**Statement of**

**Brian Martin, Ph.D.**
**Director of Biometric Research**
**MorphoTrust USA**

**Before**

**Senate Judiciary Committee**
**Subcommittee on Privacy, Technology and the Law**

# "What Facial Recognition Technology Means for Privacy and Civil Liberties"

**July 18, 2012**

Good afternoon Chairman Franken, Senator Coburn, and other distinguished members of the Subcommittee. Thank you for asking MorphoTrust USA to discuss the capabilities of facial recognition technology today.

My name is Dr. Brian Martin and as the Director of Biometric Research for MorphoTrust USA, my primary responsibility is the research and development of our company's biometric search engine. After earning a Ph.D. in Physics from the University of Pittsburgh, my career in biometrics began at a startup company called Visionics, which pioneered software-based face recognition technologies. Over the last 15 years, my teams have played an integral part in the research and development of world class algorithms and search engines for face, iris, and fingerprint matching. These technologies are used by the U.S. Departments of Defense and State, the FBI, over 30 face recognition-enabled drivers' license systems, and by several large international biometric systems including the world's largest in India. With a decade and a half of real world experience in developing biometric systems, I will address the following in my testimony:

- An explanation of how facial recognition algorithms work;
- A statement on the accuracy and limitations of current state-of-the-art face recognition technologies;
- An overview of the different categories of face recognition applications; and
- Some comments on face recognition technology/design as it relates to privacy.

## MorphoTrust USA's History and Role in Identity Solutions

MorphoTrust USA was formed when L-1 Identity Solutions was acquired in July 2011 by Safran, a global technology powerhouse in aerospace, defense, and security. Headquartered in Billerica, Massachusetts, MorphoTrust has over 1,100 employees across the country, including a biometrics facility in Bloomington, Minnesota.

MorphoTrust USA is the leading domestic provider of U.S. driver licenses, passports and passport cards. We provide solutions for border management, public safety, law enforcement, retail, travel and applicant vetting. We develop the technology for and deliver some of the largest, most complex biometric systems in the world, which are used for searching large databases to prevent identity fraud, provide criminal investigative leads, and fight terrorism. Our accomplishments range from introducing the first face recognition powered de-duplication of driver's license databases to providing the first commercial face detection technology to digital camera manufacturers. Under previous names, Visionics, Identix, Viisage, and L1 Identity Solutions, we have been at the forefront in the adoption of face recognition systems used by states and the federal government for over a decade.

## How Face Recognition Works

Automated face recognition algorithms were first studied in the late 1980's and became popular in the mid 1990's. Over the last 20 years, the technology has matured to the point where it can be used as a tool to help prevent identity fraud, to provide leads in criminal investigations, and fight terrorism. The technology is based on pattern recognition techniques used in the field of computer vision. Though there are several different approaches to face recognition, each with its own merits, most modern commercial grade algorithms follow these general steps:

1. Detection: First, patterns in an image are extracted and compared to, or tested against, a model of a face. When these patterns are determined to closely resemble the face model, the assumption is that a face is present in the image. This is called face detection, and in itself, can be a challenging research problem due to the large variability in what a face could look like in an image. Changes in the pose of the face, the expression on the face and the lighting (shadows) on the face make what is seemingly trivial for the human brain to accomplish an active area of research for computer vision scientists. Furthermore, algorithms have to discriminate between items that look similar to faces, but are not human faces (think of how one sees the man in the moon).
2. Registration: The next step following face detection is called 'feature registration'. The algorithms focus on the area of the image where a face was detected and attempt to determine the locations of a common set of facial features that will be used as key points when extracting the binary template or faceprint. The most commonly used registration points are the center of the eyes, but algorithms can use others, such as the tip of the nose, the corners of the mouth, etc. Once the algorithm determines the location of these points of interest, the features of the face are said to be registered or localized. If the

algorithm cannot find suitable features in this stage, feature registration will fail, causing the face to 'fail to enroll' in the system. Note that a human can aid in these first two steps, and in fact, face detection and feature registration were performed manually in early face recognition algorithms.

3. Feature Extraction: Once the features are registered, various forms of image processing can be performed to normalize the image, reduce noise in the image, reduce lighting and expression variations, and even normalize the pose of the face. This image processing helps to remove variations in the face that the matching algorithms cannot easily deal with. For instance, algorithms may not be able to match the same face in two different images if they are simply at a different scale (e.g., one is more zoomed in than the other). This stage would normalize the face in the image to ensure that it is the same size as the other faces it may be matched against. After this image processing is complete, features are extracted from the face into a binary representation appropriate for classification and/or matching. This is often referred to as a facial template or a faceprint. The feature extraction step is usually quite complex and can vary drastically from algorithm to algorithm. That is, the faceprint from one approach is rarely if ever compatible with a different approach or implementation.
4. Classification: An optional step is face classification. With the faceprint in hand, algorithms can be trained to classify the face into any number of categories that can be used to aid face matching or can just be informational. Some examples would be to use a classifying algorithm to estimate the gender or age of the face or even estimate if the extracted features of the face are of sufficient quality to support an accurate match of the face.
5. Matching: Finally, after the features of a face have been generated for two presentations of a face, an algorithm can be applied to match the two faceprints against each other to produce a single score value that represents the amount of similarity between the two faces. Depending on the features used and the efficiency of the representation of the features, the complexity of the match can be extremely high and CPU intensive taking on the order of a second per match or the complexity can be very low allowing 10's of millions of matches per second on a modern server computer. An example of a simple matching algorithm would be one that simply counts how many times the 1's and 0's are different between two binary faceprints. When the count of differing bits is low, the two faces are given a higher similarity (match) score compared to a case when the counts of differing bits are high.

Though the general recipe for face recognition is similar for many approaches, the details can vary dramatically. For example, the facial features used for matching could be texture-based, such as the pattern of a hairline or eyebrow, or the features could be shape-based, where the curvature of facial features is used for matching. 3D features can be estimated from the information in an image or can be directly measured from the image capture system. Features used for matching can be global, such as the shape of the head; they can be local, such as the shape of the eye, or they can be nearly microscopic such as the pores and wrinkles in the skin. In all cases, the approaches are vastly more complicated than the commonly perceived notion that face recognition systems simply look at geometrical distance measures between local features of the face such as the eyes, nose, and mouth.

Despite the complexity, the technology is currently at a state where these face recognition algorithms can be deployed in anything from cell phones to large multi-server search engines capable of searching over 100,000,000 faces in just a few seconds with operational accuracy.

## Accuracy of Facial Recognition Technology

For almost two decades, the U.S. Government has benchmarked the accuracy of automated face recognition systems. The National Institute of Standards and Technology (NIST) is currently viewed as the worldwide leader in independent benchmarking of state-of-the-art biometric technology. In NIST's 2010 face recognition report (NIST Interagency Report 7709 - Report on the Evaluation of 2D Still-Image Face Recognition Algorithms), it was shown that the best face recognition algorithms have improved by two orders of magnitude (over 100 times better) over the last decade. That is, the best algorithms can correctly determine if two faces belong to the same person 99.7% of the time while at the same time only making a mistake by falsely matching a face to the wrong person 0.1% of the time. In 1997, the algorithms could determine only if faces belonged to the same person about 50% of the time at this 0.1% false matching rate. Current state-of-the-art algorithms can in fact match faces as accurately as humans who are not trained to be experts in face matching.

This high accuracy is realized when the faces are captured in a controlled or staged environment with cooperative subjects. When the face is not looking directly at the camera, when there are strong shadows on the face, and when the image resolution is low (as one would expect from convenience store surveillance video) the accuracy of facial recognition can approach that from the late 1990's, making these scenarios an active area of current face recognition research. It is likely that recognition of faces captured in uncontrolled environments will dramatically improve over the next couple of decades, and as camera technology improves, it may allow face matching at accuracies close to what we see now in ideal conditions. Nevertheless, even when accuracy is relatively low, face recognition still proves to be a valuable tool for investigative searches. For instance, if face recognition can provide an investigative lead to a crime only 50% of the time, it is still helping to solve crimes whereas in the past it would have been unlikely to find any leads on suspects.

## Face Recognition Applications

In my earlier remarks, I addressed the concept of accuracy, but to be fair, there are different measures of accuracy for different types of applications. There are generally two main types of applications: verification and identification. The first, verification, is where the face recognition system is used to verify that you are who you say you are. This applies to the scenario of using a biometric to open a door or using your face to unlock your phone.  The accuracy numbers mentioned above reflect precisely this type of application.  Arguably more demanding (and more useful) is the application of identification where face recognition is used to determine an unknown identity from a gallery of known identities.  This would include applications where a photo from a crime scene could be compared to a database of known offenders to generate investigative leads or where faces in the database are compared to the database itself to detect people committing identity fraud by enrolling twice under different identities.

With identification, the requirements on an algorithm are very demanding since in order to perform a single identification, the recognition system must, in an oversimplified description, perform multiple verifications – one for each member in the gallery.   Identification from a database of one million faces is, in effect, the same as performing one million verifications, and consequently the algorithms should be one million times more accurate to ensure similar false positive match rates compared to the verification application.  Though most modern identification systems cannot be simplified to performing several verification attempts, the point is that as the database of identities grows to the size of millions of records, the ability to perform accuracy matching becomes that much harder, and requires that much more matching power.   With today's computers and state-of-the-art face recognition technology, tens of millions of records can be accurately matched per second on a single computer enabling very large scale identification applications while only taking up a relatively small hardware footprint.

Unlike verification where the system runs with very little human intervention, most large face identification systems require an expert human operator  to be 'in the loop' since the chance of getting a false match in identification is directly related to the size of the gallery of faces.  These expert operators either help narrow down and direct the facial search to a smaller target set of individuals or they are used to validate the potential match candidates from a facial search. This would be the case where a criminal investigator could narrow the search to suspects who live in a specific neighborhood where a crime was committed and because the pool of candidates is smaller, there is a better chance of finding a correct face matching result.  Similarly, the investigator would review the recommended list of candidates from the face recognition technology after the search is complete to validate or correct the face algorithms decisions. This is just what one is doing when they correct the automatic face labels generated by photo organization software that uses face recognition.

An offshoot application of facial recognition is simply the ability to classify or characterize faces.  In this application, there is no actual matching of faceprints to known or unknown identities. Instead, the features of the face are analyzed to estimate any number of aspects of the captured face. These could include gender, age, expression, at what direction the subject is looking, etc… with the most obvious use being to collect demographic information about the subjects in front of a camera.  Since this can be performed in real-time with a video camera, advertisers can use this information to tailor the advertisements based on the faces looking at them by dynamically changing the content of digital billboards.

### Face Recognition System Design and Privacy

MorphoTrust USA face recognition search engines are designed to store and search faceprints anonymously. That is, the faceprint is identified only by a system-specific number, which the customer of the system must link to the person's identity. This makes a faceprint match, by itself, relatively useless unless the attached identity information is available. Customers that implement face recognition systems take the responsibility for the security of their identity data, since it is this customer-owned identity database that connects various metadata to the individual, such as account numbers, home address, etc. These connections to other metadata are what can be exploited in unexpected ways.

In terms of face recognition technology, the faceprints themselves contain no more information than what was in the original images from which they were derived. If a faceprint database were compromised, the faceprints could not be reverse-engineered to recreate the original face image since the faceprints are stored in a proprietary format. The faceprints are also vendor specific and are of little use outside the system. Therefore, the accessibility of the original face image is typically the gating factor in preventing use of the face for unforeseen applications in the future.

### Conclusion

Face recognition is a mature technology capable of searching millions of faces in less than a second on modern computer hardware. The usefulness of the technology has been validated over the last decade by several government customers as a tool for fighting identity fraud, crime, and terrorism. Over the next decade, we expect that face recognition will dramatically improve in matching faces from uncontrolled capture environments and additionally improve match efficiency in the controlled cases, further broadening the scenarios where face recognition can be used effectively.

Thank you for the opportunity to address the Subcommittee on these important issues. I look forward to answering your questions.