

EXECUTIVE BRIEFING

Biometrics, fairness, and inclusion

7 ways to know you're minimizing demographic bias
when using biometrics for onboarding new customers




Steve Ritter

Chief Technology Officer, Mitek

February, 2021

Mitek





We've all been hearing about the potential for demographic bias in AI-based biometrics, particularly when used in law enforcement, at airports, and in other security situations. These concerns are justified, as poor implementation of the technology or deliberate misuse could result in discrimination and exclusion.

But to my thinking, it's just as important that biometrics be fair and inclusive in myriad other everyday life situations. Are biometrics keeping me safe or just getting in my face when I open a new credit card account, sign up for a grocery shopping service, or apply to become a seller in an online marketplace?

This is a really big deal because by 2022, according to Gartner, AI-based face comparison will be used by 80% of organizations for document-centric identity proofing in the onboarding of new customers. The technology will be a gatekeeper, helping to determine whether an individual has access to essentials like credit and a whole range of wonderful new digital services for enhancing and simplifying life.

From competitive advantage to **core requirement**

The pandemic-accelerated shift to digital transactions and services has created an immense challenge for businesses and government agencies: Amid the torrent of new applications for digital accounts and services, how do you verify applicants' identities without real-world interaction?

Clearly, we need a way to verify applicant identities remotely. This capability, once seen as a competitive advantage in terms of customer experience, has become, according to [Gartner](#), a core requirement to continue operating.


You can see the magnitude of the problem in Europe. I've read reports that [some 40% of consumers](#) couldn't access financial services during the early months of the pandemic—presumably because branches were closed and banks weren't ready with digital onboarding and identity verification.

You can see it in the US, where difficulty verifying applicants' identities to US federal unemployment insurance programs created a bonanza for fraudsters.

[The administrator of the programs put it bluntly](#): "We literally have billions of dollars at this point walking out the door under these programs due to identity theft and lack of ability to deal with that verification."

That's the downside of being unprepared to verify identities remotely. But there's also an immense upside for businesses and organizations that implement reliable, proven remote identity verification solutions. This is a moment when huge opportunities are being created for the growth of digital products and services. But success will depend on providers being able to trust consumer digital identities and consumers being able to trust providers are recognizing who they are.

TRUST IS ESSENTIAL FOR
EXPANSION
OF DIGITAL PRODUCTS AND SERVICES



95%
of businesses are confident
in their ability to recognize
their customer



55%
of consumers don't agree

[Experian 2020 Global Identity and Fraud Report](#)

Dependable biometrics for creating trusted digital identities

Document-centric identity proofing, incorporating face comparison biometrics—one of the fastest growing biometric use cases in the B2C space—has already proven to be one of the least biased.

A Nov 2020 article in *Nature*, [“Is Facial Recognition too Biased to be Let Loose?”](#) concludes that this type of biometrics “has become extremely accurate.” When you’re using face comparison to verify the rightful owner of a government-issued document such as a driver’s license or passport, “artificial intelligence is as skillful as the sharpest-eyed human.”

I wouldn’t go quite that far, but I believe this technology is one of the fairest commercial use cases of biometrics. It’s also very, very fast. (In the UK, Mitek and partner Digidentity have processed more than a million new GOV. UK Verify accounts—at up to [400 applications a minute](#).)

Identity proofing is also quite versatile. Best-in-class solutions have flexible modular architectures, including dozens, even hundreds of AI and computer vision algorithms performing specific tasks, as well as 100% configurable software development kits (SDKs) for creating user experiences. As a result, identity proofing can play different roles at various points in the customer journey. For example, you could use it to provide corroborating evidence for password-less authentication or as a second authentication factor in step-up security based on risk.




- Gartner 2020 Market Guide for Identity Proofing and Affirmation



What is document-centric identity proofing?


Use of facial biometrics, computer vision and other AI to determine if an identity document submitted via digital channels is legitimate and belongs to the applicant

How does it work?



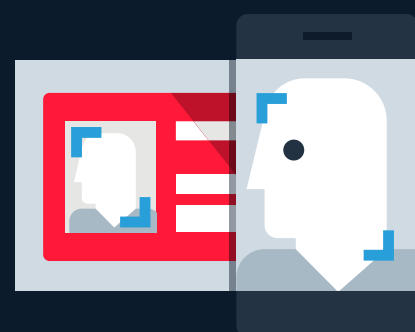
CAPTURE IDENTITY DOCUMENT

User interface (mobile app, mobile web or online onboarding) guides new account applicants to snap high-quality image of physical ID in first attempt.



VERIFY AUTHENTICITY


AI and computer vision algorithms recognize and classify the ID document, extract data from it, and evaluate authenticity. Confirm is genuine and unaltered.



PROVE REAL-WORLD IDENTITY

User interface guides applicant to take high-quality selfie. AI face recognition biometric compares selfie with ID photo. Is the person in the selfie live? Is the same person as in the ID?

Where is it used?



ONBOARDING

Identity proofing often used at beginning of onboarding process to establish the digital identity and:

- Auto-fill of application forms for fast, low-friction process
- Pass extracted information to other software for additional background checks

Some organizations use it later for proofing only those identities that cannot be verified by data-centric means.



RISKY TRANSACTIONS

Identity proofing may be invoked for added security where transaction risk is high, such as in:

- High-value transfers of funds
- Account settings changes (phone number, address, etc.)
- Sudden activity of dormant account

REVERIFICATION

Identity proofing triggered by:

- Identity corroboration requests from other security components
- Scheduled periodic checks (such as comparing a current selfie against a stored document image, where customer has granted permission)
- Random security checks (for added customer protection)

Threading identity awareness through the customer journey, with a verification "switch" at certain points

How does it cooperate with other types of security?

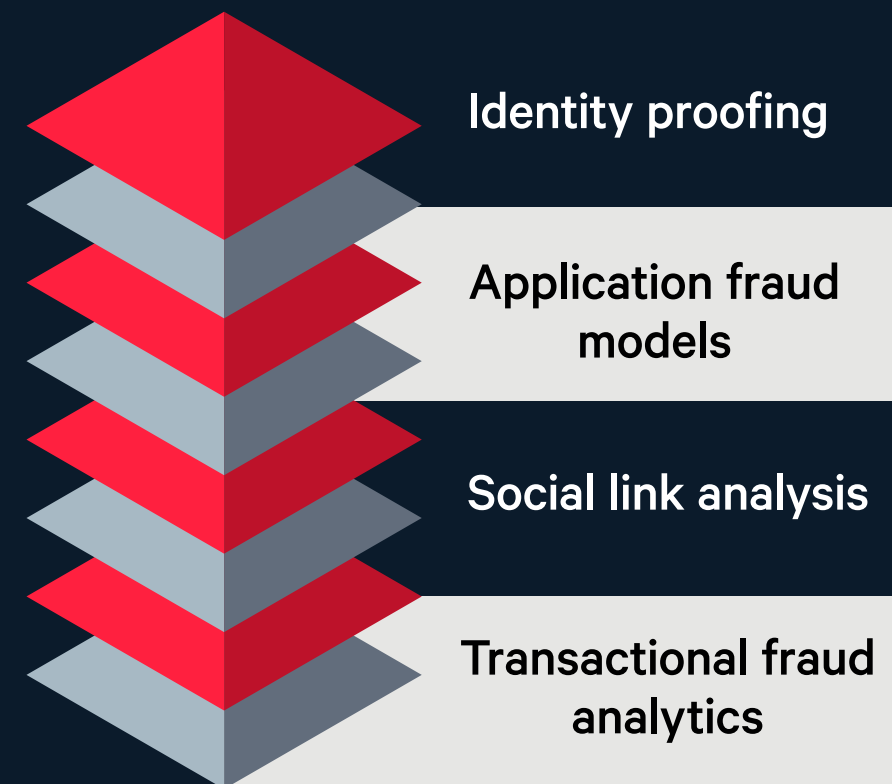
DEVICE INPUT

Identity proofing can use NFC to access stored images and data from devices and cards equipped with RFID chips.



PLATFORM ORCHESTRATION

Best-in-class identity proofing solutions can be implemented on digital identity platforms that dynamically direct interactions between layers of security.



Dynamic, adaptive approach activates appropriate layers of different situations, levels of risk and customer contexts

Flexibility and modularity in identity proofing is also one of the keys to minimizing biometric bias. Impartial biometric performance across demographics can be improved through:

Internal cooperation among identity proofing components. How is the solution helping applicants capture high-quality images that lend themselves to biometric analysis? How is it extracting and using other information from the physical ID in the identity proofing process?

Can it use both facial and voice biometrics to increase identity verification accuracy?

External cooperation with other software layers. Can the solution use NFC to access additional data and image(s) stored on an RFID chip embedded in a physical ID? Can it pass the data it extracts from ID snapshots, as well as granular analytic results, in near real-time to third-party components performing other types of identity, fraud detection, and security checks?

NEXT: Detailed guidance on what to look for in a document-centric identity proofing solution that minimizes demographic bias.

How do you know your biometrics aren't biased?


Some people say it's impossible to build AI-based biometric face recognition technology that is accurate across all demographic groups. I don't believe that's true. Problems have to do with how these biometrics are developed, used, and governed—all solvable.

Much of the concern is in response to a [2019 report from the US National Institute of Standards and Technology \(NIST\)](#). Testing face recognition algorithms from the majority of commercial developers, NIST acknowledged that accuracy in face matching—both for one-to-many searches and for the one-to-one verifications used in onboarding—had massively improved over the previous year's testing.

Even so, results show we need to keep up that pace of improvement. For instance, in one-to-one matching, NIST found that the false positive rate (incorrectly matching two images of different faces) was 10 to 100 times higher for African American or Asian faces than Caucasian faces.

In real-world use, that's a security problem since a positive match means the applicant—possibly a fraudster—would likely be onboarded.

NIST testing also found that female and younger faces tended to have higher rates of false negatives (failing to match two images of the same person). That's a discrimination problem since a negative result means the applicant would likely not be onboarded.



Performance varied among the algorithms tested, of course. So my top recommendation for organizations seeking to minimize biometric bias is:

How to know #1


Choose an identity proofing vendor that uses one of the top five facial recognition algorithms in one-to-one matching. Ask them for their results from the NIST Face Recognition Vendor Test (FRVT). You are looking for no more than 0.005 in face matching results across race/ethnicity and age demographic groups. Also, ask for internal tests and other third-party testing (like iBeta Liveness Certification). Make sure you get updated test results annually from all sources.

It's important to realize that the NIST FRVT tests face recognition algorithms only—not how they perform within a document-centric identity proofing solution. This matters because, as NIST acknowledges, one of the key factors for demographic bias is image quality. It's affected by how good the solution's image capture software is and how effectively the user interface guides applicants through the onboarding process.

With that in mind, here is my wider take on the issues around biometric bias you'll want to pay attention to and recommendations for how to know you're following best practices to mitigate them.

Pay attention to algorithm accuracy and consistency

The best face recognition algorithms are high-performing in terms of both accuracy and consistency. An example I gave in a recent [Forbes article](#) is that if an algorithm can regularly identify 90% of white males faces but only 40% of black female faces, it's accurate but not consistent, and that perpetuates the problem of bias. Alternatively, suppose it's able to consistently identify 70% of all faces 80% of the time. In that case, it is slightly less accurate but is ultimately a better, more equitable tool and more beneficial for your business.



How to know #2

Make sure you're assessing both dimensions of performance—and that the algorithm will be able to deliver those metrics *for your specific use case and target populations*. For instance, face recognition for document-centric ID proofing is a more difficult technical challenge than face recognition for device access. When I set up Face ID on my iPhone, I'm guided to move my head in slow circles so the software can record lots of images it will later use for comparison with my real-time face. In document-centric onboarding, the software sees both the selfie and the face image on the ID for the very first time. The algorithm has to be properly trained to perform well under those circumstances.

Choose an identity proofing solution that uses the right algorithm for what you're trying to do, and you'll reap major rewards. Not only will you be reducing biometric bias, but you'll also be better able to serve your entire user base—not just a single demographic—and make the onboarding process easier for all your new customers.

Pay attention to the composition of training and testing data

To achieve the highest face recognition accuracy and consistency across your demographics, the algorithm you use has to be trained on biometric data that closely aligns with the data it will be analyzing after deployment. When that doesn't happen, you get bias. For instance, in 2018, the MIT Media Lab publicized research showing that some commercial face recognition systems had [trouble recognizing people of color and women](#). Not surprising, the Lab pointed out, since one widely used training image dataset was 75% male, 80% white. Conversely, NIST tests using its Mugshot Identification Database, found that the top 20 performing algorithms were [most accurate on Black male faces](#). It could well be that some of these algorithms were trained with that very dataset—which likely contains numerous mugshots of Black males, who are arrested at a much higher rate than other demographics.



How to know #3

Analyze the demographic breakdown of your target population. Give this information to any vendor of identity proofing solutions you are considering. Require them to provide evidence their training and testing datasets reasonably align with your population demographics. Don't just rely on the assurance of choosing an authoritative data source. (This past summer, [MIT was embarrassed](#) to confirm it was taking down one of its own online training datasets when third-party researchers found that some of the images of Black, Asian, and female individuals had been labeled with derogatory language!). Training and testing datasets should also include environmental variation (the same people wearing glasses, wigs, hats, and so on) as well as techniques fraudsters use to thwart liveness detection (such as 2D and 3D masks, printed photos and cutouts, and screen and mobile replays).

Pay attention to real-world captured image quality

Face recognition performance is affected by differences between the controlled conditions in which training dataset images may have been produced and the messier real-world conditions of user-produced selfies. Naturally, biometrics developers tend to choose datasets suited to their purposes, such as images taken by professional photographers in studios. That's not what the algorithm is going to see once deployed. In the real world, different angles, backgrounds, and lighting conditions will make it harder to match faces.




How to know #4

Look for an identity proofing solution that incorporates state-of-the-art image capture software. It should have multiple levels of liveness protection, including active and passive. You want to be able to guide users to take optimal selfies (image size, distance from camera, lighting, focus, angle) *on the first try*, and make this process easy and fun. The vendor should provide a software development kit (SDK) enabling you to blend this guidance into your branded mobile app, mobile web, or online onboarding flow.

Pay attention to model governance

There's concern about letting biometrics and other AI loose to operate "in the wild" without human oversight. The problem goes beyond potential demographic bias. There's also the question: Is the algorithm becoming more accurate or less accurate over time? Is it being updated in a timely way to incorporate technology advances, support new document formats, and reflect changing population demographics?



How to know #5

Have a systematic way of tracking biometric performance and comparing it across demographics over time. As an engineer and a manager, what I want most is to be able to measure performance at a granular level so I can understand how algorithms are behaving and why. Clearly, since we're talking about machine learning, I expect performance to improve—but I want to make sure improvements occur across my demographics. *All boats should rise.*

I also recommend selecting an identity proofing vendor whose model governance practices include strong human oversight. At Mitek, for instance, teams of human experts spot-check statistically significant samples of algorithmic face matching results. (Some clients also opt to have these experts review identities that cannot be decided by the face recognition and other algorithms with a high enough level of certainty to meet their thresholds for risky situations.) The work these people do provides a second signal stream, adding to algorithmic learning, for continued improvement. It also helps us continually confirm or challenge assumptions about client population demographics to make sure algorithms continue to be trained appropriately for updates and releases.

Pay attention to data storage and privacy

Consumers are suspicious of biometrics. [Research by PYMTS and Mitek](#) found that less than 30% of consumers were comfortable providing biometric information, such as fingerprints, voiceprints, or photos for facial recognition. (Although 64% said they would be more comfortable if they understood their biometric data would not be shared with third parties). This distrust is not confined to biometrics but extends in general to the data consumers are being asked to provide when opening online accounts. Another study of [consumers in the UK, US, France, and Germany](#) found that the top two reasons people abandon sign-up processes for online shopping were fear their data would be passed on to third parties (53%) and concern about their information not being secure (50%).



How to know #6

Make sure your identity proofing solution has, by default, a minimum data retention policy so that it's not building up a repository of personally identifiable information (PII)—and, of course, never provides such information to third parties. Solutions should also incorporate multi-layered encryption to protect customer data both in transit and at rest. For instance, Mitek uses a customer-specific master data encryption key to generate a limited transaction-level encryption key, all based on industry-standard technologies. We've also implemented a formal Information Security Management System ("ISMS") based on international best practices (ISO27001:2013) and have AICPA Service Organization Controls 2 ("SOC 2") certification.

That said, I don't want to imply that identity proofing should be just process-and-forget. You want to look for a provider that has invested in a streaming infrastructure enabling them to extract maximum insights in real-time from identity proofing instances as they happen—that's essential for accelerated machine learning. It's also helpful if the solution can capture anonymous metadata: What was the result of identity proofing (matched, unmatched)? What kind of document was it—format, year, etc.? And even what is the likely demographic category of the identity—information that can be very useful for measuring and improving cross-demographic performance over time.

Pay attention to **transparency, readiness to provide evidence, and attitude toward regulations**

Organizations deploying or considering biometrics are coming under increasing scrutiny from both internal and external parties. Reaching out to vendors, they're sometimes having difficulty getting answers to questions and evidence regarding bias and performance. That's not going to be tolerated—[Gartner predicts](#) that by 2022, 95% of RFPs for document-centric identity proofing will have clear requirements for minimizing demographic bias.



How to know #7

Find identity proofing vendors that can readily provide you with the metrics and other information you need to bring to your stakeholders and decision makers, auditors, and legal counsel. Gartner has called for “openness, transparency, and responsibility,” and I couldn't agree more.

Also, very likely in the near future, you'll need this information to satisfy regulators (numerous governments have established or are in the process of establishing identity systems or, at least, standards commercial systems must adhere to). Look for the leaders in identity proofing who are operating ahead of the regulatory curve by already thoroughly documenting demographic impartiality and PII security.

Ask them about their attitudes toward regulations and if they're actively working toward seeing them enacted.

Never lose sight of why we're doing this

Certainly, we need metrics. Rigorous measurement of cross-demographic performance is the only way we can continue moving biometric bias closer and closer to zero. But let's never lose sight of what metrics represent and why we measure in the first place. A [2019 Harvard Business Review article](#) made the point better than I can:

Of course, we all know that metrics are inherently imperfect at some level. In business, the intent behind metrics is usually to capture some underlying intangible goal—and they almost always fail to do this as well as we would like. Your performance management system is full of metrics that are flawed proxies for what you care about. [my emphasis]

What do we care about as we work to reduce biometric bias in document-centric onboarding? In my view, it's freedom. Our software plays a key role in deciding who is free to access essentials like credit and the expanding range of digital services.

You might ask: Do all individuals really have an intrinsic right to access digital services? Is this a basic human freedom? In a world that's suddenly become digital-centric, I think the answer is clearly yes. And I expect we'll soon see this right codified in numerous ways by law and regulation. But let's not wait for that—now's the time for us all to move forward and make sure we're on the right side of history.

This document is for general information purposes only and is not intended to be and should not be taken as legal and/or regulatory advice on any specific facts or circumstances. All information provided in this document is provided "as is" without warranty of any kind, whether express or implied. Contents contained in this document may not be quoted or referred to for any purpose without the prior written consent of Mitek or its affiliates.

Mitek

Mitek Systems, Inc. (NASDAQ: MITK) is a global leader in mobile capture and digital identity verification built on the latest advancements in computer vision and artificial intelligence. Mitek's identity verification solutions enable organizations to verify an individual's identity during digital transactions to reduce risk and meet regulatory requirements, while increasing revenue from digital channels. More than 7,000 organizations use Mitek to enable trust and convenience for mobile check deposit, new account opening and more. Mitek is based in San Diego, California, with offices across the U.S. and Europe.