NATIONAL ACADEMY OF SCIENCES
NATIONAL ACADEMY OF ENGINEERING
INSTITUTE OF MEDICINE
NATIONAL RESEARCH COUNCIL

REPORT IN BRIEF

BIOMETRIC RECOGNITION:
CHALLENGES AND OPPORTUNITIES

COMPUTER SCIENCE AND TELECOMMUNICATIONS BOARD
SEPTEMBER 2010

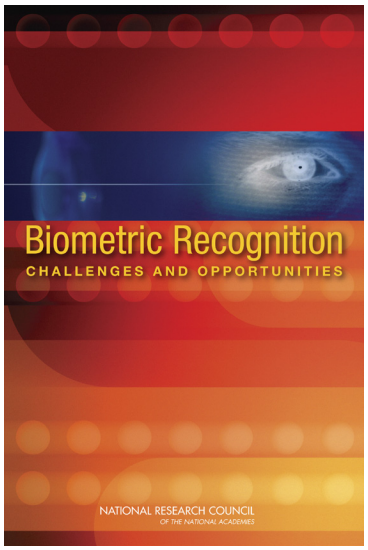

Download This Report

Biometric recognition is the automated recognition of individuals based on voice, fingerprints, or other biological and behavioral characteristics. Biometric identification systems are currently used to better control access to facilities and financial accounts, identify criminals, track medical patients, and control access to social services, among other applications, and they are likely to become more widespread in the future. Although useful in many circumstances, more research is needed to gain a complete knowledge of their strengths and limitations, particularly in very large-scale applications. This report assesses current biometric technologies and explores the technical and policy challenges associated with the development, evaluation and use of systems that employ biometric tools. It seeks to correct the common misperception that biometrics can be used to identify individuals with absolute certainty, and that rates of error can always be reduced to insignificant levels (or even known, in some cases). Finally, the report presents five principles intended to guide the appropriate and effective use of biometrics in the future. In particular, there is a need to understand the variables that contribute to error in systems, to apply scientific principles to the research of biometric recognition mechanisms, and to take a systems-level approach to the design and evaluation of biometric tools.

Most public understanding of biometric identification comes from television and film, where futuristic high-tech instruments scan a trait and then definitively match it to a single corresponding identity. In practice, the process of measuring and analyzing human traits is far more complex, and less certain, than it appears on the screen. Unlike password-based systems, which rely on input that can only be entered in one way, biometric systems are probabilistic – meaning that they assess the likelihood that each unique recording of a trait or behavior belongs to the same user as previously recorded references in a database. If the recording is sufficiently similar to one already stored, the user is considered a "match." If the differences are too great, that person is deemed a "nonmatch."

Uncertainty in Biometric Recognition

Uncertainty is inherent in each step of biometric identification. Readings of traits or actions such as face, voice, or gait may vary depending on the angle at which they are captured, the frame of mind of the subject, and other factors that are difficult or impossible to control. Some of them may change in the same individual over time. Sensor age and calibration, how well the interface mitigates extraneous factors, and the level of sensitivity to changes in light levels and other environmental variations can also affect performance. Biometric characteristics cannot be compared directly; stable and distinctive "features" must be extracted from them for comparison. Algorithms used for feature extraction can differ, contributing to performance variations. Data itself is also prone to degradation resulting from legitimate manipulation as well as corruption, mismanagement, or user error. Understanding these variables and other factors that affect system performance can drastically reduce the incidence of error, but they cannot eliminate it altogether.

A fundamental source of uncertainty that merits further investigation are the characteristics of biometric traits themselves. Little is known about the stability or distinctiveness of traits across individuals and groups. Furthermore, these traits are generally only observed through the filters imposed by measurement processes and feature extraction, making it difficult to undertake a comprehensive study of their qualities. Thus, the development of a science of human individual distinctiveness, enabled by careful data collection and analysis, is essential to the effective use of biometric recognition, especially at scale.

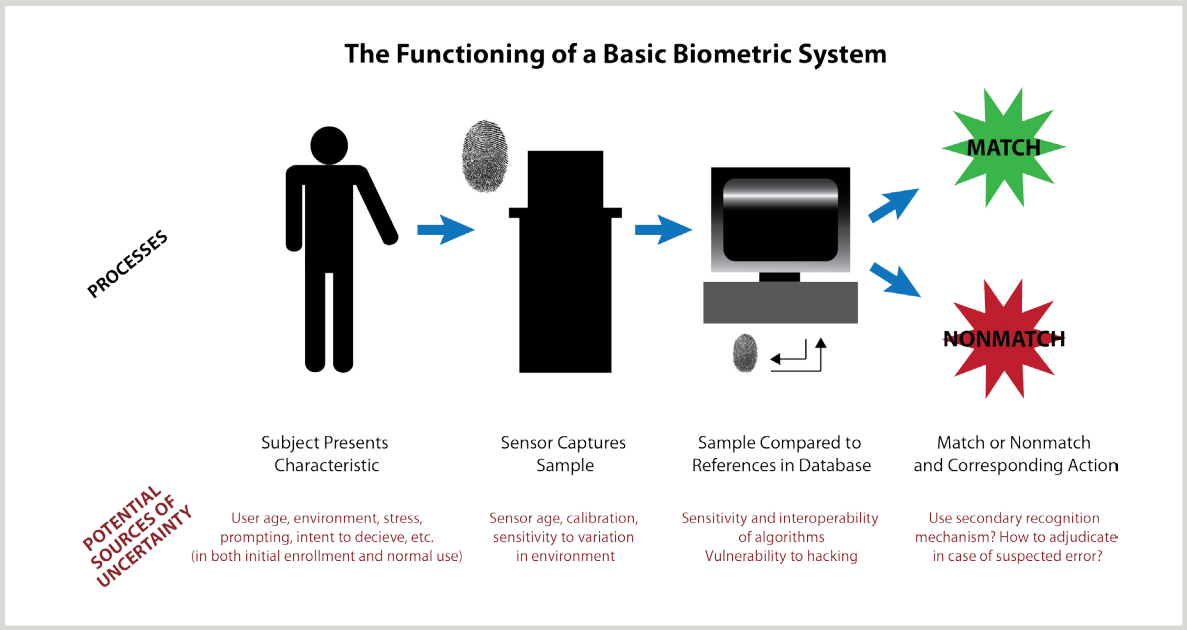

Rates of Error

Error rates in biometric identification systems are frequently misunderstood. Although a system with a false match rate (FMR) of 0.1% and false nonmatch rate (FNMR) of 0.1% might seem to have an inconsequential probability of failure, the percentage of "correct" decisions actually depends on the percentage of impostors expected in the population of users of the system, not just on the error rates of the technology.

Even a biometric system with a very low stated rate of false matches or false nonmatches can have a significant number of incorrect matches when used on a large scale. Moreover, because in most cases the number of impostors encountering a system in a real-world scenario cannot be known, it may be impossible to accurately specify the expected percentage of incorrect results, making it difficult to know how much confidence to place in a "nonmatch" result. The common assumption that a false match or false nonmatch rate of 0.1% means one can have great confidence in system results is not only incorrect, but could have dangerous legal or social consequences if, for example, biometric measurements were used as evidence in a criminal case without properly contextualizing the results.

Principle I: Users and developers of biometric systems should recognize and take into account the limitations and constraints of biometric systems—especially the probabilistic nature of the underlying science, the current limits of knowledge regarding human individual distinctiveness, and the numerous sources of uncertainty in biometric systems.

Are Error Rates a Good Measure of Reliability?

It seems intuitively obvious that a declared nonmatch in a biometric system with both FMRs and FNMRs of 0.1 percent is almost certainly correct. Unfortunately, intuition is grossly misleading. The following series of examples illustrates how the expected percentage of "right" decisions by a biometric system depends upon the percentage of impostors that the system actually encounters, not just the error rate of the technology.

Imagine that we have installed a biometric verification system to control entry to a college dormitory. The system has a 0.1 percent false match rate (FMR) and a 0.1 per cent false nonmatch rate (FNMR). How often does a nonmatch represent an attempt by a nonresident "impostor" to get into the dorm? The answer, it turns out, is "it depends."

In Scenario 1, the impostor base rate is 0 percent—that is, no impostors ever try to get into the dorm. In this case, all of the people using the biometric system are residents. Since the system has a 0.1 percent FNMR, it will generate a false nonmatch once every 1,000 authentication attempts. All of these nonmatches will be errors (because in this case all the people using the system are residents).

In Scenario 2, one nonresident impostor tries to get into the building for every 999 times a resident attempts entry. In this case, the system generates one false nonmatch for each 1,000 recognition attempts *and* it generates a nonmatch for the one nonresident impostor. On the average, therefore, every 1,000 recognition attempts will include one impostor (who will likely generate a correct nonmatch) and one resident who will generate an incorrect nonmatch. Of the two nonmatches, 50 percent of them will be correct and 50 percent of them will be incorrect.

Scenario 1
Impostor rate 0%; low nonmatch accuracy

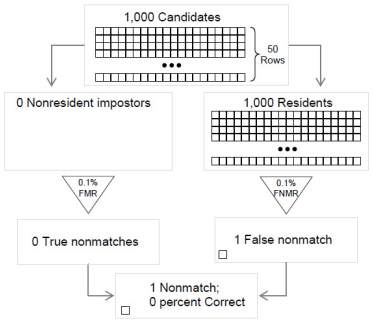

Scenario 2
Impostor rate 0.1%; moderate nonmatch accuracy

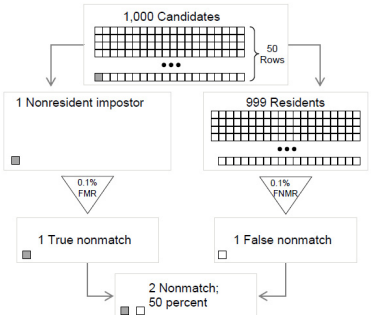

Scenarios 3 and 4 calculate confidence in the truth of a nonmatch in cases where 1 percent of the people trying to get into the dorm are nonresident impostors, and in cases where half the people trying to get into the dorm are nonresident impostors. Note that confidence in the correctness of a nonmatch approaches 99.9 percent (the true nonmatch rate of the system) only when at least half the people trying to get into the dorm are impostors!

In fact, FMR and FNMR alone are not accurate measures of how often the system gives the right answer in an operational environment. In many cases, they will greatly overstate the confidence we should have in the system.

Scenario 3
Impostor rate 1.0%; very high nonmatch accuracy

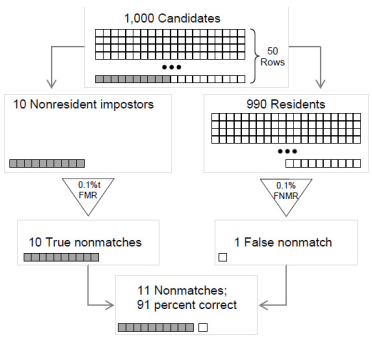

Scenario 4
Impostor rate 50%; very high nonmatch accuracy

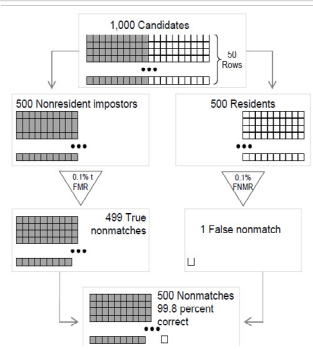

Biological characteristics and behaviors are revealed in daily life, and they cannot be easily replaced like passwords. This creates unique challenges for the design and use of large-scale biometrics systems.

The biometrics industry currently lacks well-defined best practices based on a body of solid peer-reviewed scientific research. As concerns about security, confidentiality of proprietary information, and fraud in general drive the adoption of biometrics as a routine method of recognizing individuals, it is increasingly important that the development of systems be based on a thorough understanding of the components of biometric systems and the contexts in which they are used. Basic research should be done on the distinctiveness of various biometric traits or behaviors, their stability over time, and their variability among various demographics. Research on user interactions with systems is also critical, as is inquiry into the social, legal, and cultural frameworks in which biometric systems are embedded.

As with the deployment of any security system, it is important to predict possible security threats. The aim of a threat assessment is not only to determine feasibility of threats against the resource being protected, but also against the system doing the protecting. This is a matter of particular importance for biometric systems because biometric traits, unlike tokens or passwords, cannot be easily replaced. If the same trait is used by different systems, weaknesses in one system could compromise that trait for use in all other systems. Furthermore, our biometric traits are not secret – we reveal them throughout the course of everyday life. We leave fingerprints on surfaces, faces can be photographed, and voices can be recorded. A threat assessment must also try to predict the motivations and capabilities of three distinct types of users: clients (who should be recognized), imposters (who should not be recognized but imitate those who should), and identity concealers (who should be recognized but attempt not to be), in the context in which the system will be used. All of these factors need to be carefully considered when evaluating the merits and risks of a biometric system in comparison with other security systems.

Principle II: Efforts to determine best practices for testing and evaluating existing and new biometric systems should be sustained and expanded. Careful consideration should be given to making the testing process open, allowing assessment of results and quality measures by outside parties when appropriate. The evaluation of a system's effectiveness needs to take into account the purpose for which the system was developed and how well field conditions were matched.

Principle III: Best practices are needed for the design and development of biometric systems and the processes for their operation. To scale efficiently to mass applications, these best practices should include requirements for system usability, initial and sustained technical accuracy and system performance, appropriate exception handling, and consistency of adjudication at the system level. Best practices should allow for incorporation of scientific advances and be auditable throughout the life of the system.

Contextual Factors That Impact System Performance

Environments and applications vary in ways that can affect the performance of a system. Just a few of the factors that should be considered during system design and analysis are listed here:

User context:

- What are the users' motivations for using the system?
- Are users aware of their interactions with the system?
- Do users need to be trained to use the system?
- Does user habituation affect results?

Application context:

- Is the system under live supervision?
- Is the goal to recognize users that match the database (membership) or users that do not match (watch list)?
- Is the user population an open group (e.g. the general public) or closed group (e.g. residents of a building)?
- Does testing the claim require one comparison, or many?

Technology context:

- Is the environment in which the trait is recorded controlled for consistency (e.g. lighting)
- Is user input passive or active? Is the system covert or overt?
- How quickly do users need to be processed?
- What are the required bounds on the error rates?
- What other systems will be networked or otherwise integrated with the system and what impact will they have on its vulnerability?

Far from being a simple process, a single biometric recognition occurs through a series of automated and human decisions which can interact in complex ways to affect the final outcome. The larger technological, operational, and social contexts within which each system operates must be understood and accommodated if that system is to be useful and robust. Analyses of biometric systems should take a broad systems perspective that incorporates all of these elements.

Because biometric systems are used in a multitude of contexts and for a variety of reasons, their mechanisms and components also vary. The design and analysis of a biometric system must not only take an integrated systems approach, but account for the context in which it will be used.

Because these contextual factors are subject to variation based on the nature of the problem being solved, it is essential that the problem be defined precisely and that appropriate requirements are selected for each system. A one-size-fits-all approach is not appropriate.

Principle IV: Requirements have critical implications for the design and development of human recognition systems and whether and how biometric technologies are appropriately employed. Requirements for systems can vary widely, and assessment and evaluation of the effectiveness of a given system need to take into account the problem and context it was intended to address.

Social, Cultural, and Legal Considerations

The connection between biometric traits and identity records raises a host of social, cultural, and legal concerns. For example, in contexts where individuals are claiming enrollment or entitlement to a benefit, the inability or refusal to use a system on a physical or cultural basis could result in disenfranchisement. Such cases illustrate the need for careful oversight and alternative methods of identification in order to minimize negative social consequences and avoid violating individuals' privacy or due process rights.

The success of large-scale or public biometric systems is dependent on gaining broad public acceptance of their validity. To achieve this goal, the risks and benefits of using such a system must be clearly presented. Public fears about using the system, including fear of stigma or punishment for refusing to use the system, and concerns about theft or misuse of information, should be addressed. Covert systems, designed to track and identify individuals without their knowledge, are also a cause of deep public concern. The issues of system alternatives, adjudication, authority, and privacy should be addressed by legal specialists and other experts in appropriate fields.

Principle V: Social, legal, and cultural factors can affect the acceptance and effectiveness of biometric systems and should be taken into account in system design, development, and deployment. Notions of proof related to biometric recognition should be based on solid, peer-reviewed studies of system accuracy under many conditions and for many persons reflecting real-world sources of error and uncertainty in those mechanisms. Pending scientific consensus on the reliability of biometric recognition mechanisms, a reasonable level of uncertainty should be acknowledged for biometric recognition. There may be a need for legislation to protect against the theft or fraudulent use of biometric systems and data.

Understanding the science and technology that form the basis of large-scale biometric systems will be critical if they are to be used in the future for addressing important public security needs. At present, many questions remain to be answered. Although biometric systems perform well in many existing applications, biometric capabilities and limitations are not well understood in very large scale applications involving tens of millions of users. The distinctiveness and stability of traits under a variety of conditions and within large populations is a subject that merits further research, and will require extensive collection of personal data. Sensors should be developed for maximum affordability, reliability and accuracy. A study of behavioral characteristics within user populations, including the possible behaviors of potential system "hackers" is needed. The results of this and other related research should be published in open, peer-reviewed scientific literature and used as the basis for industry-wide best practices.

Principle VI: As biometric recognition is deployed in systems of national importance, additional research is needed at virtually all levels of the system (including sensors, data management, human factors, and testing). The research should look at a range of questions from the distinctiveness of biometric traits to optimal ways of evaluating and maintaining large systems over many years.

Whither Biometrics Committee: Joseph N. Pato, Chair, Hewlett-Packard Company; **Bob Blakley**, Gartner; **Jeanette Blomberg**, IBM Almaden Research Center; **Joseph P. Campbell,** Massachusetts Institute of Technology, Lincoln Laboratory; **George T. Duncan**, Carnegie Mellon University; **George R. Fisher**, Prudential-Wachovia (Retired); **Steven P. Goldberg***, Georgetown University Law Center; **Peter T. Higgins**, Higgins & Associates, International; **Peter B. Imrey**, Cleveland Clinic and Case Western Reserve University; **Anil K. Jain**, Michigan State University; **Gordon Levin**, The Walt Disney World Company; **Lawrence D. Nadel**, Noblis; **James L. Wayman**, San Jose State University
**Deceased*

Staff: Lynette I. Millett, Senior Program Officer

Support for this project was provided by the Defense Advanced Research Projects Agency (Award No. N00174-03-C-0074) and by the Central Intelligence Agency and the Department of Homeland Security with assistance from the National Science Foundation (Award No. IIS-0344584). Any opinions expressed in this material are those of the authors and do not necessarily reflect the views of the agencies and organizations that provided support for the project. This report brief was prepared by the National Research Council based on the committee's report. More information can be obtained by contacting the Computer Science and Telecommunications Board (http://cstb.org).

Copies of the full report can be purchased from the National Academies Press, 500 5th Street NW, Washington DC, 20001; (800) 624-6242; www.nap.edu