EXPERIENCE-BASED RECOMMENDATIONS FOR FACE RECOGNITION

Face recognition is one of the most powerful and misunderstood technologies in modern times. It can help determine or help verify identity and then link that identity to specific times, locations, activities, and other persons present. In the hands of trained and experienced experts it is properly and successfully used every day to make us safer. In examples such as fighting child pornography, screening at major security events, and enabling the integrity of international border crossings, face recognition has been a successful tool in identifying suspicious actors intent on harming others or breaking laws. Simpler versions are also being used for personal electronics, access control, and time/attendance applications.

Well-trained and experienced users with nefarious intents could also use the technology for improper purposes however, such as general citizen surveillance. The bigger problem in the U.S. has been untrained and inexperienced entities who attempt to: (1) use the technology for valid purposes, but end up confusing their users and producing inconsistent results; and (2) design and promulgate policies and legislation that limit its use in positive applications or accidentally make it more difficult to ensure appropriate use.

Like many forms of intelligent systems, it is a tool that can be used for good or for ill, and it is incumbent on each user to leverage that tool correctly.

Staff of the MITRE Corporation have studied biometric face recognition algorithms, tested and evaluated systems, and guided federal agencies as they installed and operated systems that use it. This experiential knowledge is augmented with (a) robust privacy teams that work to ensure privacy protection in the engineering and federal policy domains and (b) a former Assistant Director of the White House Office of Science and Technology Policy who led government-wide policy on identity technologies during the Bush and Obama administrations. The combination of these in-depth and unique experiences inform the analysis and recommendations in this paper.

BIOMETRICS

is a general term used alternatively to describe a characteristic or a process:

As a Characteristic:
A measurable biological (anatomical and physiological) and behavioral characteristic that can be used for automated recognition

As a Process:
Automated methods of recognizing an individual based on measurable biological (anatomical and physiological) and behavioral characteristics.

FACE RECOGNITION

is a biometric modality that uses an image of the visible physical structure of an individual's face for recognition purposes.

Taken from *Biometrics Glossary*, a 2006 publication of the National Science and Technology Council.

MITRE

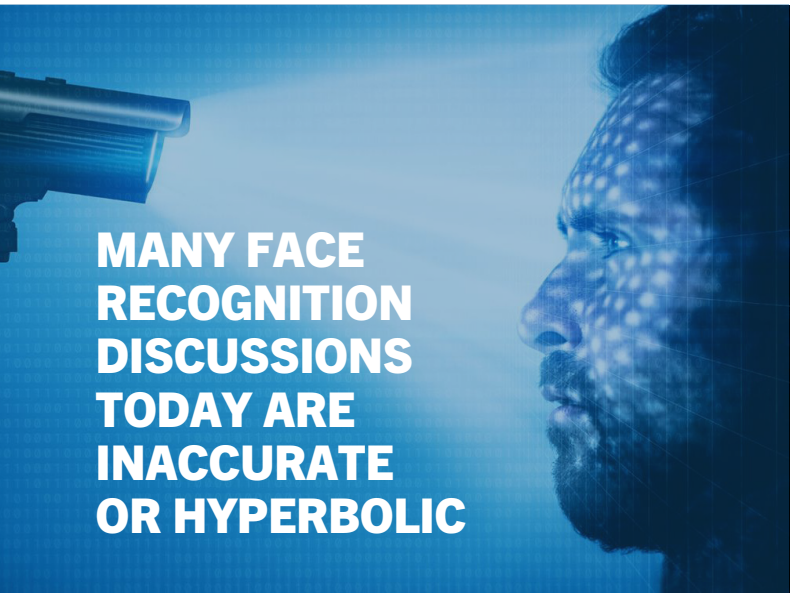

The Face Recognition Literacy Gap Continues to be a Significant Issue

Best practices, policies, and oversight are needed to ensure the ultimate objective that we all desire: enabling appropriate applications with equitable outcomes while protecting civil rights and civil liberties, and an absence of nefariously- or improperly-used applications. These ultimate outcomes are achievable, in the current day. Unfortunately, a significant literacy gap about the technology is keeping us from meeting this objective.

Most news articles and opinion pieces about face recognition that are available today contain inaccurate or biased information, and many include hyperbolic assumptions about outcomes that are disconnected from reality. Consider:

- Most references to NIST's recent FRVT[1] demographic report[2] in news articles and opinion pieces have stated that it proved that all face recognition systems are ethnically and gender biased and would definitely lead to false arrests of historically-impacted populations. The report itself, as well as NIST's subsequent briefings and Congressional testimony, state that such generalizations should not be made (as they'd often be incorrect). Relatedly, results from some other tests that help promote preferred policy outcomes are being widely championed, even though they don't meet basic statistical significance requirements and an expert could easily design similarly-scaled tests that produce either worse or opposite results.

- There are multiple other technologies that also use face images and advanced software analytics but are attempting to perform tasks that are different from biometric face recognition, such as estimating age or gender or identifying medical conditions. Yet the accuracy and issues within these other technologies are commonly (and usually incorrectly) cited as issues with biometric face recognition.

- Poorly conceived and managed implementation of face recognition by novice operators with limited training on the technology or privacy safeguards, or an individual operator's disregard for policies already in place, are often highlighted as being emblematic of the entire community.

These inaccurate messages are unfortunately serving as the foundation for many of the deliberations our nation is currently undertaking. Contrast these messages with the accurate insights below:

INSIGHT #1: FOCUS ON THE ENTIRE OPERATIONAL SYSTEM, NOT JUST THE ALGORITHM SUBCOMPONENT. Face recognition algorithms are one component of a face recognition system, which is usually itself one component within a complex human-machine system of systems in operational contexts. The face recognition algorithm plays a key, but not overwhelming, role in determining the decisions and eventual outcomes of the complex system; thus, equating the algorithm to the overarching system is an inaccurate oversimplification. It is often the case that system components and steps taken prior to and after[3] the face recognition algorithm often have a greater influence over the usability and accuracy of the overall system than the algorithm itself. (Note: NIST's FRVT evaluations assess the accuracy of face recognition algorithms only.)

INSIGHT #2: UNDERSTAND THE SIMILARITIES AND DIFFERENCES BETWEEN BIOMETRIC FACE RECOGNITION AND FACIAL ANALYTICS ALGORITHMS. Fundamentally, biometric face recognition algorithms compare an image to one or more images and produce a "similarity score" that shows how similar the identity-based attributes of the faces in the images are to each other. These algorithms have not explicitly attempted to determine the gender, age, or ethnicity, or recognize the facial expression – these latter capabilities are best classified as facial analytics algorithms, and they do not attempt to determine the identity of the individuals. There are growing instances where biometric face recognition algorithms and facial analytics algorithms are being leveraged together, or even closely combined, to help increase overall system accuracy. (Note: if studying ethics within Artificial Intelligence research, it can be useful to group these all together into a "face" bucket. Simply equating them within a policy deliberation on the appropriate use and associated issues of biometric face recognition is usually misleading, however.)

INSIGHT #3: FACE RECOGNITION IS INHERENTLY PROBABILISTIC, AND NATURE HAS AN IMPACT AS WELL. Face recognition algorithms will likely never be 100% accurate as there is no evidence that our faces are sufficiently unique, and it is

mathematically impossible to prove (or test to) 100% accuracy[4] . There are also naturally-occuring demographic variances that make it more difficult, but potentially not impossible[5], to ensure similar algorithm accuracy across different demographic groups. While this can understandably make many uncomfortable, policymakers' focus really needs to be on ensuring accurate and equitable outcomes from the total operational system rather than the algorithm subcomponent only. (Note: We use systems every day that have subcomponents that do not perform properly 100% of the time, because other aspects of the system account for it.)

INSIGHT #4: USE CASES VARY - DRAMATICALLY. Policy deliberations on face recognition must be extremely specific to a use case. Issues such as error rates, demographic performance, privacy and civil liberties, ethical appropriateness, and management and oversight requirements will vary substantially across different use cases. General investigations or improper mixing of multiple use cases will undoubtedly lead to incorrect assessments and decisions. (Note: As an example, considerations for using face recognition to control access to a secure facility are vastly different than those attempting to identify individuals in open, public spaces.)

ANALYSIS & RECOMMENDATION: Recent legislation has called for creation of a bipartisan group to study and develop recommendations for ensuring appropriate use and oversight of face recognition. That can be a useful approach, but only if the members of this group have the requisite technical training and personal experience with face recognition technology. Absent this fundamental requirement, the group will be unable to properly guide Congress, the Executive Branch, and the nation on the technology, associated issues, and operational considerations. This group should also not be limited to deep technical experts. Operational, legal, social justice, and policy experts, again with sufficient training and personal experience with this technology, are equally needed to ensure necessary safeguards are embedded within the system infrastructure.

A Framework for Proper Policy Deliberations

Generic, high-level conversations about face recognition do not accurately present the complexities associated with use of the technology and produce little benefit. Given its nature, face recognition will never exist without legitimate associated concerns. The technology itself is neither inherently evil nor out of control that an outright ban is the only, or even a wise, solution. Instead, discussions and deliberations must be specific and nuanced to be credible and beneficial to the challenges they are attempting to address.

A framework to enable these discussions and deliberations does not currently exist but can be rapidly developed. Insight #4 (above) tells us that guidelines for the use of this technology must be specific to individual use cases. While this is true, a framework must further break down the discussions into manageable components of the problem space. The high-level concept of face recognition may be easy for most people to inherently understand, as we each mentally perform that task as a part of our daily lives[6]. Computer-enabled face recognition, on the other hand, is extremely complex and difficult to understand without deep knowledge of the technology and systems design concepts. A notional framework to enable these beneficial discussions is provided below:

	Research	Testing	Initial Operational Consideration	System Design & TTPs	Operations & Maintenance	Termination
Data	●	●	●	●	●	●
Privacy and Civil Liberties Risks	●	●	●	●	●	●
Safeguards & Oversight	●	●	●	●	●	●
Reporting	●	●	●	●	●	●

Each gray cell would open to provide a detailed analysis & guidance on how to reason about that specific aspect of the issue. Components would include:

- An introductory overview that provides fundamental knowledge on the sub-topic & known issues;
- Existing best practices and standards; and
- A list of questions for potential operators and policymakers to consider

Variances by use cases would need to be included within certain cell's discussion or added as a third dimension.

ANALYSIS AND RECOMMENDATION: A completed framework along these lines would mature the policy and legislative deliberations surrounding face recognition, while also providing learned guidance to researchers and potential users of the technology. Both will help lead to proper policies that ensure face recognition is used in a proper, equitable manner going forward.

While a public-private collaboration has begun developing an initial draft of such a concept, the framework will need to become a community-wide activity that is continually enhanced with the latest insights – much like the existing MITRE ATT&CK framework[7] for cybersecurity.

Recommended Interim Congressional Action

It is clear that many in Congress want to do *something* on the face recognition front in the near-term, which is admirable. However, this action must be balanced with the realizations that (a) doing so properly will take time and effort, and (b) bans and moratoriums would likely create more harm than good, as well as be more difficult to overturn in the future. An interim step to consider would be to issue a resolution expressing the sense of Congress on concerns and expectations with the technology.

- In trained hands, biometric face recognition is beneficial in a variety of applications. In untrained hands (or trained hands of those with nefarious intent), use of the technology can potentially lead to undesirable outcomes. Additional technical and ethical training, guidelines, policies, and perhaps legislation will likely be required to help ensure wise and properly managed usage of the technology

- Face recognition is a complicated technology, with error rates, data, policy, and privacy concerns that vary significantly by use case. Guidelines, policies, and legislation must therefore be specific to individual use cases and system elements for them to deliver intended outcomes.

- Recognizing that face recognition will likely never be 100% accurate, well-conceived requirements and oversight will be required to ensure that results from face recognition algorithms are leveraged within proper contexts and the overall system's results are both proper and provide equitable outcomes across demographic variances.

- Our nation's privacy construct, which is decades old, has been overtaken by technological innovation unforeseen at the time of its construction. While this is manifesting in current conversations about face recognition technology, it also applies for many other technologies and issues. The nation needs to perform a fundamental assessment of privacy considerations, in both government and commercial applications, so that a modern construct can be developed.

- The use of face recognition to identify individuals as part of targeted surveillance or in support of law enforcement investigations should continue to have a "human in the loop" review of the output of face recognition algorithms by trained face examiners before they are used in decision making by other system components.

By issuing such a resolution Congress would be signaling that more detailed regulations and oversight are forthcoming, thus creating a near-term deterrent to unwise or unlearned implementations while it continues to study this issue and craft legislation that is evidence- and outcome-based, actionable, equitable, and measurable.

For information about Experience-Based Recommendations for Face Recognition, contact Duane Blackburn, dblackburn@mitre.org.

[1] NIST has evaluated performance of face recognition matching algorithms through the Face Recognition Vendor Test program for the past 20 years.
[2] https://nvlpubs.nist.gov/nistpubs/ir/2019/NIST.IR.8280.pdf
[3] Including a "human in the loop"
[4] There is no common definition of "accuracy" as that measure will vary by use case. We should be discussing these issues in terms of error rates, such as a false reject rate at a given false accept rate.
[5] NIST's 2020 FRVT report on demographic variances did not detect measurable variances for the top-performing algorithms.
[6] Note: studies have shown that humans tend to perform this task well on individuals that we know but are poor at recognizing faces of unfamiliar individuals.
[7] https://attack.mitre.org/

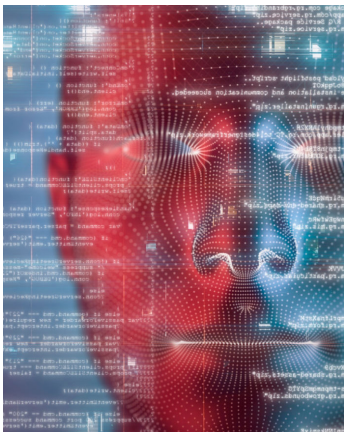

Legislation or regulation of face recognition technology needs to be evidence- and outcome-based, actionable, equitable, and measurable.

MITRE's mission-driven teams are dedicated to solving problems for a safer world. Through our public-private partnerships and federally funded R&D centers, we work across government to tackle challenges to the safety, stability, and well-being of our nation.

Our objective is to help inform policy development by providing evidence- and data-driven information on impacts, implications, and implementation considerations. These contributions will drive development of impactful and actionable policy that helps to address significant national issues.

MITRE | SOLVING PROBLEMS FOR A SAFER WORLD®