# Matty McMattface: Security implications, mitigations & testing strategies for biometric facial recognition systems

Prepared by:
Matt Lewis, technical research director

# Table of contents

# 1. Matty McMattface – Security implications, mitigations & testing strategies for biometric facial recognition systems

## 1.1 Summary, business context & target audience

Biometric facial recognition is becoming an increasingly popular mechanism for authenticating users in online and mobile environments. In addition, it is continually being adopted for physical access control, whether at border controls such as airports or within secure facilities to enforce strict access control (and/or time and attendance tracking) to sensitive rooms or areas. In this paper, we look at the current state of play with facial recognition and its potential applications as a secure authentication mechanism. We examine peripheral security issues around the use of facial recognition that need to be considered and understood by those implementing and using such technology. Specifically, from a security perspective, we look at the potential attacks against facial recognition systems, with focus on the spoofability of this technology and its different implementations. We provide an overview of research performed around the spoofability of two modern facial recognition applications for logical access control - Windows Hello [1] facial recognition on the Surface Pro 4 tablet and how we worked with the Microsoft Security Response Centre (MSRC) [2] in helping it implement improvements to Hello to mitigate issues found around spoofing. We also show the same research against Trusted Face (Android's predecessor to Face Unlock [3]) which allows Android phone or tablet users to unlock their devices by simply looking at them.

Alongside the documented attacks we provide analysis of the complexities involved in performing them such that we can understand both strengths and weaknesses of the technology in order to accurately assess risk when considering application of facial recognition to different authentication scenarios.

This paper is aimed at IT practitioners tasked with implementing, testing, or looking to use facial recognition biometrics as an authentication mechanism in physical and/or logical systems or applications. The paper should also be useful to anyone interested in learning more about facial recognition in general, with specific focus on the security merits and limitations of facial recognition systems.

## 1.2 Overview

Facial recognition is a physical biometric, meaning it is concerned with physical, intrinsic properties of faces which are fairly consistently measurable. The technology has been evolving for some time, with initial research performed as far back as the 1960s [4]. As camera and sensing technology has improved, miniaturised and reduced in cost in recent years, suddenly the potential applications of facial recognition are huge due to the necessary technology (and software) being readily available in modern smartphones, tablets and laptops. Some banks are already trialing facial recognition as part of customer authentication [5] and fraud detection functions, while various border controls are already using it and are looking at how the technology can help transform the way in which we enter and exit those borders [6]. From a vendor perspective there are hundreds, if not thousands, of small to large enterprises in this space, be they startups with new approaches to, or innovation in biometric technology [7], or large tech companies such as Google (Face Unlock/Trusted Face in Android [8]) and Microsoft (Hello face authentication in Windows 10) providing device and/or enterprise-grade authentication through facial recognition [9].

Naturally, as with any advances in and adoption of new technology, concerns exist around the security of that technology. For facial recognition several questions arise, non-exhaustively including:

- In which applications are facial biometrics a good choice for authentication?
- When should facial biometrics not be used?
- Can facial biometrics be easily spoofed or faked?
- Do environmental factors such as lighting affect the performance of facial biometric systems?
- How do we test the security of facial biometric systems and validate vendor claims?

This whitepaper aims to address the points above yet is by no means exhaustive; instead it serves as a baseline of understanding for those seeking to implement facial recognition and/or those seeking to test the security of facial recognition systems. The methodology is important as it allows for consistency during feasibility study, security and risk assessment and risk mitigation for facial recognition system deployments.

## 1.3 Preliminaries

Before we begin we set out some preliminaries. Biometrics in general have been met with scepticism over the past few decades while inaccuracies in the media and some vendor claims have contributed to misinformation around this technology. Here we establish some facts around biometrics and facial recognition in order to inform and educate the reader to an accurate baseline of understanding.

### 1.3.1 How facial recognition works

There are two types of biometric – physical and behavioral. Facial recognition is considered a physical biometric since it works by measuring physical, intrinsic properties of the face. The typical modus operandi of a facial recognition system is to first find the face within an image or video stream, then identify key points on the face such as the minima and maxima of features such as the mouth, eyes, nose. Finally, measurements between key points are taken to create a geometric model of a user's face which should be fairly unique to the individual.

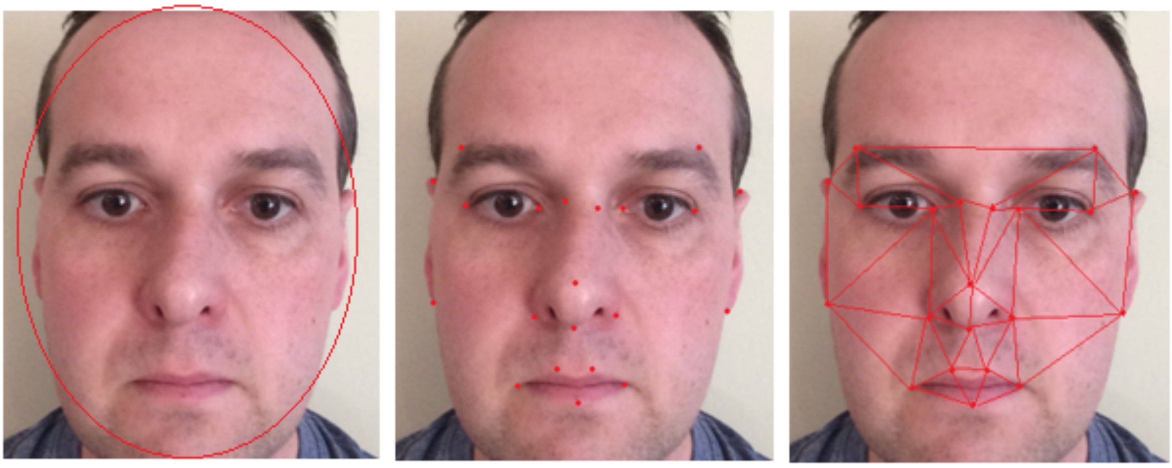

*Figure 1 - First the face is found  then specific points on the face located. Finally  distances and geometry between points are calculated*

There is no standard (or mandate) for how facial recognition systems should be implemented algorithmically, in terms of facial point localisation, geometry definition and template construction. As such, various implementations have been developed and deployed over time with varying levels of success around performance and accuracy. For facial recognition, the ideal is for systems to focus only on measuring aspects of the face that are least variable. For example, some older systems used to capture facial skin tone as an identifier, however, skin tone is an example of something which can change dramatically due to tanning/paling and under different lighting conditions. Similarly, the mouth is a variable facial aspect, changing depending on a myriad of conditions such as talking, smiling, pouting etc. while hair is also variable due to different styling choices, growth and/or cutting. As such, the most effective facial recognition systems tend to avoid capturing and comparing colour-based aspects of the face and might commonly avoid capturing data points around the mouth. For colouring and lighting reasons there is a growing trend for facial recognition systems to capture facial images under infrared lighting conditions.

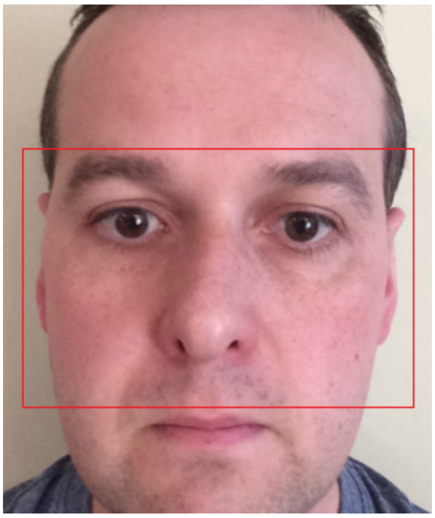

*Figure 2 - To minimise variance  facial recognition systems usually focus on capturing data around from above the mouth and not beyond the eyebrows*

Note that facial recognition is a separate biometric mode to iris recognition. While it is possible to combine both biometric modes in a multi-modal system, facial recognition in isolation does not extract iris patterns and perform iris recognition as part of the overall facial recognition matching process. Some facial recognition algorithms may take measurements of the eyes and distances between them etc. but this does not include any analysis of a user's iris pattern(s).

### 1.3.2 Biometric matching is a probability calculation

Biometric matching algorithms operate on a statistical model, whereby the system is pre-configured with a minimum threshold which must be satisfied by the presented sample before a positive match is made. This is done to cater for variances in presentation of biometrics each time and also to account for other factors that can disrupt sample capture; for example, people's faces will change due to a number of factors (such as alertness, physical condition, ageing, mood, weight loss/gain etc.) while environmental factors such as different lighting conditions can impact on the quality of image(s) captured by a facial recognition system (see later).

From a security perspective, the main threshold to consider is what's called the False Accept Rate (FAR). This is the probability that a biometric system incorrectly matches the input pattern of a person to that of a different person. In facial recognition for example, two or more people may look similar and have sufficiently similar measurements in their facial features (e.g. identical twins) and might therefore be able to successfully authenticate as each other, though depending on the sophistication of the facial recognition system, it may be able to measure very discrete differences allowing for differentiation of even identical (to the naked eye) twins.

The False Reject Rate (FRR) is therefore the trade-off of the FAR. This is the probability that the system fails to match a valid biometric sample presented to a system. In real-world applications, the FRR will likely frustrate users as it will mean that the system will sometimes deny legitimate users access to their information or resources.

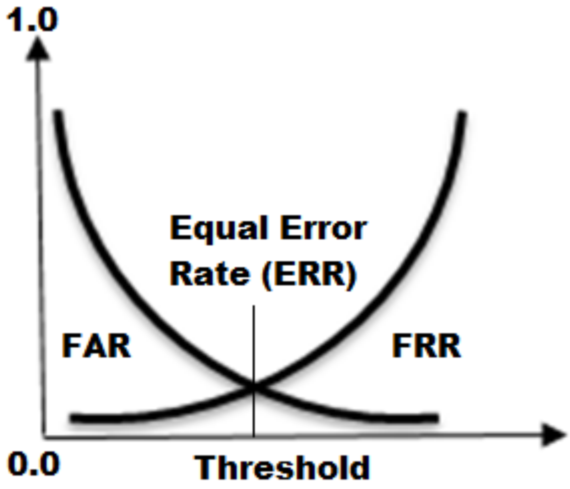

*Figure 3 - FAR  FRR and their intersection which is the Equal Error Rate (ERR)*

The FAR and FRR values are typically configurable within biometric systems and the choice of these values needs to be carefully considered depending on the application, and the nature of the data that application seeks to protect. The two rates are inversely proportional to each other meaning that an increase in one error rate decreases the other. A good rule of thumb for FAR configuration [10] and different levels of security is:

- 1 user in 100 users is falsely accepted – Basic Security
- 1 user in 10,000 users is falsely accepted – Medium Security
- 1 user in 1,000,000 users is falsely accepted – High Security

Depending on the system/application, the threshold may be adjustable (e.g. by an authorised system administrator), allowing for fine tuning of a system over time and to the features of its specific user population. Other systems may have the matching threshold pre-configured or hard-coded by the system developers or manufacturers and thus cannot be changed post deployment.

### 1.3.3 Biometric data is personal information

Under many jurisdictions biometrics are considered personal information. In the UK, the Data Protection Act 1998 [11] and Protection of Freedoms Act 2012 [12] present principles for protection of biometric data, while General Data Protection Regulation (GDPR), which comes into force in May 2018, introduces a new category of data that is specific to biometrics [13]. Any biometric data captured and stored within a biometric system (such as face images and templates) must be secured according to best practice. Where biometric systems or applications employ the use of end user equipment (e.g. smartphones) to capture (and perhaps store) biometrics, this introduces an interesting question around data controller and ownership. As such, legal counsel should be engaged early on when considering implementation of facial recognition biometrics to understand exactly what data is captured, when, how, and where/how it is stored and secured.

### 1.3.4 The need for fallback processes

It is often the case that users won't be able to use biometric systems – this includes registered users and/or new people at the point of biometric system enrolment.

If a registered user in a facial recognition system was to at some point in the future suffer unfortunate and significant disfigurement to their facial features in ways that would stop acquisition algorithms from being able to identify and extract facial features, then the user may not be able to successfully authenticate to the system. In biometrics this is referred to as a Failure to Acquire (FTA) condition as the system cannot acquire a valid sample from the user. Similarly, prospective users may be able to satisfy image acquisition, but due to potential issues such as facial disfigurement not satisfying the quality control aspects of the enrolment process, may not be able to enrol, resulting in what's known as a Failure to Enrol (FTE) condition.

Whether FTA or FTE, what this means is that biometric systems will need to cater for these conditions – i.e. one cannot preclude individuals from using IT systems or from passing through border control simply because the technology cannot extract the requisite features of those individuals. As such, biometric systems need to implement fallback or secondary authentication processes that do not rely on biometrics but allow for authentication of FTA-prone or FTE-prone individuals.

The potential issue with this is that if the fallback function is available to all users, then that function may become the preferred or de-facto mode of use, thereby rendering the presence and implementation of the biometric component potentially useless. The biometric function of authentication systems needs to be sufficiently convenient and easy-to-use to avoid fallback functions being used in preference. With facial recognition, a fallback function might include other authentication factors such as PINs, passwords or different biometrics such as fingerprint or voice recognition.

> "The biometric function of authentication systems needs to be sufficiently convenient and easy-to-use to avoid fallback functions being used in preference."

### 1.3.5 Verification vs. identification

In biometrics, authentication can be performed in two different modes:

- Identification – this is a one-to-many matching process, whereby a user makes no claim of identity. In biometric identification, users present their biometrics to a system and are essentially asking the question ‘tell me who I am’. The system then performs a one-to-many search across all enrolled templates and returns either a positive match or notification of a non-match if the user’s biometrics are not recognised from the database of registered users. Due to the statistical matching process of biometrics, some systems may return an entire match list, including one or more individuals who are positive matches to the sample presented. Returning a match list may be preferable with facial recognition systems that operate in identification mode and on large enrolment databases (e.g. millions of users) – for example, in one study [14], National Institute of Standards and Technology (NIST) found an error rate increase of 1.2 times was experienced in research where the number of facial images in the study increased by a factor of ten – from 160,000 to 1.6 million. In [15], experimental research results demonstrated that both commercial and non-trainable algorithms consistently had lower accuracies on the same demographic cohorts

- Verification – this is a one-to-one matching process, whereby the user makes a claim of identity. Here, a user tells a biometric system that they have an identity (e.g. through a username or some other unique identifier). The user then presents their biometric (face) and the system performs the matching process with the presented sample and a stored template belonging to that user. If there is a match then the individual is verified. Failure to match means that the user is not verified. This is similar to border control applications where the image of a traveller is scanned from their passport (or read from an embedded chip within the passport) and compared with the facial image taken by the border camera at that point in time. Note that specifically in border control systems, at the same time of this one-to-one verification process, a secondary background one-to-many identification process may occur whereby the traveller’s facial image is checked against a watch list of individuals wanted for outstanding warrants for example.

### 1.3.6 Negative identification

A powerful feature of biometrics which no other technology can offer is negative identification. This is where a system can check if a new user is already enrolled in the system. This function is well-suited to applications that aim to detect and thwart fraud. In a facial recognition system for example, when users enrol for the first time, the system could perform a negative ID check to see if that user is already enrolled under a different identity. No other technology can achieve this type of checking without relying on trust of the information provided by users; the application of negative identification can be rewarding in the combat against fraud in areas such as banking and government systems that process applications for state benefits or identity documents such as passports and driving licenses.

## 1.4 Considerations for facial recognition systems

In this section, we document key areas for consideration for those seeking to use biometric facial recognition within physical or logical access control systems and applications.

### 1.4.1 Faces are not secret

Unlike passwords or PINs, which typically remain known only to ourselves, faces (and biometrics in general) are vulnerable to copying through a variety of methods. For facial recognition, taking a picture of someone's face can be done easily, either cooperatively or covertly – such images might later be used in spoofing attacks against the victim's presence within facial recognition systems. In the 21st century our faces are particularly prone to image capture or exposure owing the vast amount of imagery that we put online about ourselves through social media for example, making it substantially easier for motivated attackers to obtain facial images of victims for later spoofing attacks.

### 1.4.2 Facial privacy & surveillance

Preservation of the privacy of human faces may need to be considered in some facial recognition systems. For example, some religions may have restrictions on the level of exposure permitted of faces – this may either mean that some individuals may object to using such technology on religious grounds, or this may at least mean there is a requirement to provide private spaces for individuals to be able to use facial recognition systems in ways that do not expose their faces in public places. If moral or religious objections to the use of facial recognition are particularly strong or deemed lawful within jurisdictions, then this may require fallback processes to be implemented as mentioned in section 0 above.

Note that facial recognition is gaining rapid adoption in surveillance-based applications, from border control (e.g. coupled with CCTV at airports) to marketing applications that can (via CCTV) track

> **"Unlike passwords or PINs, which typically remain known only to ourselves, faces (and biometrics in general) are vulnerable to copying through a variety of methods."**

customers as they walk around stores and perhaps even detect their gender and mood (this information might be used in targeted advertising on electronic displays in store for example) [16]. Unless special exemptions are in place for surveillance-based facial recognitions systems (e.g. matters of national security), then system integrators will need to be very careful around clandestine use of such systems. Explicit opt-in permissions will need to be sought from users that make it clear what the system is, how it works and how data is processed, otherwise there is risk of breach of several data protection laws and regulations. Note also that users will need to be provided with clear opt-out choices for such systems.

### 1.4.3 The need for attention grabbing

With surveillance-based systems in particular, there may be a need to grab people's attention in ways that ensure that their faces look directly into cameras, such that facial images can be detected and acquired. Examples here include crowd surveillance where perhaps hundreds of faces need to be captured and processed every minute. Techniques to be employed here will be application-specific, though common methods involve use of visual aids or stimuli such as animated advertisements just above or below the camera, use of LEDs/flashing imagery [17] around the camera or even sound projections in specific locations. Without such visual cues, such surveillance systems are likely to exhibit poor FTA rates.

### 1.4.4 Monozygotic (identical) twins

To the human eye, identical twins may indeed appear to be the same person. This may also be true for some (typically older) facial recognition systems that may not be able to distinguish between identical twins. This does present the potential for an actual 'evil twin' scenario, whereby one twin may be able to successfully authenticate as the other on systems where only one of them is enrolled. While to the human eye the distinction between identical twins may seem impossible, to facial recognition systems that produce precise measurement models of faces, the task of differentiating between identical twins may be trivial since the underlying geometric models of identical twin faces may be substantially different – e.g. a millimetre difference in distance between eyes, nose, mouth etc. would not necessarily be noticeable by the human eye but would create quite different geometric models from input images of sufficient high quality and resolution [18].

### 1.4.5 Environmental factors

The operating environment can have a major impact on the performance of a biometric system. Where users present their face to a border control camera, smartphone camera or laptop webcam for example there are a number of things to consider:

- Different lighting conditions – variances in lighting conditions might have an impact on the samples captured during enrolment and authentication. Note that systems operating under infrared lighting will be less susceptible to impact here (such as Windows Hello). Where infrared is not used for image capture, some level of guidance will typically need to be provided to end users; such as informing them not to use the system in dark or extremely lit (perhaps with sun glare) conditions. In end user device facial recognition applications, end users should be given guidance on optimum lighting conditions under which to use the facial recognition function. Sophisticated systems may be able to detect sub-optimal conditions and result in a FTE/FTA error accordingly.

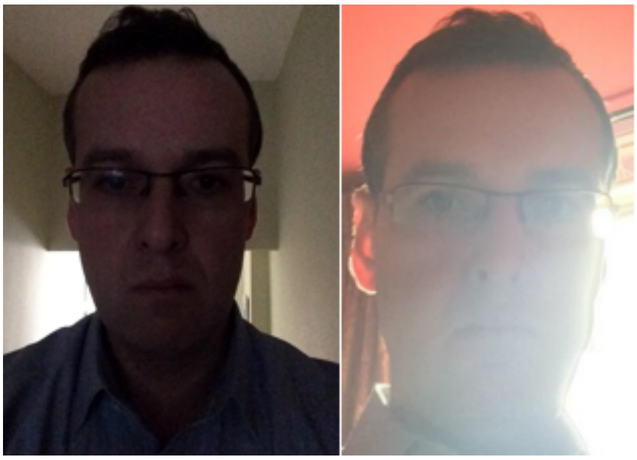

*Figure 4 - Examples of dark and glare lighting conditions that might affect system performance*

- Faulty equipment – if users use their own devices (e.g. smartphones) to present their facial biometrics then broken/faulty components such as cameras could have an impact on the quality of biometric sample captured by the device.

- Dust or dirt on cameras – this could have an impact on the quality of facial image captured and could affect the FAR/FRR or FTE/FTA rates of the system. For end user device applications of facial recognition, system integrators may want to consider periodic issuance of guidance to end users on how to clean their devices and cameras to maintain optimum performance. For border control systems, operators should be given guidance on safe, regular cleaning of the camera and sensing equipment in order to maintain optimum performance.

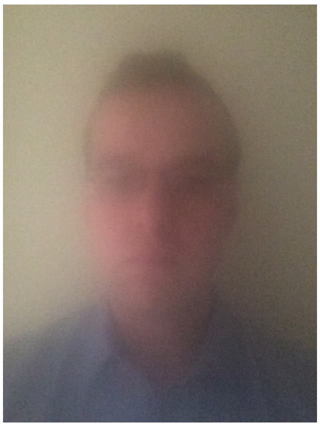

*Figure 4 – Image from smartphone with hand cream smeared (perhaps accidentally) over the camera lens*

- Electromagnetic interference – either from other components within smartphones/laptops or ambient electronic devices, there could be some negative impact on a capturing camera's performance due to electromagnetic interference.

- Different equipment operating with matching algorithms – some facial recognition developers may produce the facial matching algorithms only, leaving the hardware and sensing technology choices to end user device manufacturers who choose to embed the algorithms within their own technologies. This means that the algorithms may be used with a myriad of different camera and sensing technologies across different devices – as such, there may be a big difference in the performance of what is essentially the same backend biometric system due to the differences in quality/composition of different camera/sensing technology used by different end user device manufacturers. In a crude example, one might expect different performance of the same facial recognition algorithms deployed in smartphone devices that implement a two megapixel and eight megapixel camera respectively [19].

- Facial presentation – while guidance may be given to end users on optimal facial presentation to cameras, variance will arise in the angles and distances between face and camera. If presentation guidance is not precise enough and the capturing process is too lenient on angle/distance from camera, then this will likely have a big impact on the performance of the facial recognition system.

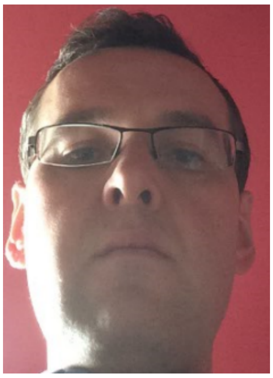

*Figure 6 - User looking down into smartphone camera - not an optimal presentation angle*

In end user device applications of facial recognition, effects on FRR could result in a denial of service for legitimate users, leading to user frustration when not being able to access their data or resources. In physical access control systems such as automated border control at airports, a high FRR could increase queuing time, traveller frustration and cause overall disruption to effective airport operation. For these reasons, the environmental factors above should be considered and addressed from design through to ongoing maintenance phase of the underlying biometric system.

### 1.4.6 Biometric ageing in facial recognition

As we age as humans, our physical features change. Our faces will change (albeit subtly) over time and will often become weaker. This could affect the performance of biometric systems and as such biometric systems should take new samples/re-enrol users every so often in order to maintain templates that are indicative of a user's current state.

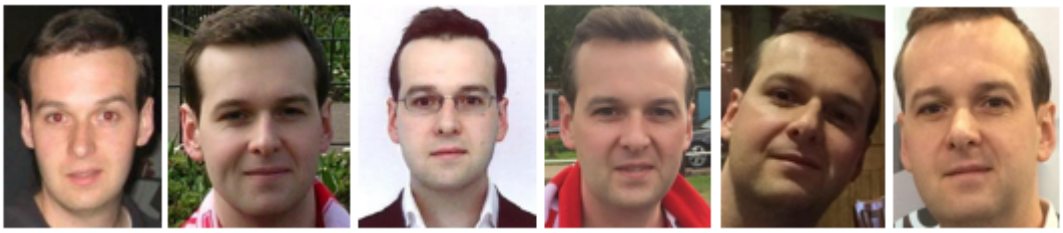

*Figure 5 - The ageing face of Matty McMattface – 2004 - 2017*

As use of facial recognition continues to be used over time, further research across large, snapshotted datasets of faces is likely to reveal interesting insights into optimum parameters and technology implementations/strategies for managing facial ageing [20].

### 1.4.7 Minimum age

Related to ageing is consideration of minimum age for use of facial recognition systems. From infancy up to adolescence the faces of individuals can change radically within short amounts of time as the individuals grow and develop. Various studies have been performed in this area which tend to conclude a minimum age of at least 12 years old before facial recognition becomes a viable authentication mechanism for those users [21].

### 1.4.8 Effects of weight loss or gain on facial recognition systems

Similar to the effects of ageing, weight loss or weight gain might have a noticeable effect on a person's face, potentially increasing the FAR for a person who loses or gains weight after initial enrolment. The same considerations around periodic enrolment should be made by system designers and integrators – to cater for changes to the face due to changes in weight, regular updating of the enrolled template may be beneficial to the overall performance of the system. Studies have been performed that conclude negative effects on system accuracy due to changes in weight, however, many of these studies have been performed as a precursor to further research in how to develop systems that can cater for changes in weight without negative effect on the system's performance [22].

### 1.4.9 Effects of cosmetic surgery (rhinoplasty) on facial recognition systems

Another way that faces might change subtly or even radically is through cosmetic surgery, perhaps either performed as a result of direct medical need or simply from an aesthetics need [23], [24]. From an attack/bypass perspective there are two interesting angles in this domain:

- A registered user wanting to avoid detection (achieving a new identity) by cosmetically changing their face

- An individual seeking to masquerade as a registered user by cosmetically altering their face to look like the victim

The latter point is likely to be most difficult in terms of complexity and with a very low (to potentially impossible) success rate since the measurements, as opposed to visible appearance would need to match those of the victim. The former point is more realistic and likely – those wanting to avoid detection on watch lists or to achieve new identities might be able to do so through cosmetic changes to their faces.

### 1.4.10 The need for periodic re-enrolment

Due to the number of different ways in which faces might change over time (naturally or forcibly), there will be a need for facial recognition systems to regularly re-enrol people to cater for these changes. Some systems may choose to update templates upon each new successful authentication, rendering the template as current as possible and thus as close as possible to the current state of the user's face. Other systems may perform less frequent template updates, or may not update at all – in this latter case, the system's FRR is likely to increase over time due to the enrolled user's inevitable ageing and changes in features. Consideration of periodic enrolment is therefore quite key to the consistent performance of facial recognition systems.

### 1.4.11 Achieving liveness checking in facial recognition systems

Liveness checking is a requisite for most biometric systems, particularly when operating in unsupervised environments such as on end user devices. For facial recognition, liveness checking techniques will seek to determine if the face presented to a camera or sensor is real and alive. There may be a number of methods available to achieve this depending on the technology available to the underlying system. Common methods include:

- Challenge-response gestures – the system may ask users to move their heads or make different facial gestures while looking at the camera, such as smiling, pouting, raising eyebrows etc. In primitive 2D systems this technique may not be able to distinguish between a video feed presented to the camera and a real person in front of the camera, so this method does not necessarily guarantee liveness. Another issue with requiring facial gestures is the impact on self-conscious users and the general practicality of the requested gestures. Many users might feel silly having to make gestures when using systems in public places which may affect the performance and or user acceptance of the system.

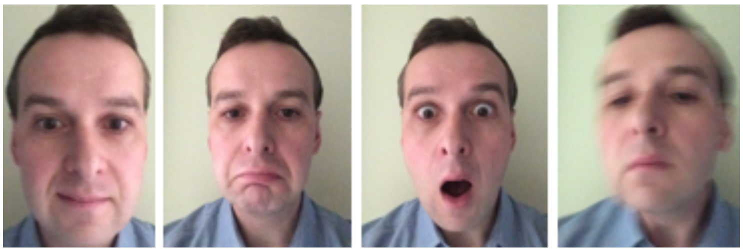

*Figure 6 - Requesting users to move/make gestures might help detect liveness but can look a bit silly*

- 3D and depth-sensing technologies – some systems may employ multiple cameras for 3D imaging and depth-sensing technology to identify depth and contours. These can be used to at least determine that a face presented to a system is 3D and not a 2D image. Such technologies may not be able to distinguish between real faces and 3D mask equivalents.

- Reflective techniques – use of different lighting projected from the sensing equipment may be done in order to assess the reflective properties of the image presented to the camera – the resulting data may provide indicators of liveness/fakeness as different materials such as glass, paper, plastic, latex etc. may present different reflectivity which is detectable within imagery.

- Thermal imaging – thermal imaging equipment might be used to identify heat patterns across faces and thus determine whether or not those patterns are indicative of a live face. This might help detect use of hard-material masks for example.

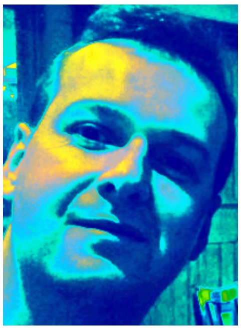

*Figure 7 - Thermal imaging to detect heat and liveness might defend against image/mask spoof attacks*

## 1.5 Attacks against facial recognition systems

This section examines some of the potential attacks against facial recognition systems. The aim of this section is to present a high-level methodology for testing and understanding the potential attacks with a view on the level of complexity involved in mounting each attack. Suggested mitigations are also presented for each attack.

We consider attacks against facial recognition systems in two scenarios - at enrolment and authentication, since these two scenarios are subtly different. Note that in this section we do not focus too much on attacks against underlying technology – e.g. attacks against people's smartphones in order to subvert facial recognition applications in some way – we are more concerned with attacks against the biometric aspect (acquisition and processing functions) more generally.

| Attack | Attack Outcome | Resources Required | Level of Complexity | Mitigation |
|---|---|---|---|---|
| Enrolment or presentation of an image, mask or non-face image.<br>For systems that allow 2D image enrolment, it may be possible to enrol an image or fake artefact of a face. | This attack might be performed either with the intent of providing an image/artefact to somebody else at a later date (authentication factor reuse), or for masquerading as someone at enrolment in an identity-theft scenario. | An attacker would need good quality imagery of a target/victim, taken with photographic equipment either with the target's consent, or covertly without consent, or perhaps imagery of the victim can be found in online resources such as social media.<br>If making a 3D mask from a 2D image (see later) of a target then 3D printing and software capabilities will be required. | The complexity involved in obtaining a good quality image of a target is relatively low owing to potential online imagery of a target or modern photographic hardware and lenses allowing for good quality image capture at long ranges.<br>The complexity with 3D printing continues to reduce as consumer-based 3D devices continue to appear on the market and the cost of these devices decreases. | One or more liveness checking mechanisms can help mitigate these types of attack. 3D systems may help mitigate these issues – perhaps not with mask attacks but at least attacks using 2D printed images of victims.<br>For physical access control, human supervision at enrolment can be very effective in reducing the potential for attack in this area. |
| Disguises used to enrol multiple identities or in attempts at masquerading as a victim for identity theft. | If the motive is identity theft, the outcome might be the same individual being able to enrol as two or more distinct identities through use of disguise methods. | Different options are available for disguising/ modifying facial appearance. Examples include prosthetics (e.g. fake nose), accessories (such as | If the motive is identity theft, then the complexity of attack may be low since there are infinitely many ways in which someone might disguise themselves. | While disguises may alter the general appearance, measurements may remain similar to the original face and so the actual biometric matching algorithm provides the best defence here, i.e. the ability to detect |

| | | | | |
|---|---|---|---|---|
| | | glasses, piercings, etc.), facial hair and makeup. | Disguising oneself to look like somebody else would be difficult since the measurements of the disguised face would need to match those of the victim, more so than the general appearance/likeness. | similarities and determine through negative identification if a user has previously enrolled.<br>Systems that operate under infrared illumination would render makeup and facial hair disguises fairly useless, while techniques such as thermal imaging might easily detect prosthetics and other inorganic materials. |
| Composite faces - Enrolment or presentation of 2D imagery or 3D masks created from composites of two or more faces. | The aim of this attack would be creating a template that can match two or more distinct faces (see later). Perhaps one of the composite images is that of an unbeknownst victim.<br>At authentication, a composite image might be used to increase the chances of a false match should the system be operating in identification mode. | Few resources would be required – good quality and similarly-posed images of two or more individuals and morphing/editing software to create the composite face. For 2D systems lacking liveness checking, a simple printout of these composite images might be sufficient to present at enrolment, or at further cost, a 3D printed mask of the composite might be created. | The complexity involved in obtaining a good quality images is low owing to potential online imagery of a target or modern photographic hardware and lenses allowing for good quality image capture at long ranges.<br>The complexity with 3D printing continues to reduce as consumer-based 3D device continue to appear on the market and the cost of these devices decreases. | One or more liveness checking mechanisms can help mitigate these types of attack. 3D systems may help mitigate these issues – perhaps not with mask attacks but at least attacks using 2D printed images of victims.<br>For physical access control systems, human supervision at enrolment can be very effective in reducing the potential for attack in this area.<br>For systems working in identification mode, where they might return a match list of different identities that are sufficiently similar, this might be an indication of a |

| | | | | |
|---|---|---|---|---|
| | | | | composite face enrolment. |
| Cosmetic surgery. | Extremely motivated attackers may undergo cosmetic surgery to avoid watch list detection and to assume new identities at enrolment in new systems, or to attempt to masquerade as a victim. | This would require access to skilled plastic surgeons. | The complexity involved is high, as is the likely cost and depending on the level of surgical alteration, with no guarantees of success for either evasion or masquerade. | Cosmetic surgery is typically a non-reversible procedure and so this in itself might serve as a strong deterrent to all but the truly motivated or desperate attackers.<br>If cosmetic surgery is not too drastic then the recognition algorithm may still correctly identify the individual if previously enrolled. |
| Weight gain or loss. | Extremely motivated attackers may seek to put on significant weight or lose significant weight to change their facial features in ways that might avoid watch list detection and to assume new identities at enrolment in new systems. | The main resource required for either weight gain or loss is time in order to over or under consume and/or exercise/not exercise accordingly. | There's no real complexity other than a likely strong will needed by an attacker in over or under eating for significant periods of time in order to change their physical appearance. This approach while possibly unhealthy is at least safer than the cosmetic surgery approach (and is reversible). | A strong will is typically needed to under or over consume for long periods of time thus this in itself might be a deterrent to all but the truly motivated or desperate attackers.<br>The underlying recognition system may still be able to detect the attacker if previously enrolled, depending on the sophistication of the system and the techniques it employees to cater for changes in features. |
| Eyeglass frame attack<br>This attack was researched and performed by [25]. | This attack involves using a Machine Learning approach to generate eyeglasses with patterns that can either allow for masquerade as a | To maximise effectiveness, this attack requires access to the source code or good understanding of the matching | The complexity is medium to high, owing to the need for a good understanding of machine learning and neural networks, in | A simple request for glasses wearers to remove them at enrolment might mitigate this issue, in addition to human supervision asking for removal at |

| | | | | |
|---|---|---|---|---|
| | victim or for evasion of recognition. | algorithm's matching/ classification process. It also requires good quality images of a target/victim and/or an attacker seeking to evade detection. | addition to the ability to be able to print specific patterns onto eyeglasses. | physical access control systems.<br>For unsupervised systems, techniques for detecting the presence of glasses and not proceeding with enrolment until they are verified as removed might be considered. |
| Presentation of imagery or artefacts that exploit coding errors in the image acquisition or processing functions in ways that might result in code execution. | An attacker may be able to craft an image or facial artefact in ways that when captured, then encoded/processed exploits software flaws in the running process resulting in code execution or hijacking of the running process.<br>In a crude example, suppose a mask is presented to a camera with a ridiculously large nose – depending on the image processing and template generation function, this could conceivably cause some sort of buffer overrun condition should the measurements be way outside expected (and possibly hardcoded) norms. | This would require the ability to manipulate input imagery and also access to the known target process to test different inputs (e.g. through fuzzing) and or code review of the image acquisition functions for potential processing flaws. | This would be very difficult to achieve owing to the variation in image capture and resulting finite digital pattern that is created. Even if a remotely exploitable vulnerability was found, it would be unlikely that the same attack payload could be generated 100 per cent the same due to variances in lighting and face presentation to the camera. | A Secure Development Lifecycle (SDLC) for all facial recognition software functions (image capture, processing, quality control, feature extraction and template generation)<br>Additional defence-in-depth and OS controls around the running processes of the software (ASLR, DEP etc.) would help mitigate exploitation of software flaws. |
| Impersonation/ lookalikes. | Attackers may naturally alter their facial expressions when facing a | For lookalike impersonation, an attacker would need to | The complexity involved in altering facial expressions is | Attacks in this space are essentially against the FAR of the system, therefore |

| | | | | |
|---|---|---|---|---|
| | camera, or employ the services of lookalikes in attempts at masquerading as other known enrolled users. | employ the services of someone who looks similar to the target/victim. | minimal. The complexity of finding someone who looks sufficiently like a target/victim and employing them to masquerade as that victim is likely high, though if an attacker has access to a large database of face and identities (perhaps scraped from social media) then they might be able to use facial recognition to quickly find such similar-looking individuals.<br>While to the naked eye lookalikes may look similar to victims, to the facial recognition system their facial geometries may be significantly different rendering this type of attack likely ineffective on most modern systems. | the best mitigation is a robust matching process with a low FAR. |
| Manipulation of lighting conditions to affect performance. | Attackers that can control the lighting conditions around image capture devices may seek to do so in attempts at increasing the system's FRR (for evasion) or perhaps even increasing the | This would require access to lighting controls to be able to dim ambient light and or directionally focus light sources into the camera to cause glare and other | In unsupervised environments, there is no complexity here since an attacker can modify the surrounding lighting conditions in a number of ways. | Attacks in this space are essentially against the FAR/FRR of the system, therefore the matching algorithm itself is the best mitigation if sufficiently robust. |

| | | | | |
|---|---|---|---|---|
| | FAR for masquerade. | lighting effects that would affect image acquisition. | The complexity is high should an attacker wish to change lighting conditions in supervised, high security locations such as border control. | The quality control function after of the image acquisition function is also paramount – advanced systems should be able to detect undesirable lighting conditions and reject those images accordingly.<br>Systems operating under infrared illumination will be less prone to attack via lighting manipulation. |
| ‘Evil’ or ‘colluding’ twin | One twin of two identical twins may be able to successfully authenticate as the other on systems where only one of them is enrolled. | No real complexity – just the scenario requires an ‘evil’ or ‘colluding’ identical twin to attempt authentication as their pre-registered sibling. Approximately one in every 285 births results in identical twins [26]. | While to the human eye the distinction between identical twins may seem impossible, to facial recognition systems that produce precise measurement models of faces, the task of differentiating between identical twins may be trivial since the underlying geometric models of identical twin faces may be substantially different. | Attacks in this space are essentially against the FAR of the system, therefore the best mitigation is a robust matching process with a low FAR.<br>Advanced systems would likely easily distinguish between identical twins. |

## 1.5.1 3D masks from 2D images

Spoofing is typically of significant interest to biometric security researchers. For facial recognition, spoofing may involve all manner of techniques aimed at masquerading as pre-enrolled users, or in negative identification systems, looking at ways of disguise to become unidentifiable from existing enrolments. For 2D systems that lack effective liveness checks, usual spoofing attacks involve presenting 2D images (or video recordings) of a victim to a camera, perhaps printed onto paper or projected from a laptop or smartphone screen. For 3D systems and/or systems employing more advanced liveness detection, the use of 3D artefacts such as masks is a common area for exploration. In this section, we document our success in generating 3D-printed masks from 2D images to successfully spoof Android's Trusted Face and an earlier implementation of the Windows Hello facial recognition algorithm.

Our aim with this research was to understand the feasibility of taking a set of 2D images of an enrolled user and turning that into a 3D mask to use in spoofing facial recognition systems. As a potential real-world attack, typical scenarios here might include attackers taking 2D images of a victim user either covertly or from public online sources such as social media. The author was first enrolled as a user of a Surface Pro 4 tablet running Windows Hello and an Android phone configured with Trusted Face.

We identified a US-based manufacturer that could perform 2D image to 3D mask creation through use of 3D printing techniques [27]. The manufacturer provides this service primarily for novelty reasons and requires at a minimum one, but preferably three images of an individual in order to create a 3D composite and eventual 3D mask. When supplying three images it is recommended that one profile image is supplied along with two side-profile images of the same face (left and right). The following three images were taken by the author as selfies, using a plain background with an iPhone 5s camera:

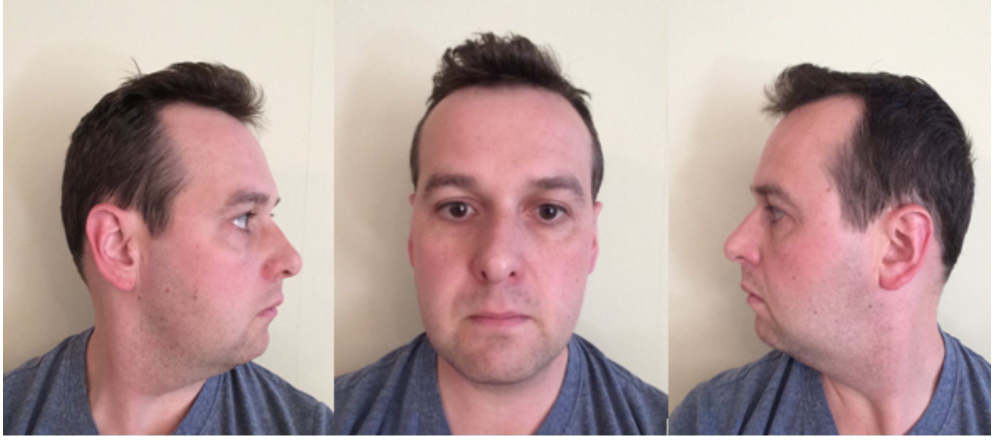

*Figure 8 - Three profile images taken as selfies on an iPhone 5s*

These images were then uploaded to [27]. When received and processed by the manufacturer, customers are provided with an image of the 3D composite for approval before manufacture. The following example previews were received and approved for 3D printing:

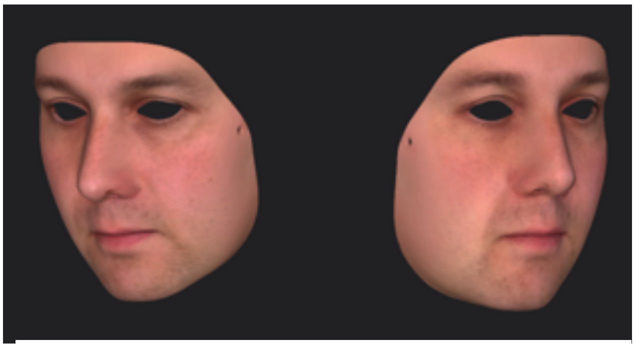

Figure 9 - 3D composite of the three 2D images supplied

The manufacturer provides the option for masks to be fully solid, or as wearable masks with cut-out eyes and nostrils. For research purposes both options were purchased at a cost of $299 (£232) each. In terms of the material and manufacture the vendor explains that:

*"The 3D portraits and heads are made of a hard resin composite in full 24-bit color with a matte varnish. The colors are not painted on but are actually part of the make-up of a thick color outer layer of the model. Due to the manufacturing process, the surface is slightly rough to the touch (sugar-coated) and some layering can be visible."*

The following images of the finished articles show the fully solid and cut-out masks as received. No facial measurement details were provided to the manufacturer, however, the resulting 3D masks are approximately life-size and in proportion to the author's real face but with a main human observation that the nose on the masks is visually different to that of the real face (straighter and more angular in composition):

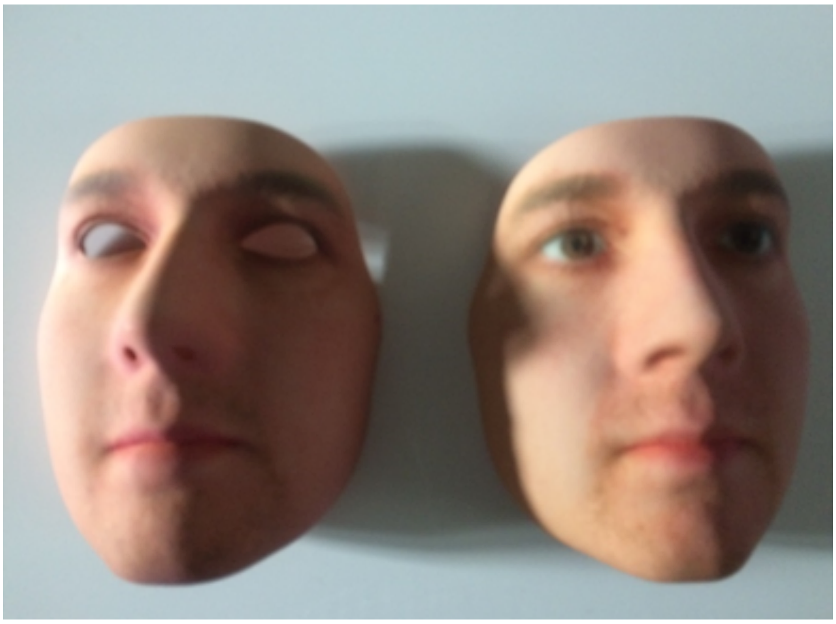

Figure 10 - The resultant 3D-manufactured masks

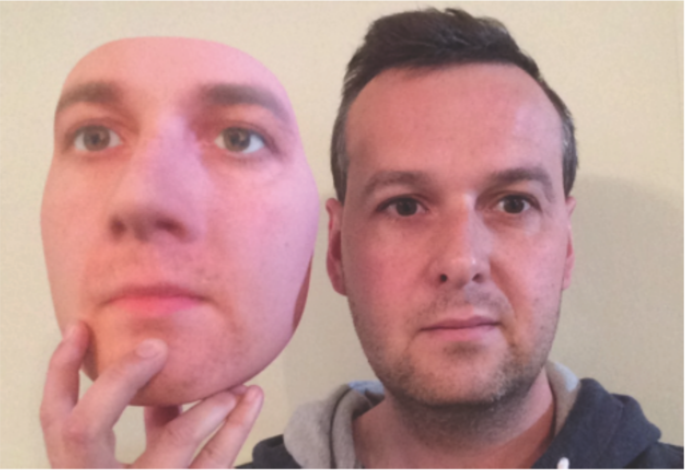

*Figure 11 - One of the masks next to the original face for comparison*

Out of the box, both masks were successfully used to masquerade as the author and successfully unlock the Surface Pro 4 tablet and Android Trusted Face. To remove the potential for error in testing, the masks were given to other individuals to present to the device cameras for authentication bypass tests. These were successful, in addition to simply presenting the masks to the camera without having to wear them.

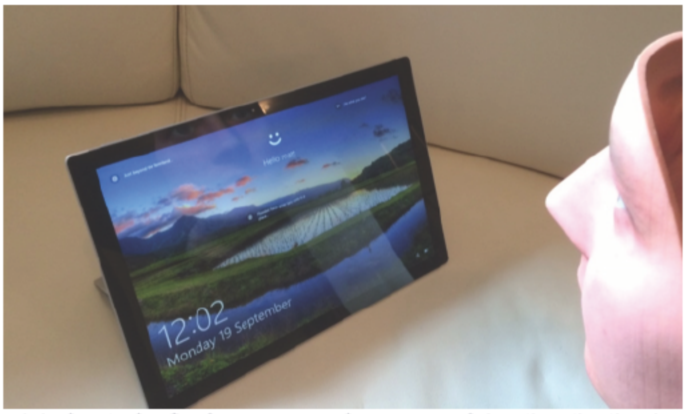

*Figure 12 - Holding the mask in front of a Surface camera for successful authentication bypass (older Hello algorithm)*

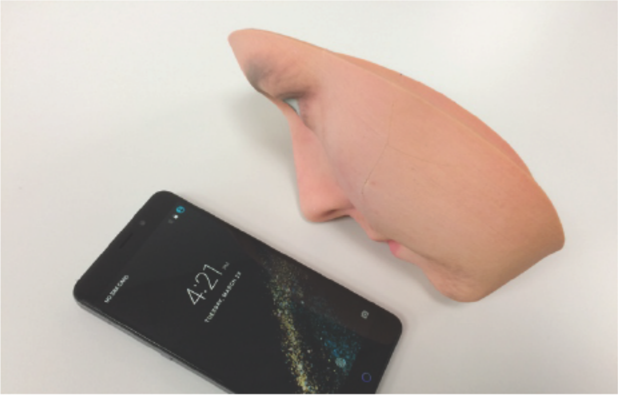

*Figure 13 - Holding the mask in front of an Android camera for successful authentication bypass*

Successful spoofing of Trusted Face with the mask was expected since:

1. There is no 3D aspect or real liveness checking to the image acquisition

2. Spoofing attacks against earlier versions of Trusted Face (Face Unlock) have previously been researched and reported [28]

3. Google's guidance on use of Trusted Face [29] reads that:

   *"This facial recognition is less secure than a PIN, pattern or password. Someone who looks similar to you could unlock your phone."*

Success with Windows Hello was more unexpected owing to:

1. Use of infrared to minimise spoofing successes

2. An FAR of 1 in 100,000

3. Machine-learned threshold for positive patching

Upon confirmation of this ability to spoof Hello we immediately reported our findings to the Microsoft Security Response Centre (MSRC) in September 2016. Over a period of a few months we worked with MSRC in performing various tests and captured telemetry from the camera and lighting components of the underlying Surface Pro 4. Following analysis, MSRC's theory on why the mask was successful is that the enrichment algorithm was too liberal in which samples it chose to 'enrich' the stored template, leading to it being eventually weakened and allowing the mask to unlock.

Since these tests and subsequent analysis, Microsoft has updated its recognition and enrichment algorithms with enhancements that specifically target the issue identified. Any Windows 10 build newer than version 15014 now has the updated algorithms. NCC Group re-enrolled on the newer build and from various tests was unable to use the same masks in successfully spoofing Hello. While this test with one user mask is not exhaustive, it does show that changes made to the Hello algorithms appear to have mitigated the issue.

### 1.5.2 Composite faces

An interesting area of research in biometrics is composite samples. This is where attackers may seek to combine two or more biometric samples to create a hybrid biometric or template that can positively match two or more individuals. Image morphing or editing software can be used to create composite facial images.

There are a few strategies available for combining faces. In a crude form, we might simply take say two images of two distinct individuals:

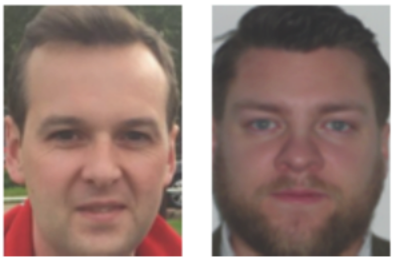

*Figure 14 - Source images of distinct individuals prior to composition*

Then create composite faces by splicing the respective faces in the X and Y axis.

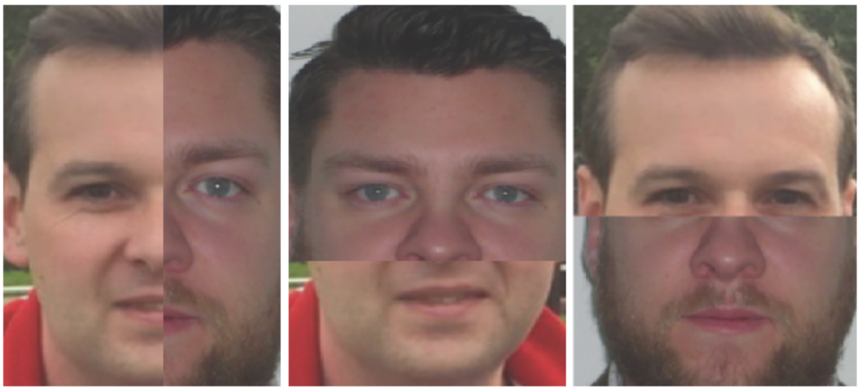

*Figure 15 - Crude X/Y-axis composite faces*

A subtler approach might be to morph [30] both images at different levels of composition to achieve a range of different faces of mutual similarity. Potentially, an image somewhere in the middle of this range might successfully match as either individual:

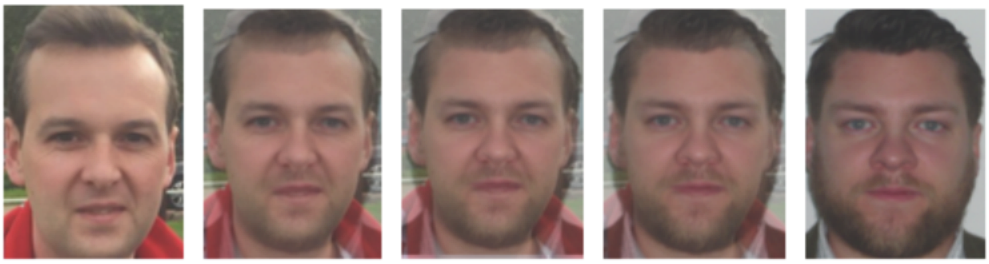

*Figure 16 - The disturbing love-children of Matty McMattface & Craigy McCraigface*

In terms of exploitation of facial recognition systems, the supposition is that if 3D masks of these composite images could be made then conceivably they might be used to allow for positive matching as either individual within the composition. A further avenue of research here might be composition of n faces, where $1 < n < \infty$. I.e. how many source images of discreet individuals might we be able to successfully combine to create a ‘master-key’ facial biometric? Our research in this field continues…

### 1.5.3 Risk assessment of facial recognition spoofing attacks

It is important to appreciate various factors involved in facial recognition spoofing in order to contextualise the risk and to understand the level of difficulty/complexity involved in facial spoofing.

In the first instance, as mentioned in section 0, we are reminded that biometrics are not secret. Unlike passwords that are (ideally) kept secret, biometrics can be copied – this is not a weakness, but rather a property of biometrics, yet with the consequence that biometric copying can in some instances lead to the production of artefacts (in this case masks) that might be used to successfully spoof a biometric system.

Hello and Trusted Face are 2D-based solutions. Note that for Hello [31] the reference reads “However, if a device is capable of providing depth in data in addition to IR, Windows will use the provided depth data to supplement the integrated anti spoofing countermeasures”. It is possible that systems employing further depth data and 3D image acquisition are less vulnerable to spoofing attacks, or at least harder to spoof that 2D versions.

Facial recognition systems such as Microsoft Hello and Trusted Face offer the ability to improve or enhance recognition through acquiring further samples under different lighting conditions and possibly facial expression or configuration (e.g. with/without glasses) which serve to further enrich the user’s template. A more feature-rich model/template of a user’s face will likely result in a stricter matching threshold for a user during subsequent authentications, therefore these enhancing techniques should wherever possible be utilised by users to minimise the risk of spoofing.

For successful attack of someone’s physical device that might be enabled with facial recognition for authentication or authorisation, an attacker requires physical access to that device. This reduces the exposure of spoofing vulnerabilities since attackers would need to:

- know or be able to identify valid victims
- be able to obtain good quality images of victim faces for production of good quality 3D masks; and
- gain eventual physical access to their devices in order to mount an attack

The factors above, coupled with ever-enhancing liveness checking (or anti-spoofing) techniques facilitate lower residual risk in this domain.

## 1.5.4 Assessing facial recognition algorithm quality & performance (Black-Box)

When performing security assessments of systems, ideally the source code will be available for review in order to help understand the composition of that system and any inherent vulnerabilities. Facial recognition algorithms are typically sensitive in terms of intellectual property for the respective vendor, therefore assessing these algorithms from an implementation perspective can typically only be done via reverse engineering and/or physical experimentation/testing through use of the system.

For performance testing, a detailed methodology for achieving this can be found in [32]. One of the challenges with biometric performance testing in general is the requirement for access to large sets of individuals in order to obtain large sample/reference sets. For 2D systems, if there is a capability to bulk enrol users in offline mode (e.g. working from images) then access to large databases of facial images can assist here. References to online facial imagery database can be found in [33]. Where systems are 3D-based and/or proprietary in how they operate (e.g. Machine Learning for matching and analysis), this may prove more of a challenge in gaining access to a database of sufficient size for testing; in cases like this the only solution may be to physically enrol live individuals. This process can perhaps be expedited through crowd-sourcing across networks if the capturing equipment is readily available on different endpoints in those networks.

Another interesting area of research around facial recognition performance, particularly where Machine Learning approaches may be used for matching and analysis, is how training models are created and from what ethnicities. For example, people of different ethnic backgrounds can look different. Suppose a matching model is trained in a lab environment on a large image set of Caucasians, but is then deployed in Africa or China. Without a sufficiently diverse initial training set there is the potential for some implementations to perform poorly when put in use for an ethnically diverse user base.

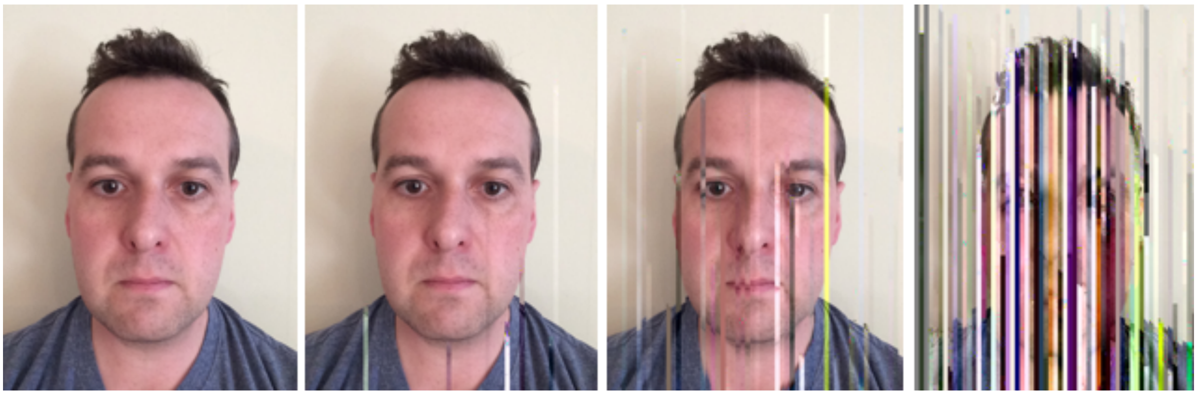

*Figure 17 - Fuzzed images using FuzzyFace (from source image: left to different levels of manipulation)*

Another angle to explore when researching facial image performance and quality is manipulation of digital images than can perhaps be bulk-enrolled. This might involve digitally altering large datasets of facial images. A very simple method for digital bit-flipping (fuzzing) source images can be found in three – this is where we use the fuzzing tool zzuf [34] and a short python script (FuzzyFace) to modify image files in ways that distort the image. While this simple script is fairly crude it could be extended to fuzz source images in many different ways – this is useful when testing quality controls and also the potential for exploiting any flaws in the acquisition function as the malformed imagery is processed (such as memory corruption or overflow/underflow flaws in code).

## 1.6 Biometric data protection – face images & templates

As mentioned in section 0, biometrics are considered personal information – for facial recognition this means that suitable controls need to be deployed to protect facial images, templates and any other facial identifiers/attributes captured when stored and rest and when transmitted across networks. In facial recognition systems (particularly physical access control applications), images are often stored/backed-up in case of a future need to re-enrol/re-create templates due to data loss or corruption. By storing original images this does not require users to be physically present at an enrolment station in order to re-enrol.

At rest, facial biometric data might be stored within local storage of an end user's device such as a laptop or smartphone app. Server-side, a backend database might be employed to store images and templates. Those implementing facial recognition systems therefore need to understand exactly where all face data is captured and stored within the system and apply suitable at-rest protections accordingly.

In transit, there are many options for encrypting facial data between client and server. Typically this will involve use of strong ciphers across the TLS protocol.

Depending on what and how facial imagery is captured, there may be scope for information leak within that imagery. So for example, while the face is the primary element to be captured and extracted by a facial recognition system, if that system acquires images or video clips of users then there is the potential for personal identifiers to be included within the captured images. Examples might include landmarks in the background of the image which expose an individual's location, or if image acquisition is done in the home, there is the potential for all manner of inadvertent background image capture such as imagery of other household members. Where imagery might be uploaded to a server-side component there is also the potential for information leakage through the image metadata that might be captured by the underlying device, such as GPS coordinates of where the image was taken.

"One of the challenges with biometric performance testing in general is the requirement for access to large sets of individuals…"

## 1.7 Conclusion

Biometrics are application-specific. For some applications, they are an obvious choice and can work well for providing authentication services. For other applications, they are the wrong choice and so before implementing any biometric system or augmenting existing authentication systems with biometrics, a feasibility study should be performed to understand the biometric advantages and limitations.

For facial recognition, spoofing will remain a topic of interest for researchers and attackers owing to the various techniques available for acquiring images of victims and turning those into 3D artefacts to attempt spoof attacks. This type of attack will be more applicable (and likely) for unsupervised applications, such as use on end user devices for device and/or app authentication. As noted however in 0, the residual risk of these attacks can be low owing to a number of factors that need to be satisfied for successful attack.

Supervised environments will have the luxury of human inspection and certainly this should be diligently implemented on critical systems such as border or building access control.

For systems with high assurance requirements, facial recognition as a sole authentication factor might not be sufficient. Its use as an additional authentication factor (e.g. used in conjunction with valid PINs/Passwords/secure tokens) would help improve assurance around user authentication.

The role of liveness checking or spoof detection will be important to the continued adoption and effectiveness of this biometric. The technology and methods for achieving this continues to develop and more crucially, reduce in cost; it is anticipated that face spoofing will over time become very difficult owing to techniques such as thermal imaging.

Particularly in financial applications, when used in negative identification mode, facial recognition can provide useful insight into potential fraud on the underlying system which is a powerful feature of biometrics that should be adopted where possible.

Where facial recognition is used for providing access to sensitive/critical data or functions, ideally the matching and template storage should all be done server-side to minimise the potential for tampering and abuse at the client-side. For example, a facial recognition Internet-banking application which performs the biometric matching on the user’s handset itself would not be a good approach as this would present the opportunity for the matching process and subsequent authentication outcome to be modified or subverted.

Facial templates and images should be suitably protected in transit and at rest – system integrators and owners need to understand where all biometric data is captured, stored and transmitted within biometric systems. Despite the onus on integrators to protect biometric data, end users also play a role in the security of their facial imagery. Users need to be cognisant of the fact that biometrics are not secret – exposure of their facial images on public websites or social media may increase the likelihood of them falling victim to spoofing attacks. End users with concerns in this area should seek to limit their online facial image exposure. Similarly, end users should ensure that their devices are not left unattended for extended periods of time, since it is when their devices are accessible to attackers that the potential for spoofing attacks against them increases.

Finally, the performance claims of face recognition system vendors need to be independently verified for maximum assurance. For example, it is not sufficient for a vendor to claim an FAR of one in 1,000,000 without being able to back this up with independent statistical research comprising cross-comparison tests across a large sample of users. System integrators should also work with vendors to understand the matching thresholds available for configuration within facial recognition systems and choose a threshold that provides the right level of FAR vs. FRR, commensurate with the overall security and assurance requirements of the underlying system.

# 2. FuzzyFace

```
# FuzzyFace - Use to generate fuzzed facial image files
# Matt Lewis, NCC Group 2017
import os
# preserve the image header by adding it back to the fuzzed file
def fixupfile(header, filename):
       f = open(filename, 'r+b')
       f.seek(0)
       f.write(header)
       f.close()

jpg_header =
'\xFF\xD8\xFF\xE1\x21\x09\x45\x78\x69\x66\x00\x00\x4D\x4D\x00\x2A\x00\x00\x00\x08\
x00\x0B\x01\x0F\x00\x02\x00\x00\x00\x06\x00\x00\x00\x92\x01\x10\x00\x02\x00\x00\x0
0\x0A\x00\x00\x00\x98\x01\x12\x00\x03\x00\x00\x00\x01\x00\x06\x00\x00\x01\x1A\x00\
x05\x00\x00\x00\x01\x00\x00\x00\xA2\x01\x1B\x00\x05\x00\x00\x00\x01\x00\x00\x00\xA
A\x01\x28\x00\x03\x00\x00\x00\x01\x00\x02\x00\x00\x01\x31\x00\x02\x00\x00\x00\x06\
x00\x00\x00\xB2\x01\x32\x00\x02\x00\x00\x00\x14\x00\x00\x00\xB8\x02\x13\x00\x03\x0
0\x00\x00\x01\x00\x01\x00\x00\x87\x69\x00\x04\x00\x00\x00\x01\x00\x00\x00\xCC\x88\
x25\x00\x04\x00\x00\x00\x01\x00\x00\x03\xF0\x00\x00\x04\xE6\x41\x70\x70\x6C\x65\x0
0\x69\x50\x68\x6F\x6E\x65\x20\x35\x73\x00\x00\x00\x00\x48\x00\x00\x00\x01\x00\x00\
x00\x48\x00\x00\x00\x01\x39\x2E\x33\x2E\x33\x00\x32\x30\x31\x36\x3A\x30\x38\x3A\x3
1\x32\x20\x31\x34\x3A\x31\x35\x3A\x33\x37\x00\x00\x1F\x82\x9A\x00\x05\x00\x00\x00\
x01\x00\x00\x02\x46\x82\x9D\x00\x05\x00\x00\x00\x01\x00\x00\x02\x4E\x88\x22\x00\x0
3\x00\x00\x00\x01\x00\x02\x00\x00\x88\x27\x00\x03\x00\x00\x00\x01\x00\xFA\x00\x00\
x90\x00\x00\x07\x00\x00\x00\x04\x30\x32\x32\x31\x90\x03\x00\x02\x00\x00\x00\x14\x0
0\x00\x02\x56\x90\x04\x00\x02\x00\x00\x00\x14\x00\x00\x02\x6A\x91\x01\x00\x07\x00\
x00\x00\x04\x01\x02\x03\x00\x92\x01\x00\x0A\x00\x00\x00\x01\x00\x00\x02\x7E\x92\x0
2\x00\x05\x00\x00\x00\x01\x00\x00'
ll = 0.00001                     # lower limit fuzz ratio
ul = 0.00002               # upper limit fuzz ratio
basefile = 'face.JPG'            # the base file to fuzz (face image)
numfiles = 1000               # number of fuzz files to generate
progress_count = 100             # report on progress every progress_count

basecommand = 'zzuf -s CHANGEME -r ' + str(ll) + ':' + str(ul) + ' < ' + basefile
+ ' > ' + 'CHANGEME_' + basefile

print 'Generating fuzz files...'

for i in range(0, numfiles):
        newcommand = basecommand.replace('CHANGEME', str(i))
        os.system(newcommand)
        if i % progress_count == 0:
                print 'Completed ' + str(i) + ' of ' + str(numfiles)

       fixupfile(jpg_header, (str(i) + '_' + basefile))

print 'Done.'
```

# 3. References

[1] https://msdn.microsoft.com/en-us/windows/hardware/commercialize/design/device-experiences/windows-hello-face-authentication

[2] https://technet.microsoft.com/en-us/security/dn440717.aspx

[3] http://www.androidcentral.com/face-unlock-explained

[4] https://en.wikipedia.org/wiki/Facial_recognition_system#History

[5] http://fortune.com/2016/09/06/hsbc-facial-recognition-biometrics-digital-revolution/

[6] https://www.theguardian.com/australia-news/2017/jan/22/facial-recognition-to-replace-passports-in-security-overhaul-at-australian-airports

[7] https://angel.co/biometric

[8] https://support.google.com/nexus/answer/6093922?hl=en-GB

[9] https://msdn.microsoft.com/en-us/windows/hardware/commercialize/design/device-experiences/windows-hello-face-authentication

[10] http://www.andrewpatrick.ca/essays/fingerprint-concerns-performance-usability-and-acceptance-of-fingerprint-biometric-systems/

[11] http://www.legislation.gov.uk/ukpga/1998/29/contents

[12] http://www.legislation.gov.uk/ukpga/2012/9/contents

[13] https://iapp.org/news/a/top-10-operational-impacts-of-the-gdpr-part-3-consent/

[14] https://www.nist.gov/news-events/news/2014/06/nist-performance-facial-recognition-software-continues-improve

[15] http://ieeexplore.ieee.org/document/6327355

[16] https://www.theguardian.com/media-network/2016/aug/17/facial-recognition-a-powerful-ad-tool-or-privacy-nightmare

[17] https://www.youtube.com/watch?v=LpIZ0m0y2FU

[18] http://spectrum.ieee.org/computing/software/can-biometrics-id-an-identical-twin

[19] http://www.theverge.com/2016/10/18/13304476/google-pixel-vs-iphone-7-samsung-galaxy-s7-edge-camera-comparison

[20] https://link.springer.com/chapter/10.1007/978-3-642-33868-7_19 aging

[21] http://subs.emis.de/LNI/Proceedings/Proceedings245/15.pdf

[22] http://ieeexplore.ieee.org/document/6996243/

[23] https://link.springer.com/chapter/10.1007/978-3-642-37444-9_44 plastic

[24] https://link.springer.com/referenceworkentry/10.1007%2F978-1-4899-7488-4_9108 plastic

[25] https://www.cs.cmu.edu/~sbhagava/papers/face-rec-ccs16.pdf

[26] http://www.twinstwice.com/twins.html

[27] http://thatsmyface.com/

[28] https://www.wired.com/2016/08/hackers-trick-facial-recognition-logins-photos-facebook-thanks-zuck/

[29] https://support.google.com/nexus/answer/6093922?hl=en

[30] https://3dthis.com/morph.htm

[31] https://msdn.microsoft.com/en-us/library/windows/hardware/mt282188(v=vs.85).aspx

[32] http://www.idsysgroup.com/ftp/BestPractice.pdf

[33] https://www.kairos.com/blog/60-facial-recognition-databases

[34] http://caca.zoy.org/wiki/zzuf

# 4. About NCC Group

NCC Group is a global expert in cyber security and risk mitigation, working with businesses to protect their brand, value and reputation against the ever-evolving threat landscape.

With our knowledge, experience and global footprint, we are best placed to help businesses identify, assess, mitigate & respond to the risks they face.

We are passionate about making the Internet safer and revolutionising the way in which organisations think about cyber security.

Headquartered in Manchester, UK, with over 35 offices across the world, NCC Group employs more than 2,000 people and is a trusted advisor to 15,000 clients worldwide.