Why the mobile biometrics surge demands true liveness

John Wojewidka

John Wojewidka, FaceTec

In just a few years' time, at least half of the world's consumers are expected to be using mobile devices as their main form of ID, and for primary access to their banking, medical, insurance and other confidential accounts. Devices and networks will continue to become more powerful, and service providers will have to meet these market demands with more secure mobile apps.

Increasingly, their answer to this need is strong, reliable mobile biometrics. In fact, forecasts by Acuity Market Intelligence[1] suggest that by 2019 all smartphones will have at least some kind of biometric technology on board, and by 2020 the same is expected to apply to wearable tech and tablets (*see Figure 1*). But, the rapid rise in mobile access increases people's exposure to breaches to such a level that consumers will require more than just another form of user identification to safely log into confidential, sensitive accounts. Truly secure login demands a much higher level of certainty – and that can only be achieved with what has been considered the holy grail of combined biometric authentication, concurrent identification and liveness verification.

Current technology lacking

To date, nearly all mobile account login solutions, including legacy biometrics, have failed to offer effective and reliable security. Password, PIN, fingerprint and 2D facial recognition are all actually convenience features – well-suited to opening devices and, at most, facilitating small mobile wallet transactions, but highly vulnerable to being hacked or spoofed. Likewise, fingerprint sensor hardware is inconsistent in terms of false acceptance rates (FAR), has usability limitations when cold or humid, and cannot be used with gloves. Eye/iris scans using infrared are ineffective in direct sunlight and can even hurt some users' eyes; and 2D face scans have challenges with dark conditions and can false-reject when the user changes hair, makeup or puts on glasses. Even a 3D-depth camera implementation does not inherently prove liveness, and can also struggle to identify users when there are changes in hair, makeup and glasses.

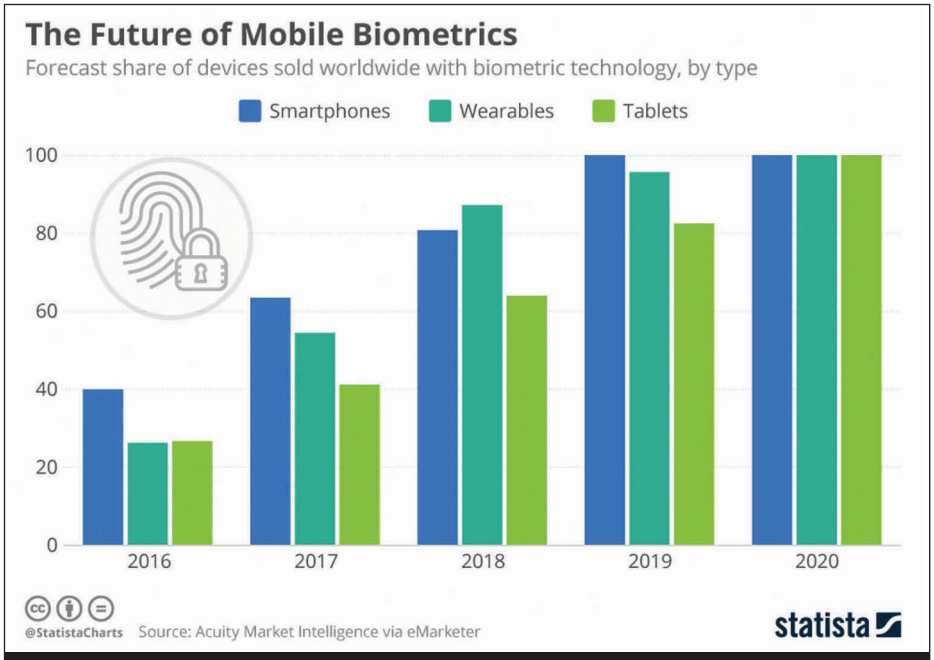

By 2019 it's predicted all smartphones will have at least some kind of biometric technology on board, and by 2020 the same is expected to apply to wearable tech and tablets.
Source: Acuity Market Intelligence via Statista

> ***"Password, PIN, fingerprint and 2D facial recognition are all actually convenience features, highly vulnerable to being hacked or spoofed"***

Spoofing occurs when a sensor, like a fingerprint reader or camera, is presented with a non-human representation of a human trait, and it is prevalent. Fingerprint imprints can be acquired from recently touched objects and even developed from photos. Most commercial readers, particularly those in mobile devices, can then be fooled with a recreated imprint. Apple's Touch ID was spoofed within 24 hours of the iPhone 5's introduction. Again, almost all face recognition solutions using 2D cameras can be scammed quickly with a photo, crude animation, video, 3D non-human representation or even a video projection. The same is true for eye/iris scans. Multi-factor authentication (MFA) is adopted primarily because legacy single-factor methods are not secure enough, but it is inconvenient for everyday use. And although MFA gives the impression it is more secure than single factor, if any one of the factors is easily spoofed it does little more than present some extra hassle to a hacker.

Not to be confused with MFA, multimodal approaches offer alternative biometric choices to the user but these do not actually increase security. By allowing the user to choose one option over another, they also allow hackers to choose the most easily spoofable modality

as well. Typically, multimodal solutions are offered via badly informed biometrics integrators who only create the apps, not the actual authenticators. By giving the user the choice of fingerprint, voice or face, they effectively triple the attack surface. Meanwhile, behavioural biometrics and continuous authentication are effective in highly controlled environments, but generally only detect a breach after it occurs, meaning the damage is done and the breach will simply get reported.

Recently of course, Apple made a bold leap into a future built on more secure biometrics by replacing its Touch ID fingerprint sensors with Face ID. This is a 3D face recognition system that uses expensive proprietary hardware – notably dual cameras, infrared technology and a neural-network AI chip – which together determine three-dimensionality and identify the user. But while this approach has raised the biometric security bar, it has not yet been verified by third parties as non-spoofable. Another issue is that no matter how effective the solution is, the exclusively-priced iPhone X can only provide Face ID security to a fraction of total smartphone users: the $999 starting price is high and production will, according to reports, be constrained for at least a year. And it's unclear whether the complex hardware will find its way into Apple's less expensive phones in the next major refresh.

Converging trends

As users everywhere have become inseparable from their mobile devices, and are more aware of the potentially devastating effects of account breaches, several important global trends are converging and contributing to the acceptance of liveness as necessary and viable.

Software and hardware development over the past few years has produced better and faster solutions, yet they may still fail to meet increasingly demanding market requirements. The industry is painfully aware of the problem, but is stuck re-purposing old technologies and testing standards. Focused on metrics like false acceptance rates (FAR), suppliers often don't even mention false rejection rates (FRR), let alone discuss liveness detection. FAR-based metrics are projected on to screens during presentations but they don't mean much in the real world. In biometrics, FAR can't be looked at in isolation; when the FAR is increased, the FRR is too. Meanwhile, the move toward more sophisticated, and of course more expensive, hardware-based depth sensing technology is good for determining three-dimensionality, but not necessarily for proving liveness. Depth sensing cannot alone determine whether the sensor sees a presently-live image, even if it can discriminate between, for example, a 3D head and a photo or video.

Third party testing and certification is critical as we move forward, but in the past there have been very few if any commercial presentation attack certifications offered for biometrics. The reason why liveness verification has been so difficult to achieve is there hasn't been an urgent need for testing, because every legacy biometric would have failed. Liveness detection has been in the crosshairs of researchers worldwide for over a decade, but has not been successfully achieved until very recently.

"There is a long list of failed attempts to build effective liveness detection, but we are now seeing expensive and complex solutions starting to enter the mainstream"

Vendors have been making claims that may seem plausible, and some have even 'demonstrated' their product's effectiveness. But it can be nearly impossible to tell if what is presented is actually doing the job the company claims. Several biometric products have made it to market only to be spoofed within hours of release. The good news is that several highly regarded verification and testing companies are hard at work to catch up, and we should start to see more sophisticated, and much more relevant, solutions within a year.

The best results for detecting liveness have come from artificial intelligence, and more specifically neural networks and machine learning. These methods represent a fundamental change in the way biometric solutions are developed. AI is quickly coming of age, and as with other complex tools, it works best in the hands of those who both understand their limitations and know how to work around them to tease out the best results. Developers indicate that models and development cycles are getting dramatically faster, while providing much higher levels of accuracy. All told, biometrics powered by AI could not have come at a better time for a market that is increasingly desperate for effective solutions.

Redefining liveness

There is a long list of failed attempts to build effective liveness detection, and they have taught us two things. First, this area is important enough to plough serious resources into, being now understood as the defining factor in true authentication as users and vendors seek password replacements and legacy biometric alternatives. Second, while this area has been hotly debated for over a decade, we are now seeing expensive and complex solutions starting to enter the mainstream.

The existing methods claiming to represent liveness – including eye blink, on-screen prompt movement, face or head movement – can only prove a non-current relationship between a human and what the sensor sees. But visual human attributes can be separated easily from a living, breathing, actual human by spoofing with photos, videos, projections or 3D inanimate representations, like full-head masks and fake busts. YouTube, the world's repository for the contemporary history of just about everything, has hundreds of examples of how sensors were spoofed with chewy candies, cats' paws or easily acquired photos and video from the internet. Even fingerprint images lifted from a publicly available photo taken from several metres away have been used to access a dignitary's personal files.

At these levels of spoofability, nearly all current access methods are simply a form of identification – still only a convenience feature – and cannot be considered truly secure. For liveness to transform identification to true authentication, a more exact definition is required. Bearing in mind that liveness cannot be determined by sensing an inanimate representation of the correct user (as above), the solution must accurately and quickly identify the image of the current user as correct by matching images. But it is imperative that liveness verification – determining if that user is a real, awake, living human with those matching physical attributes – must happen at the same time.

To maintain the highest security levels throughout the authentication process, it is vital that the data gathered from all this activity must be processed and stored securely. This requires images to be processed into biometric data, encrypted and then sequestered in the most secure zone on the device itself. This in turn means the biometric data only lives inside that single device's content, not in an on-premises database or in the cloud along with millions of other users' personal data. And there's a nice performance bonus – because it doesn't require several calls to a server over a mobile service, the solution will consistently deliver much faster results to the user. In summary, true authentication demands accurate identification of the correct user, and verification that they are alive and at the controls at login time. Otherwise, it's just another form of identification/recognition, which has clearly been proven to be spoofable.