with high but *reduced* similarity scores.

- If both subjects exist in the gallery, any combination of subject A-only, subject-B only, or subject A and B could be returned at top rank, but in all cases, with high but *reduced* similarity scores.

Under a **new application** scenario where subject B is trying to obtain a new ID document

- For a bona fide (non-morph) search of subject B, we would expect no photos of the subject to be returned and any photos retrieved at top rank would be returned with very low similarity scores, indicating that no existing matching identity was found.
- For a morph search
  - If the other contributing subject (subject A) exists in the gallery, prior photos of subject A is expected to be returned at top rank with high by *reduced* similarity scores.

The use of such an approach would require similarity score threshold calibration for when to trigger that an image might be morph. The effectiveness of this approach will also be highly dependent on the underlying face recognition algorithm being used in the system. This is discussed in much more detail in NIST IR 8430: FATE Part 4A - Utility of 1:N Face Recognition Algorithms for Morph Detection [22].

## 6. Procedures for investigation of candidate morphs

When an organization deploys morph detection in operations, the software may at some point yield a candidate morph. An investigation, if successful, will reveal that this flagged photo is a morph or a legitimate photo i.e., a false detection.

The following two subsections suggest steps and procedures to help investigators. The first pertains to applications in which an S-MAD detector has flagged an image; the second applies to applications in which a suspect morph accompanied by a live-capture constituent photo (known-non-morph) have caused D-MAD software to flag the pair.

### 6.1. Procedures after an S-MAD detection

Given a suspect morph M, and a positive indication from an S-MAD implementation -

1. Inspect the photo for visible artifacts – see section 5.1.1.
2. If the photo was submitted digitally, inspect the EXIF header – see section 5.1.3.
3. If M is a photo being submitted for renewal of an ID document and if a database of prior applicants and photos is available.
  - a. Execute a text database query to retrieve all N prior photos ( $X_1, X_2 \dots$ ) of the claimed identity.
  - b. If a D-MAD implementation is available, compute a D-MAD score between M and all  $X_n$  priors (and look for any positive indications).

- c. If a 1:1 face matcher is available, execute a 1:1 comparison between M and all  $X_n$  priors (and look for any rejections). Additionally human examination of all photos may reveal facial morphology differences due to morphing.
- d. If a 1:N face recognition engine is available, execute a face recognition search of M against the database and see section 5.3.2 and [22].
4. Reject the subject's application (but retain all documentation and photos as evidence of potential attempted fraud) and request a new trusted photo from the applicant.
5. If the subject can be asked to go to a trusted photo capture site, schedule a session and require live trusted capture. This may be costly and inconvenient.

## 6.2. Procedures after a D-MAD detection

Given a suspected morph M, a live capture photo X, and a positive indication from a D-MAD implementation -

1. If possible, refer the traveler to a secondary inspection process.
2. Inspect suspected morph for visible artifacts - see section 5.1.1.
3. Inspect suspected morph and live capture photo for feature absences or differences – see section 5.1.2.
4. If an S-MAD implementation is available, run it on suspected morph (and look for positive indication).
5. Collect a new pristine quality “live” photo from the traveler, photo Y.
6. Using D-MAD software
  - a. Compute a D-MAD score,  $D_{MY}$ , for the pair (M, Y).
  - b. Re-compute the original D-MAD score,  $D_{MX}$ , from the pair (M, X) to avoid software version discrepancies between primary and secondary software installations.
  - c. If the D-MAD scores  $D_{MY}$  is below threshold this may indicate that M is not a morph. However, the investigator must realize that the score can be impacted by ageing – for example when the pair are collected a decade of more apart – or some other change in facial appearance, such as injury or plastic surgery.
7. Use FR software – see section 5.3.1
  - a. Compute FR similarity score,  $S_{MX}$ , between suspect morph and primary border photo X.
  - b. Compute FR similarity score,  $S_{MY}$  between suspect morph and the new pristine photo Y.
  - c. Compute FR similarity score,  $S_{XY}$ , between the two same day photos X and Y.

- d. You would expect to see  $S_{MX} < S_{MY} < S_{XY}$ . However, if  $S_{MY}$  is a lot lower than typical same-person similarity scores and the date of issuance on the ID document is much less than 10 years, this might be evidence that M is a morph. It's possible that bona fide scores from matching the same person across different photos could be lower due to ageing.

## 7. Preventing morphs from entering operational systems

One of the most effective defenses against the use of morphs in identity fraud is to prevent morphs from getting into operational systems and workflows in the first place. Methods to achieve this apply during document application and issuance.

### 7.1. Attended live enrollment

A common trusted capture process is for the live photo to be collected in-person in the presence of a trusted attendant. Live enrollment processes require in-person attendance, which may require long distance travel. One way to reduce travel is to authorize capture by certified photographers working at, for example at town halls or post offices. In such cases, staff should be trained and vetted, and photos should be transmitted directly and securely from them to the identity issuing agency.

### 7.2. Unattended trusted capture

When supervised live enrollment is not possible, there are other forms of trusted capture, but these are of a lower level of trust as they can be subject to successful attacks. These may include

- Deployment of trusted equipment that collects and transmits photos directly to the collecting agency's servers without any user intervention. An example of this is a photobooth for passport photos. The hardware should cryptographically protect the integrity and confidentiality of the data by preventing tampering or upload. Unattended photographs can only be trusted if the capture includes adequate presentation attack detection [23] – to prevent, for example, replay of a morph on a tablet display.
- Use of certified photographers
- Use of dedicated secure mobile-phone or tablet-based application equipped with defenses against, and detectors of, injection and presentation attacks.
- Others...

### 7.3. Apply morph detection

Document issuance entities should adopt a risk-based approach when determining which morph detection procedures and techniques to implement. For organizations that permit user-submitted photos as part of the application process, a practice that carries a higher risk of image manipulation, additional safeguards should be considered. These may include implementing trusted capture methods to prevent tampering, substitution, or morphing. When trusted capture is not feasible, entities should assess the risk level and consider proportionate countermeasures, such as deploying automated morph

detection tools and training staff to recognize morphing threats. Staff should also be equipped to interpret detection results and take appropriate steps to verify the applicant's true identity.

As an example, a passport agency can take a risk-based approach by layering multiple identity verification and morph detection measures based on the potential threat level. For example, an agency might implement the following checks on an application photo:

- A 1:N duplicate check of the application photo
- A 1:1 comparison of the application photo with all previously submitted passport photos of the applicant
- Single-image morph detection (S-MAD) of the application photo
- Differential morph detection (D-MAD) comparing the application photo with all previously submitted passport photos of the applicant

If all these checks consistently indicate that the photo may be a morph, the accumulated evidence raises the risk that the image has been manipulated. In such high-risk cases, the agency may choose to escalate by rejecting the application and requiring the applicant to appear in person at a trusted facility (e.g., a passport office or postal service location) for photo capture or to submit an alternate image.

## **8. Additional training**

Training for human examiners on the detection of morphs has been developed by the Norwegian ID Centre and is available to government entities. Interested government parties may apply for access to the training [here](#).

## 9. References

- [1] M. Ngan, P. Grother, K. Hanaoka and J. Kuo, "Face Analysis Technology Evaluation (FATE) Part 4: MORPH - Performance of Automated Face Morph Detection," NIST, 2025.
- [2] M. Ferrara, A. Franco and D. Maltoni, "The magic passport," in *IEEE International Joint Conference on Biometrics*, Clearwater, FL, 2014.
- [3] B. Chaudhary, P. Aghdaie, S. Soleymani and J. Dawson, "Differential Morph Face Detection using Discriminative Wavelet Sub-bands," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, Nashville, TN, 2021.
- [4] K. O'Haire, S. Soleymani, B. Chaudhary, P. Aghdaie, J. Dawson and N. Nasrabadi, "Adversarially Perturbed Wavelet-based Morphed Face Generation," in *16th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2021)*, Jodhpur, India, 2021.
- [5] "NIST Face Analysis Technology Evaluation (FATE) MORPH," [Online]. Available: [https://pages.nist.gov/frvt/html/frvt\\_morph.html](https://pages.nist.gov/frvt/html/frvt_morph.html).
- [6] "iMARS - image manipulation attack resolving solutions," [Online]. Available: <https://imars-project.eu/>.
- [7] J. Prince, "Operational Experiences with Attack Detection," in *International Face Performance Conference (IFPC) 2025*, 2025.
- [8] R. Thelen and J. Horchert, "Activists smuggle photomontage into passport," 22 September 2018. [Online]. Available: <https://www.spiegel.de/netzwelt/netzpolitik/biometrie-im-reisepass-peng-kollektiv-schmuggelt-fotomontage-in-ausweis-a-1229418.html>.
- [9] M. Torkar, "Slovenia: Morphing Cases in Slovenia," in *International Face Performance Conference (IFPC)*, 2022.
- [10] JTC1/SC37, ISO/IEC FDIS 20059 Information Technology – Methodologies to Evaluate the Resistance of Biometric Recognition Systems to Morphing Attacks, vol. 1, 2025, pp. 1-17.
- [11] "Morph Thing," [Online]. Available: [www.morphthing.com](http://www.morphthing.com).
- [12] M. Ferrara, A. Franco and D. Maltoni, "Face Demorphing," *IEEE Transactions on Information Forensics and Security*, vol. 13, no. 4, pp. 1008-1017, 2018.
- [13] M. Ferrara, A. Franco and D. Maltoni, "On the Effects of Image Alterations on Face Recognition Accuracy," in *Face Recognition Across the Electromagnetic Spectrum*, Springer, 2016.
- [14] M. Ferrara, A. Franco and D. Maltoni, "Decoupling texture blending and shape warping in face morphing," in *International Conference of the Biometrics Special Interest Group (BIOSIG)*, Darmstadt, Germany, 2019.
- [15] "Abrosoft FantaMorph," Abrosoft, [Online]. Available: <https://www.fantamorph.com/>.
- [16] H. Zhang, S. Venkatesh, R. Ramachandra, K. Raja, N. Damer and C. Busch, "MIPGAN—Generating Strong and High Quality Morphing Attacks Using Identity Prior Driven GAN," *IEEE Transactions on Biometrics, Behavior, and Identity Science*, vol. 3, no. 3, pp. 365-383, 2021.
- [17] K. Preechakul, N. Chathee, S. Wizadwongs and S. Suwajanakorn, "Diffusion Autoencoders: Toward a Meaningful and Decodable Representation," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.

- [18] Z. Blasingame and C. Liu, "Greedy-dim: Greedy algorithms for unreasonably effective face morphs," in *IEEE International Joint Conference on Biometrics (IJCB)*, Buffalo, NY, 2024.
- [19] S. Venkatesh, R. Ramachandra, K. Raja and C. Busch, "Face Morphing Attack Generation and Detection: A Comprehensive Survey," *IEEE Transactions on Technology and Society*, vol. 2, no. 3, pp. 128-145, 2021.
- [20] "Bologna Online Evaluation Platform (BOEP)," [Online]. Available: <https://biolab.csr.unibo.it/fvcongoing/UI/Form/BOEP.aspx>.
- [21] M. Ngan, "NIST Face Morph Detection Evaluation (FATE MORPH)," in *International Face and Fingerprint Performance Conference (IFPC) 2025*, 2025.
- [22] M. Ngan, P. Grother and K. Hanaoka, "Face Analysis Technology Evaluation (FATE) Part 4A: MORPH - Utility of 1:N Face Recognition Algorithms for Morph Detection," NISTIR 8430, 2022.
- [23] JTC1/SC37, ISO/IEC 30107-1: International Organization for Standardization: Information Technology – Biometric presentation attack detection – Part 1: Framework, vol. 2, 2023, pp. 1-11.
- [24] R. Thelen and J. Horchert, "Activists smuggle photomontage into passport," 22 September 2018. [Online]. Available: <https://www.spiegel.de/netzwelt/netzpolitik/biometrie-im-reisepass-peng-kollektiv-schmuggelt-fotomontage-in-ausweis-a-1229418.html>.