


paravision

Paravision's Approach to Biometric Image Quality

NIST FRVT 1:N Identification



Separating the good from the bad in face image quality

In March 2020, the National Institute of Standards and Technology (NIST) released a report assessing the efficacy of face image quality algorithms. Despite the significant impact image quality has on the performance of face recognition systems, it's public exploration is only in its nascent stages. This whitepaper attempts to shed light on this topic by discussing:

- What biometric image quality is
- Why it matters
- How NIST evaluates face image quality
- How Paravision performs relative to leading vendors

What is biometric image quality?

Biometric image quality is a number assigned to an image that is intended to correlate with an algorithm's confidence in matching results.

An image with a high quality score should deliver high-confidence matching results, while one with a low quality score may not. Many factors inform quality scores. A face with sunglasses may have a lower score than one without; an oily or oversaturated fingerprint may have a lower score than a clean, clear one; an iris hidden by an eyelid may have a lower score than one in full view.




Figure 1 - A classification of images in the order of increasing visual quality.

Why does it matter?

Image quality algorithms help improve face recognition performance by weeding out faces that are most likely to cause recognition failures. This maximizes the accuracy of face recognition systems, leading to better security and user experiences.




Security

One use of face recognition in security scenarios is to create and enforce blacklists. In such situations, face recognition algorithms that consider image quality will demonstrate superior performance to those that don't.

For instance, at a checkpoint, blacklisted individuals can evade systems without image quality thresholds by obscuring their appearance.

However, in a properly implemented system, a face with a low image quality score could trigger a secondary response, such as notifying security personnel or informing the individual in front of the camera to change their pose.



User Experience

Successful biometrics implementations require systems that work consistently in a range of environmental conditions, since a failure to do so can result in added inconvenience.

In face recognition-powered access control, not considering image quality can impact both the enrollment and matching processes. Whitelisted individuals may get stuck at the door (a false negative error) if either the enrollment photos or presented samples are not of sufficient quality.

A good image quality algorithm could be used to provide guidance on how to improve the way users present their face (pose change, lighting conditions, etc.) during enrollment as well as subsequent identification. This increases the performance of the overall system and improves user experience.

Old Approaches to Image Quality

Until recently, face recognition image quality was evaluated by how well images adhered to ISO/IEC standards.

ISO 19794-5 was the first standard that brought forward attributes - such as pose, illumination, focus, and several others - to assess image quality prior to enrollment. It stated that “besides normative requirements of size and proportion, that the face is uniformly illuminated, in focus, and captured from straight ahead with no rotation or pitching.”

ISO 29794-5 laid out methods for the calculation of image quality scores for facial images. It introduced the quantitative incorporation of concepts such as

symmetry, resolution, illumination intensity, brightness, contrast, color, exposure, sharpness, and several others.

While these criteria worked well for evaluating passport or ID photos, they were not practical in less-constrained, real-world environments. And, since these image quality standards were created to align notionally - rather than statistically - with matching algorithm performance, they were not necessarily well correlated to actual matching results.

The Modern Approach to Image Quality

Training image quality algorithms to provide utility to their recognition algorithms.


The last few years have seen an increased use of machine learning to train face recognition algorithms using large, diverse datasets. Images in these training datasets have varying levels of quality. Since different recognition algorithms calculate matching scores using different and unknown underlying calculation methods, the definition of “good” or “bad” image quality differs between algorithms. Therefore, standardized measures of image character and fidelity (like those laid out in the ISO/IEC standards) are no longer appropriate.

The mindset around image quality has shifted towards showing quality values as predictors of true matching performance. Therefore, today's approach makes use of training quality algorithms on how they can provide the highest value images to their respective recognition algorithms (i.e. based on their utility). This pragmatic approach

to testing correlates with NIST's testing, with NIST stating that ***“utility of a sample to a recognition engine is what drives outcome operationally and is of most interest to end-users.”***

For instance, in some matching algorithms, sunglasses may cause a marked decrease in performance, while in others, someone turning their face away from the camera may have more of an impact. With image quality algorithms tailored to face matching performance, the first quality algorithm would flag more faces with sunglasses, and the second would flag more faces turned away from the camera, resulting in the best possible performance for each face matching algorithm. This is not to say that one quality algorithm can't work with another matching algorithm, but that requires an intelligent look at cross-algorithm performance, such as that provided by NIST's testing.

Two technology trends are powering the adoption of this modern approach



- 1** The increased use of deep learning to train both face recognition and image quality assessment algorithms.

This is rapidly improving the performance of face recognition systems to a point where they meet and even exceed the performance of other biometric modalities such as iris and fingerprint while being more convenient.
- 2** Advancements in edge processing capabilities fuelling the push towards on-device processing of more complex workloads.

A combination of more powerful edge processors and machine learning models optimized for edge chipsets is enabling face recognition on devices like mobile phones and smart security cameras. High performing image quality algorithms only pass on images that meet a required threshold to recognition engines, thus reducing required processing power and potential network bandwidth, making them indispensable to edge use cases.

Paravision Image Quality Assessment Performance

NIST's Face Image Quality Assessment report shows how well algorithms' image quality scores improve face matching algorithm performance. The more effective an image quality algorithm is at identifying low-quality photos, the more aggressively it will remove these photos from the dataset, maximizing the performance of the associated recognition algorithm.

Face recognition accuracy is judged on two broad metrics: Type I errors (false positives) and Type II errors (false negatives). False positives occur when a new face is incorrectly matched with a known face in the dataset; false negatives occur when a known face is not matched with its correct identity.

In Figure 2, the false positive identification rate is controlled at 0.0001. The chart shows how the false negative identification rate improves as the images with the worst quality scores are removed from the dataset. An ideal algorithm would perfectly predict which photos would cause false negatives, immediately giving all of them the worst quality score and removing them from the dataset. Of the companies whose image quality algorithms were evaluated against their face recognition algorithms, Paravision's performance came closest to perfect. Paravision's image quality scores most immediately

and consistently flagged the photos that caused false negatives in the matching algorithm.

Figure 2 illustrates that when an application rejects 20% of faces with the worst image quality scores, Paravision's matching algorithm achieves a 99.98% accuracy, while CEIC's algorithm achieves a 99.42% accuracy. These differences may seem small at first glance, but large-scale face matching amplifies small differences in performance. For instance, the Hartsfield-Jackson Atlanta International Airport (ATL) has 63,000 employees that need to check in every day. With Paravision's algorithm, only 13 employees would have any trouble at the access point; with CEIC's algorithm, 365 employees would. (Without any of the worst-quality images removed, Paravision's algorithm would inconvenience 486 employees, and CEIC's would detain 677.)

Face Matching Performance with Screened Poor-Quality Images

FNIR / Perfect by Proportion of Worst-Quality Images Removed, FPIR = 0.0001

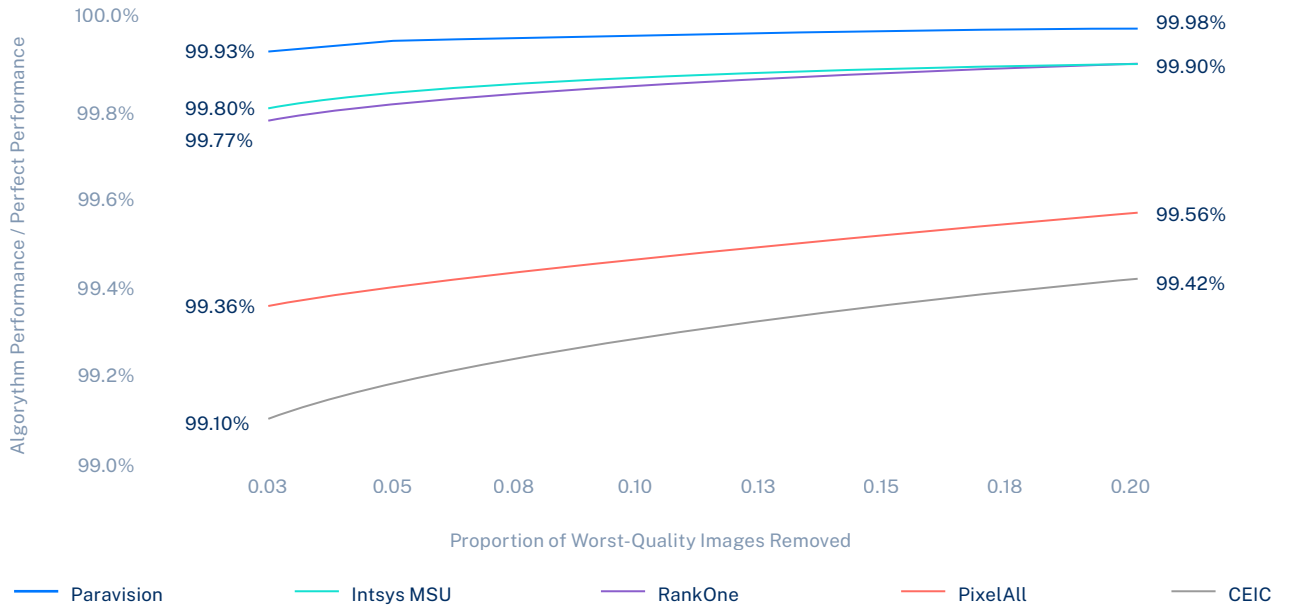


Figure 2 - This chart shows the matching performance of different face recognition algorithms when increasing percentages of poor quality images are removed by their respective image quality algorithms. This testing was performed on the Webcam Images dataset and shows matching accuracy when 2.5% to 20% of faces are removed.

Source: NIST FRVT Quality Assessment

NIST's latest testing shows that Paravision's image quality algorithm outperformed other evaluated competitors.

When considered alongside its top 5 performance on the NIST FRVT 1:1 Verification and 1:N Identification reports, Paravision offers the most complete development platform for partners building mission-critical security solutions that are accurate, scalable, and globally competitive.

Paravision's AI-powered approach to image quality and face recognition provides superior face recognition that remains accurate across a range of challenging conditions. This shift from traditional ISO/IEC standards to AI-powered image quality scoring will move the industry forward, as companies collectively improve their algorithms and underlying algorithm components.

For more information on Paravision's
industry-leading accuracy, please email
sales@paravision.ai

paravision