

The State of Facial Recognition Software in 2020

Originally published at <https://www.iflexion.com/blog/facial-recognition-software>

Published: August 24, 2020 | By Martin Anderson



No one could have predicted how much the facial recognition market would change in 2020.

Even at the start of the year, the sector was facing a growing tide of popular dissent¹ and calls from around the world for increased legislation, oversight and accountability². Then, almost immediately, the world was wearing a mask. Police officers who once would caution and fine citizens for covering their faces³ were now arresting them for *not* covering them⁴.

At the same time, the repurposing or new deployments of facial recognition systems with a COVID-19 agenda re-fermented discussion, controversies, and talk of legislative action⁵. In response, many of the biggest facial recognition providers on the planet called a moratorium on selling these services to law enforcement.

It's a year like no other, with nearly all the long-term market analysis predictions wrong-footed by the pandemic. Therefore in this article, we'll take a broader look at the current capabilities of [computer vision software development](#) for facial recognition purposes, including some of the technical, social and increasing regulatory challenges that the sector will need to overcome to regain its former momentum in a post-COVID-19 market.

There are technical, social and regulatory challenges that the facial recognition sector will need to overcome to regain its former momentum in a post-COVID-19 market.

What Is the Facial Recognition Market Worth in 2020 and Where Is It Heading?

The global market for facial recognition technologies is forecast by one source⁶ to reach \$9.93bn by 2027 from the 2020 market valuation of \$3.4bn. Another report⁷ cites a \$52bn global market value by 2023.



The Coronavirus Upends the Facial Recognition Market

However, most of these bullish projections are based on market trends and statistics prior to the advent of COVID-19 (where masks have changed the landscape of [intelligent video analytics](#)), and before the current round of controversies over privacy and governance combined with the novel disease to produce a chilling effect on the sector.

Mask usage as a precaution against COVID-19, combined with the new seclusion that characterizes the coronavirus age, not only reduces the effectiveness of facial recognition algorithms that worked well before⁸, but also radically thins down the available sample base. In most countries, there are simply fewer faces to evaluate; and most of those are occluded by masks.

Some governments have used the pandemic as a spur to install additional facial recognition infrastructure, for instance in Russia⁹, South Korea¹⁰ and China¹¹, while others have adapted existing systems to detect mask usage and social distancing¹².

In truth, these are opportunistic measures whose long-term ramifications or utility aren't yet known. Therefore, the market prospects for the facial recognition sector must remain uncertain while its fundamental model is so compromised by the practical consequences of the pandemic.


How Accurate Is Facial Recognition in 2020?

According to a recent NIST report¹³, the leading facial recognition algorithm in 2020 has an error rate of just 0.08% in comparison to the 4.1% error rate for the market leader in 2014, with over 30 other algorithms also exceeding the best score from 2014.

However, several factors stand as an obstacle to the wide applicability of any single facial recognition accuracy ranking system.

Quality of Images

NIST observes that the inconsistent quality obtained from a wide diversity of capture inputs presents a problem in terms of both achieving and monitoring accuracy across the available range of facial recognition systems.



If that data is too inconsistent, the machine learning models powering a modern facial recognition system tend to converge at a mean quality that may not accurately represent any of the input data¹⁴.

When the image data varies wildly in quality, the model must either generalize or else reject data below a certain quality threshold, depending on how it is weighted. Either way, some potential accuracy is sacrificed in the process.

For law enforcement networks, available inputs will vary from relatively crisp mugshot databases to images from blurred or indistinct video feeds, such as from CCTV systems that were outfitted decades ago and intended for manual monitoring rather than as a source of primary evidence.


Multi-Dimensional Encoding for Facial Recognition

In 2019, a study out of the University of Notre Dame¹⁵ examined the problem of varying facial capture quality and concluded that today there is no clear theoretical guidance for designing new architectures to specifically fit this task.

However, it notes that there is ongoing research into multi-dimensional models (see below) that are more capable of utilizing images of any quality without losing accuracy across the range of images¹⁶.

A second approach is to identify 'quality tiers' of data inside the neural network itself, and use this information to adapt the model's weights and to make it less reductive where data quality varies — the aforementioned 'multidimensional' method.


A research collaboration between Microsoft, Tencent, and various universities¹⁷ proposes a Feature Adaptation Network capable of distinguishing multiple scales of quality in data input via a Random Scale Augmentation (RSA).



The technique claims to achieve normalization while maintaining the feature applicability of each level of quality from diverse sources.


GAN Facial Reconstruction

Some facial recognition techniques are beginning to employ Generative Adversarial Networks (GANs)¹⁸ and adapted encoder/decoder models to produce highly detailed faces from extremely low-quality images.




This 'super resolution' technique leverages the abstraction capabilities of a neural network¹⁹: irrespective of input image quality, a machine learning model will turn a bitmapped image input into a mathematical vector formulation at the encoder stage — a unique hash that represents the essential characteristics of the input face.

The model will later use a decoder to read and reproduce that same vector information at the requested scale.



The model decoder cannot reproduce the image with greater fidelity than the original, because the original facial topography data was too sparse, and because there is usually a hard limit on input resolution for the extracted faces.

However, a GAN-trained system, or an adapted encoder/decoder configuration, can examine and explore the characteristics of high-resolution output from *other* high-resolution images, and use the vector/bitmap relationships it observes as a guide to making a detailed face, even from a blurred or indistinct input image.



Though super-resolution modelling gives the impression of restoring detail based on similar paths from degraded to high-res faces, it uses the original 'degraded' material only as a broad template. The result is essentially a 'smart guess'.



For entertainment purposes, this method has a lot of potential. For law enforcement, however, such approaches are destined to be called into question irrespective of their success, since a miscalculation can create a convincing but non-existent person who may by chance resemble someone unrelated to the input image.

How accurate is facial recognition technology in 2020?

What Facial Recognition Systems Can Do

Highly Specific Forks and Derivations

Downstream forks and adaptations of popular open-source facial learning repositories tend to be customized to add or exclude various capabilities, or else to adapt the generalized code to specific case conditions.


This can include cameras fixed at a certain height, the use of a specific resolution for input video quality, or the generation and use of datasets with a deliberate racial bias (for instance, when the system is located in a country where the vast majority of subjects will have a specific ethnicity, and where the ability to identify highly granular differences between subjects of that narrow ethnicity are essential for a functional system).

Profile and Acute-Angle Recognition

Even though side-on shots of newly enrolled prisoners became a general standard over a hundred years ago, not all facial recognition systems are capable of performing profile recognition, or processing acute face angles (the so-called '3D' facial recognition).



Some systems, such as street level and police-mounted cameras aimed at drivers, may concentrate exclusively on profile recognition, while others may include profile recognition in more general ID systems where targets may present a range of angles.



Systemic and Inadvertent Bias

Bias takes a prominent place on the list of. Machine learning models are designed to extract averages and sub-sets out of massively diverse data inputs in order to identify trends, patterns and tendencies. But if a dataset contains a majority of data points with a high proportion of a single label or classification, such as *ethnicity=white*, it will begin to see that label as a 'baseline' and anything different to it as an 'outlier'.

Furthermore, it will be better at distinguishing between examples of the dominant classified type (where the similarities form a central clique and the differences within that clique have greater room to emerge), than between the outliers, who may show only *radical* differences among themselves — and who inevitably become defined by their relationship to the core group rather than their relationship to each other.

Balancing Racial Diversity

In our recent review of [facial recognition software pros and cons](#), we described some of the major controversies of recent years regarding racial bias in machine learning, which include:

- Amazon, IBM and Microsoft temporarily withdrawing the sale of their facial recognition technologies to law enforcement entities until new legislation defines clear parameters to prevent unfair bias in derived systems (and a recurrence of the public pressure which forced them to this point).
- A late 2019 study from NIST claiming that facial recognition systems deployed by US law enforcement departments have an ID error rate 10-100 times worse for African-American and Asian faces as opposed to Caucasian faces²⁰.
- Public outcry when it was revealed²¹ that the Federal Bureau of Investigation and Immigration and Customs Enforcement (ICE) transposed state driver's license records into a facial recognition database without consent.

The popular solution for racial bias in facial recognition systems is to ensure datasets have a diverse ethnicity²². As scandals around bias have multiplied, a number of prominent companies have sought to redress this imbalance: IBM developed a more multi-ethnic million-face dataset in 2019²³, and later that year Google Contractors were even caught bribing students and homeless people of African-American origin to volunteer their faces for the Pixel 4 facial recognition model²⁴.

Well-resourced generic global facial recognition systems, such as Apple's and Google's Face ID systems, can be well-served by such diversity; but localized applications (such as a face-based entry system to a Chinese research lab, or a security clearance database for a school in Zimbabwe) may need to devote limited model space to a narrower range of facial types in order to achieve precise IDs within the broad parameters of one or two racial face-types.

The limitations of model space are rarely addressed in the heated discussions around racial bias in face datasets, but they are a hard reality for low-latency AI-based ID systems that need to run locally rather than on the cloud; for facial recognition systems that operate in monocultures; and especially for models that must run on the limited resources of [mobile-scale deep learning systems](#).

Model Development from Scratch

Part of the problem may lie with the common use of existing semi-trained models and unsuitable open-source face datasets such as CelebA, since these are frequently dominated by one racial type or another²⁵.

Unless the facial recognition development community reduces its dependence on semi-trained open-source templates, which often contain ineradicable low-level features, it may have no practical or safe way to take accountability for the software it produces.

Conclusions

Accuracy

Facial recognition seems to have reached measurable record levels of accuracy in 2020, despite the lack of uniform federal testing standards in the US²⁶, such as in various pan-global challenges, most of which revolve around the use of long-established open-source datasets, and which impose narrow parameters that don't necessarily reflect the heterogeneity of the emerging market for facial recognition, object recognition and image segmentation.

Market Health

In any given year, the health of the global facial recognition market is assessed by final analysis of its growth in the previous 12-18 months. Though the usual annual premium market reports are available, they have a reduced value in this strangest of all years. All we can know for certain is that the market was in a period of notable and consistent growth up to the advent of the novel coronavirus.

We can only conjecture how the sector would have negotiated the public and corporate dissent of 2020 if COVID-19 had not intervened and put it in a holding position.

Of all the 'real-world' industry sectors currently treading water during the wait for a vaccine, the future of commercial facial recognition most depends on a return to some semblance of normal life. While the masks remain on, it is, for the most part, a problem without a solution.

Regulatory Outlook

With the promise of new regulation from the EU²⁷ and the US²⁸, the overturning of a landmark case in the UK²⁹, and a growing criticism of the proliferation of facial recognition in Asia³⁰, the sector will clearly be forced to implement reforms around privacy in the next few years.

Prioritizing governance in line with new regulations may slow down the facial recognition sector a little, but may also provide a stable market environment in which [data science for social good](#) can thrive, and where common consent has at least been addressed, if not satisfied.

As with the financial and banking sector, storing and transforming facial recognition data will be a trivial task compared to issues around acquisition, management, governance, regulatory compliance and negotiating transparency in a field where both the subjects and the watchers jealously guard their privacy.