RAND CORPORATION

Research Report

SAMUEL PETERSON, BRIAN A. JACKSON, DULANI WOODS, ELINA TREYGER

Finding a Broadly Practical Approach for Regulating the Use of Facial Recognition by Law Enforcement

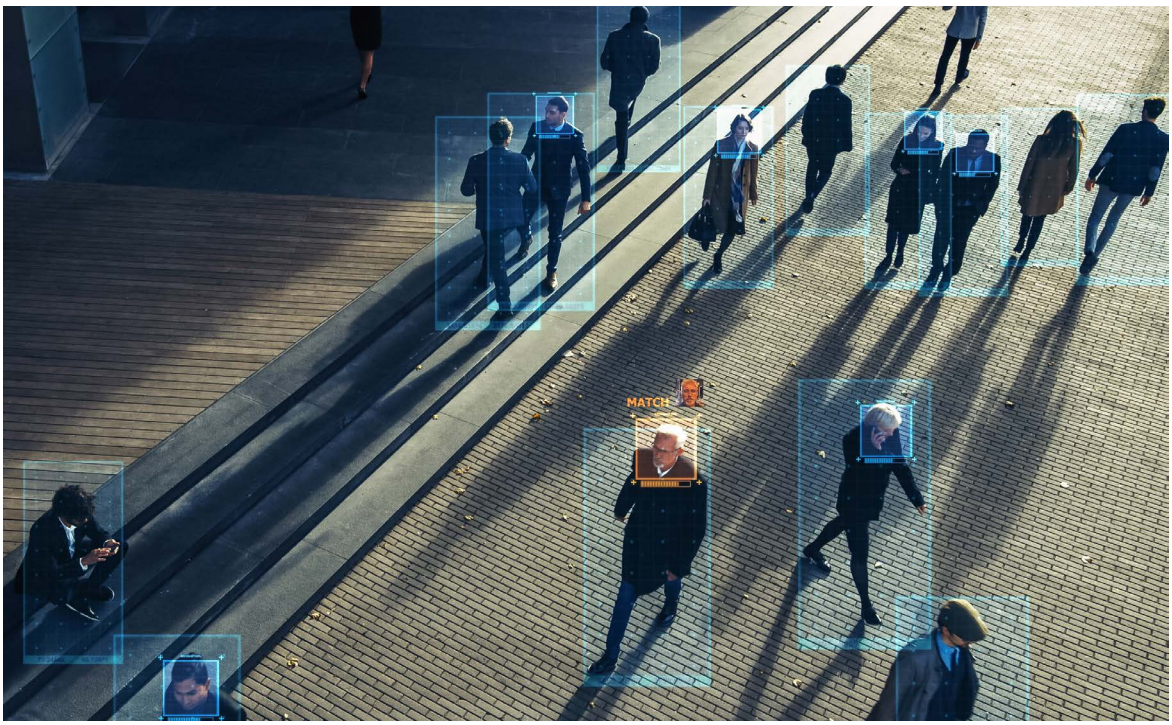

For more information on this publication, visit **www.rand.org/t/RRA2249-1**.

About RAND

The RAND Corporation is a research organization that develops solutions to public policy challenges to help make communities throughout the world safer and more secure, healthier and more prosperous. RAND is nonprofit, nonpartisan, and committed to the public interest. To learn more about RAND, visit www.rand.org.

Research Integrity

Our mission to help improve policy and decisionmaking through research and analysis is enabled through our core values of quality and objectivity and our unwavering commitment to the highest level of integrity and ethical behavior. To help ensure our research and analysis are rigorous, objective, and nonpartisan, we subject our research publications to a robust and exacting quality-assurance process; avoid both the appearance and reality of financial and other conflicts of interest through staff training, project screening, and a policy of mandatory disclosure; and pursue transparency in our research engagements through our commitment to the open publication of our research findings and recommendations, disclosure of the source of funding of published research, and policies to ensure intellectual independence. For more information, visit www.rand.org/about/research-integrity.

RAND's publications do not necessarily reflect the opinions of its research clients and sponsors.

Published by the RAND Corporation, Santa Monica, Calif.
© 2023 RAND Corporation
RAND® is a registered trademark.

Cover image: Gorodenkoff / Adobe Stock.

Limited Print and Electronic Distribution Rights

This publication and trademark(s) contained herein are protected by law. This representation of RAND intellectual property is provided for noncommercial use only. Unauthorized posting of this publication online is prohibited; linking directly to its webpage on rand.org is encouraged. Permission is required from RAND to reproduce, or reuse in another form, any of its research products for commercial purposes. For information on reprint and reuse permissions, please visit www.rand.org/pubs/permissions.

About This Report

Technological advances in society often have implications for policing. Specifically, much as fingerprinting and DNA evidence were groundbreaking advances that have been used to solve serious crimes, remote biometric identification, such as facial recognition and object identification, is rapidly developing and being used in policing. This has significant implications for legislators and policymakers, criminal justice officials, and society in general that should be understood, and further developments need to include efforts to maximize the benefits of such technology while limiting the potential risks. This report examines current and near-future uses of facial recognition by law enforcement and identifies the various concerns and related policy domains that could be used to guide public policy. This report will be useful for policymakers, criminal justice officials, and interested members of the public.

Justice Policy Program

RAND Social and Economic Well-Being is a division of the RAND Corporation that seeks to actively improve the health and social and economic well-being of populations and communities throughout the world. This research was conducted in the Justice Policy Program within RAND Social and Economic Well-Being. The program focuses on such topics as access to justice, policing, corrections, drug policy, and court system reform, as well as other policy concerns pertaining to public safety and criminal and civil justice. For more information, email justicepolicy@rand.org.

Funding

Funding for this research was provided by gifts from RAND supporters and income from operations.

Acknowledgments

We would like to thank RAND for supporting this project, as well as the RAND-Initiated Research leadership, including Vice President of Innovation Nancy Staudt and Director of RAND-Initiated Research Lisa Jaycox. This project benefited greatly from the insights of the subject-matter experts we interviewed. They include (along with others who declined to be named) Armando R. Aguilar (assistant chief of police, Miami Police Department), Albert Fox Cahn, Esq. (executive director, Surveillance Technology Oversight Project, S.T.O.P.), Joseph Courtesis (inspector, New York Police Department, ret.), Kade Crockford (director of the Technology for Liberty Program, American Civil Liberties Union of Massachusetts), Pete Fussey (codirector of the Human Rights, Big Data and Technology Project, University of Essex), Jerome D. Greco (digital forensics supervising attorney, Legal Aid Society), Patrick Grother (National Institute of Standards and Technology), Bob Gualtieri (sheriff, Pinellas County, Florida), Daniel Heltemes (sergeant, Arizona Department of Public Safety), Clark Kimerer (assistant chief, Seattle Police Department, ret.), Dave Maass (director of investigations, Electronic Frontier Foundation), Diana Maurer (director, U.S. Government Accountability Office), David Ray (chief operating officer and general counsel, Rank One Computing Corporation), Christian Quinn (retired law enforcement executive, managing principal, Fulcrum Innovation LLC), Paul Steinberg (senior vice president—technology, Motorola Solutions), Hoan Ton-That (founder and chief executive

officer, Clearview AI). We would also like to thank our peer reviewers, Kristin Warren from RAND and Katie Kinsey from the Policing Project, who provided valuable feedback that greatly improved this report. Lastly, we thank our RAND colleagues Sydne Newbury, John Hollywood, and Donell Harvin for their contributions to this project.

Summary

Communities across the United States are grappling with the implications of law enforcement organizations' and other government agencies' use of facial recognition (FR) technology. Although the purported benefits of FR as stated are clear, they have yet to be measured and weighed against the existing risks, which are also substantial. Given the variety of ways in which FR can be used by law enforcement, the full benefit-to-risk trade-off is difficult to account for, leading to some municipalities that ban the use of FR by law enforcement and others that have no clear regulations. This report provides an overview of what is known about FR use in law enforcement and provides a road map of sorts to help policymakers sort through the various risks and benefits relative to different types of FR use.

We categorize the various identified risks associated with FR technology and its use by law enforcement, including accuracy, bias, the scope of the search (i.e., surveillance versus investigation), data-sharing and storage practices, privacy issues, human and civil rights, officer misuse, law enforcement reactions to the FR results (e.g., street stops), public acceptance, and unintended consequences. The concerns are discussed in detail in Chapter 3, but they are summarized here.

Concerns About Whether Facial Recognition Use by Law Enforcement Is Justified Across a Variety of Situations

- Will police use FR for minor crimes?
- Will FR use for field identification lead to an unwanted expansion of government power (e.g., identify everyone at all times) and to potentially riskier interactions?
- Will FR use in noncriminal situations create issues with consent and error? Will the results be used against the subject of the scan (e.g., if they have an open warrant)?

Concerns About Whether There a Legitimate Law Enforcement Purpose for Scanning Someone

- Since FR is a powerful tool, will it be used to scan people in public for no legitimate law enforcement reason or track their movements?
- Will accuracy be reduced as the number of searches increases?
- Could more-expansive searches violate civil rights?

Concerns About How Accurate Facial Recognition Is in the Law Enforcement Setting

- Will a lack of common practices or standards affect understanding of accuracy at scale or of the impacts of different practices (e.g., image quality, image alterations, threshold settings)?
- Are training and gallery data high quality, and how is accuracy complicated by the true positive not being included in the gallery (testing always includes correct answer)?
- What is the additive or interactive effect of human and algorithm error on overall accuracy?

Concerns About Whether Facial Recognition Systems Are Biased Against Certain Demographic Groups

- Are training data reflective of the population?
- What are the effects of gallery data representativeness and demographic clustering?
- What are the effects of a higher algorithm error for certain demographic groups, and how do demographic error rates interact (e.g., older black females)?
- What is the disparate impact due to source of data (e.g., placement of public CCTV cameras) and use of supplemental information (e.g., criminal history) to guide FR search decisions by human operators?
- What are the effects of human bias for other races and unfamiliar faces?

Concerns About Whether the Underlying Data Are Flawed

- Are training data sufficiently representative of not only demographics but variation in image quality (both of which affect accuracy)?
- What are the issues surrounding privacy and consent of being included in gallery data?

Concerns About What Happens After Police Use Facial Recognition Software

- Will uncritical acceptance of FR results by officers lead to false arrests (i.e., lack of probable cause)?
- How are FR results handled with eyewitnesses, and does FR increase similarity enough to contribute to misidentifications?
- Will FR use ever require or prompt an officer to act immediately on FR results (e.g., a field stop)?

Concerns About Whether Facial Recognition Use Bypasses Civil Rights Protections

- Could FR be used to stifle free speech, religion, or assembly?
- Will FR use lead to unreasonable searches or violate expectations of anonymity (e.g., pervasive surveillance)?
- Does nondisclosure of FR impede the right to confront witnesses or to disclose potentially exculpatory evidence to the defense?
- What are the effects of the courts not having paid much attention to FR use so far and legislation lagging behind use?

Concerns About Whether Facial Recognition Use by Law Enforcement Violates People's Privacy

- What are the potential effects of ever-expanding FR databases with images that were not collected for any government purpose?
- What are the potential effects of more people being included in searches and search results being retained?
- What are the potential effects of data being kept indefinitely?

Concerns About Whether Facial Recognition Data Are Being Handled Appropriately by Law Enforcement

- What are the policies for oversight and data quality in FR data-sharing networks that include millions of images from multiple states and multiple sources?

Concerns About the Unintended Consequences That Might Come from Facial Recognition Use

- Could FR use change police investigations by overemphasizing video evidence or diminishing other investigative skills to solve crimes?
- Could the community view FR results as definitive or reduce reporting or cooperation because they assume video evidence exists?
- Could FR use reduce community engagement due to officer overreliance on FR or due to community mistrust?

Next, with this in mind, we outline a policy road map that identifies considerations to account for some of these concerns. Unlike other existing documents, this report does not make specific recommendations for policy, noting that the knowledge and research on FR use in law enforcement are likely to continue to change. Thus, these policy considerations provide guidance for key categories to be addressed by policy, although they may become more stringent or relaxed as the field evolves.

Contents

About This Report iii
Summary v
Figures and Tables xi

CHAPTER 1
Introduction: The Use of Facial Recognition by Law Enforcement 1
Identification in Law Enforcement 2
Current Issues in Use of Facial Recognition by Police: Regulation, Privacy, Civil Rights, Bias, and Public Perceptions 6
What Do We Know About Existing Policy and Practice? 10
The Current Study 11

CHAPTER 2
Developing a Spectrum of Use Cases as a Step Toward a Broadly Practical Approach for Regulation 13
What Do We Mean by Use (and Misuse) Cases? 13
Use and Misuse Across the Current and Proposed Implementation of Facial Recognition 14
Defining Use and Misuse Cases for the Interviews with Subject-Matter Experts 17
Subject-Matter-Expert Interviews 18

CHAPTER 3
Understanding the Landscape of Concerns 21
Thinking About Potential Benefits and Risks of Facial Recognition 21
Mapping the Landscape of Concerns 24
What Happens After the Facial Recognition Search? 34
Legal and Ethical Implications of Facial Recognition Use 35
Unintended Consequences? 41

CHAPTER 4
Developing a Policy Road Map to Help Navigate the Landscape of Concerns About Facial Recognition 43
Finding a Better Path 43
Building a Path Through the Landscape: Defining How Facial Recognition Systems Should Be Used by Police 44

CHAPTER 5
Conclusions 63
The Potential Ubiquity of Facial Recognition and Options for Narrowing 63
Mapping the Policy Road Map to Concerns 64
Preparing for the Future of Facial Recognition and Other Remote Biometric or Object Identification Technologies 66

APPENDIXES

A. Survey Methodology 67
B. Interview Protocol 69

Abbreviations 73
References 75

Figures and Tables

Figures

1.1. Facial Recognition Process and Key Features 4
1.2. Geometry-Based Facial Recognition 5
1.3. Face Recognition Process 6
1.4. Overview of Survey Responses by Trust in Police 8
2.1. Common Facial Recognition Use Steps 14
2.2. Categories of Potential Misuse 15
3.1. Three Groups with Perspectives on Facial Recognition 22
3.2. Factors in the Facial Recognition Search Process That May Affect Accuracy 30
4.1. Building a Road Map Through the Landscape of Concerns About Law Enforcement Use of Facial Recognition Systems 44
4.2. Relative Estimates for the Costs of Different Crime Types 46
4.3. Organizational Location of Approval and Performance of Facial Recognition Searches: Effects on Accountability Support 50

Tables

1.1. Overview of Survey Response by Trust in Police 9
1.2. Example State Facial Recognition Statutes 11
2.1. Hypothetical Use Cases for Interviews with Subject-Matter Experts 18
2.2. Policy Dimensions of Facial Recognition Use for Subject-Matter-Expert Interviews 19
3.1. Scope of Search and Typical Use 27
3.2. Potential Sources of Bias in Facial Recognition Search 32
5.1. Policy Road Map and Related Concerns 65

CHAPTER 1

Introduction: The Use of Facial Recognition by Law Enforcement

Facial recognition (FR) has progressed from a futuristic technology found only in science fiction literature to a capability that is commonly used in some domains. FR is now built into commercial products to protect data while increasing user convenience, being employed by companies to identify and track individuals for sales and marketing purposes, and proliferating online both by social media networks and by search engines that mine the vast quantities of photos and identifying information disseminated by people on the internet. Given the potential of FR to aid in the identification of individuals, law enforcement organizations have pursued FR to improve the efficiency and success of their crime prevention and investigation efforts by using it to identify people involved in criminal behavior.

Although a variety of uses of FR are becoming accepted more commonly (e.g., securing access to your smartphone), FR use by law enforcement remains a frequent target of controversy and continues to be regarded by some as a high-risk application of FR. Among the issues raised by civil society groups and others are concerns about the effect of FR on people's willingness to exercise constitutional rights (e.g., because of fears of state surveillance and the abuse of power), the potential for FR systems' errors to lead to miscarriages of justice, and analyses that have suggested that biases embedded in such systems could perpetuate or worsen inequities in criminal justice actions.

Developing effective and sustainable approaches for managing the use of new tools by government is never easy, especially when these tools carry concerns about the impact on society, privacy, and individual rights (American Civil Liberties Union of Massachusetts, undated).[1] The challenges grow in magnitude if time passes and technologies are implemented before policymakers address the need to develop policies and guidelines. Even though the application of FR in law enforcement is still in its infancy, efforts to develop regulations have struggled. Some organizations serve to assist local departments by developing model policies, but lack of consensus across the profession has led some of these organizations to defer developing resources or policies for using FR.[2] In some cases, concerns about FR have led to wholesale bans of police use of the technology (Sheard and Schwartz, 2022). In others, concerns about rising crime have made the potential benefits of the technology more compelling and led to reversals of restrictive legislation (Reuters, 2022).

[1] Although there is still an ongoing push for bans or moratoria on FR use, the American Civil Liberties Union of Massachusetts has highlighted possible legislation that would limit use of FR in many public places, centralize FR searches to state police, clarify due-process requirements, and give defendants notice and information about how FR was used in their case. See American Civil Liberties Union of Massachusetts, undated.

[2] In a publication by the International Association of Chiefs of Police, the organization noted that it had not yet developed a model policy for FR systems, "given some of the conflicting professional views on facial recognition programs" (International Association of Chiefs of Police, Law Enforcement Policy Center, 2022). Several years earlier (before George Floyd's death and subsequent national focus on U.S. policing), a group managed by the federal government (and including the International Association of Chiefs of Police) did publish a policy development template for localities (U.S. Department of Homeland Security and Bureau of Justice Assistance, U.S. Department of Justice, 2017).

This pendulum dynamic is fundamentally unproductive if the goal is to identify and implement appropriate and balanced use of new technological capabilities to serve society's broader goals. Swings from permissiveness to restriction risk wasting resources if technology acquisition efforts are disrupted or canceled; may create societal harms from lack of restriction on inappropriate use of these technologies, given uncertainty in management structures and oversight; and may sacrifice the opportunity for police—and government more broadly—to develop policies for responsible use of a potentially valuable tool.

Survey research from 2022 suggests about half of U.S. adults are comfortable with law enforcement using FR as long as constraints are applied, although some racial differences in favorability exist, including concerns about disparate impact and false arrests. And a large majority agree that FR results alone should not be grounds for arrest (Rainie et al., 2022).

Still, many communities across the country are currently considering whether FR should be used at all, which is a critical first step to policymaking. As an emerging area, a key question is whether FR technology can be used by law enforcement in a way that serves society's goals. Communities are still in the early phases of identifying FR's potential, optimizing its use to avoid errors, and developing policies about use that recognize current flaws. And it is not clear that policymakers have the information to know whether and how anticipated benefits or risks are being managed currently by organizations that use FR. More transparency about law enforcement use of FR (e.g., how often, for what crimes) and research to assess the costs and benefits of FR as a law enforcement tool is critical for informing this debate. Current accuracy testing being done by the National Institute of Standards and Technology (NIST) may be partially applicable to law enforcement settings, but this research happens in a more controlled environment. The real-world accuracy in law enforcement settings is not known and would help assess potential risks (e.g., situations in which the technology performs poorly, disparate impacts) (see, for example, Ho et al., 2020, p. 753). Further, research would inform whether and how much benefit is accrued through FR use (e.g., efficiency, case clearance) and help identify potential unintended consequences (e.g., decreasing clearance rates for crimes without video). This is critical information for the public in deciding whether to authorize FR use and for demonstrating the tangible benefits of FR use, as well as for policymakers seeking to regulate and minimize risks. In the near term, debates about whether law enforcement should use FR and what that should look like if approved will continue. As the state of knowledge about FR use in law enforcement continues to grow, there is a need for nuanced thinking about the meaning and options for responsible use, for better disaggregation of the concerns and potential harms regarding law enforcement's uses of FR, and for identifying options to address those concerns effectively.

This report is intended as a first step in accomplishing those goals. We examine, objectively, the various dimensions along which use of FR for law enforcement purposes can be considered. We also summarize the variety of policies and views about FR acceptability that currently exist. We emphasize that we do not take a stance on whether FR should be used (in its current state or ever) or what the ideal approach should be. Rather, we attempt to summarize a wide range of relevant considerations when making these decisions and policy options that fall between the extremes of permissive use and banning law enforcement use of the technology.

The remainder of this chapter provides background on how law enforcement conducts identification, how FR is used by law enforcement, issues in the use of FR by law enforcement, and what is known about existing policies and practices.

Identification in Law Enforcement

Identification is a critical function of law enforcement. It is important for distinguishing individuals to ensure that the appropriate person has been detained and processed through the justice system. Identification is also

critical for determining whether a particular person or object was present at the location of a crime or was present at locations connected to a crime. Methods of identification have evolved dramatically, often because of the availability of new technologies.

Technological Advances to Aid in Identification

Methods to improve identification are constantly evolving. The most-established methods fall under the forensic identification category. Three primary methods of forensic identification are fraction ridge analysis (fingerprints), forensic odontology (teeth), and DNA analysis (genetic material), a relatively recent addition. A key feature of these identifiers is that they are reliable indicators of identity when appropriately examined because they are based on a certified source that is not easily tampered with (e.g., you cannot readily morph or deepfake DNA at a crime scene; see Ngan et al., 2020; Tolosana et al., 2022; Vincent, 2022). Further, the comparison is being made to a known source that was obtained directly—and legally—from that source.

Secondary methods of identification can be useful where primary identifiers are inconclusive or missing. These may include examinations of medical information, pathology, personal effects (e.g., tattoos, piercings, jewelry, clothing, shoes), or other pieces of evidence, such as handwriting or firearms that may be used in combination to identify an individual. Forms of digital forensic evidence, such as cell phone tower pinging, GPS coordinates, or bank records also serve as secondary methods to link someone to a location at a particular time. It is important to note that each of the above methods of identification—both the primary and secondary—has accuracy limitations that we will not discuss further in this report.

FR can be viewed as a key advancement in methods of identification, with characteristics that make it distinct in some important ways from previous identification methods. For instance, although similar in some ways to primary methods, FR is somewhat different in that it can be implemented from a distance (i.e., it does not require direct contact). It is this passive ability to capture all faces in an area that makes FR potentially much more expansive—and much more intrusive—than other forms of primary identification. Although the potential breadth of FR's ability to identify people is dependent on how widely cameras and video surveillance systems are deployed, FR can analyze video data with much less human effort than a manual review. FR is also different from secondary identification in that the face is—with few exceptions[3]—a unique biometric identifier on its own. Nevertheless, a number of factors can impede the accuracy of facial image capture and FR systems, including image quality (e.g., illuminance, angle, occlusions).

How Do Facial Recognition Systems Work?

Unlike the primary identification methods discussed above, FR can be practically employed to assist law enforcement in several distinct tasks: identity verification (is this person who they say they are?), identification (who is this person?), *and* surveillance (can we identify any of these people?). The technology uses a generally similar approach for each type of use. Even when FR technology is not being used, investigators often perform a manual comparison against mugshot images to search for known individuals who match the description of a named suspect or circulate an image or a witness description and ask for public tips. Before the broad availability of FR, all of these tasks could possibly be performed by law enforcement manually, but the time and labor involved—particularly for seeking to identify all individuals in a large group—made doing so impractical.

The algorithmic FR process, which is based on the use of artificial intelligence, typically begins with a probe image, which is submitted for comparison to a database (see Figure 1.1). This database, which some-

[3] One example is identical twins. However, aging is also a known factor that impedes the ability of FR to recognize the same face over time.

FIGURE 1.1
Facial Recognition Process and Key Features

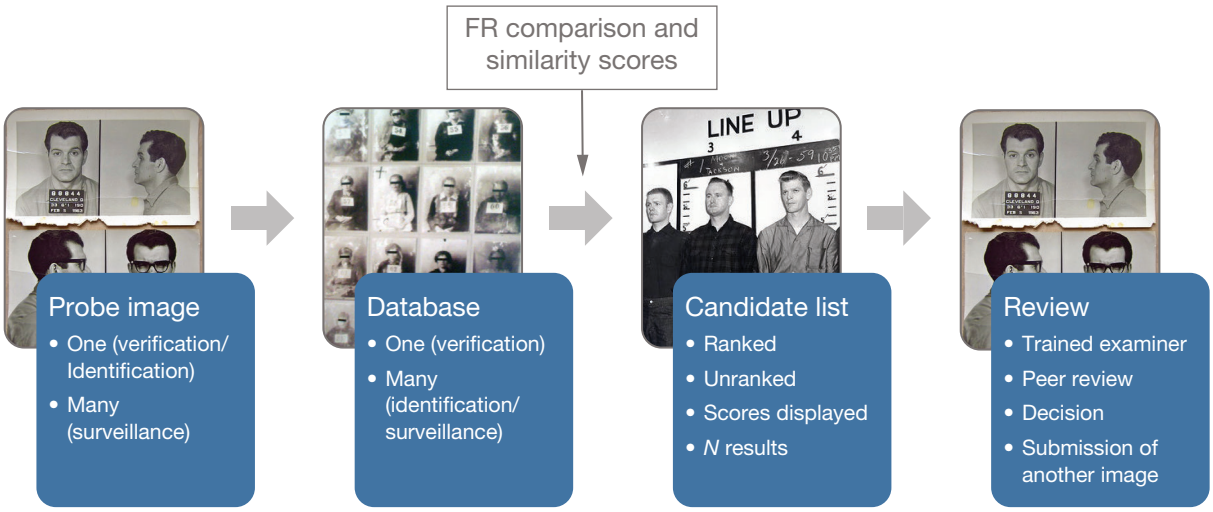

SOURCES: angus mcdiarmid, "Joseph C llacqua," October 28, 2012 (CC BY-NC 2.0, https://creativecommons.org/licenses/by-nc/2.0/); Miikka H, "21042007," April 21, 2007 (CC BY-SA 2.0, https://creativecommons.org/licenses/by-sa/2.0/); ~BostonBill~, "the usual suspects," March 26, 1959 (CC BY-NC 2.0, https://creativecommons.org/licenses/by-nc/2.0/).

times includes millions or billions of images, may be used for one-to-one image matching (for verification),[4] one-to-many image matching (for identification),[5] or many-to-many image matching (for identification or surveillance).[6] To make this comparison between the probe image and images in the database, FR algorithms calculate the similarity in faces represented by pairs of images. Early geometry-based approaches plotted points on faces (see Figure 1.2) to measure the distances between facial features. Recent advances have used deep-learning approaches (Payal and Goyani, 2020). For instance, among the best algorithms in a one-to-many setting, NIST testing has shown that false-negative identification rates for rank = 20 testing and false-positive identification rate thresholds of 0.001 have improved from 2.9 percent to 0.15 percent and from 9.7 percent to 1.6 percent, respectively, from 2014 to 2018.[7] NIST attributes this to the change from old algorithms to those using convolutional neural networks (Grother, Ngan, and Hanaoka, 2018).

[4] For example, one-to-one matching might be used to check that the identity provided to an officer corresponds to the same name in a driver's license database. Each image (one) is compared with the other (one).

[5] For example, one-to-many matching would be used to search for a person who is not yet identified, such as a criminal suspect captured on security camera. A single probe image (one) might be compared with a mugshot database to see whether any of the people in the database (many) are highly similar to the probe image.

[6] Many-to-many matching might be used to identify multiple people at the scene of the crime after the fact from stored video, or it could be used to conduct ongoing searches of multiple people within a specified area. This might be used to scan an entire area during or after an event.

[7] Notably, these accuracy rates correspond to mugshot probe images, which are generally higher quality than images obtained in actual law enforcement use (e.g., closed-circuit television [CCTV] video stills). Poorer image quality generally means lower accuracy. *Rank = 20* means the algorithm returns the actual subject among the top 20 ranked candidates. A false-negative identification in that setting would mean the actual subject was not in the top 20 on the candidate list. Other accuracy testing sets a false-positive threshold (here, 99.9 percent true positives) and examines false-negative rates based on that threshold. Thresholds set a minimum similarity or confidence score at which the FR software returns results. Any images that rank below the threshold are not returned to users in their search.

FIGURE 1.2
Geometry-Based Facial Recognition

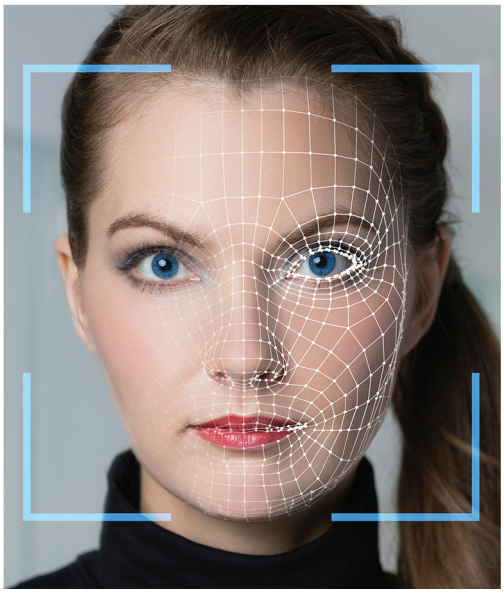

SOURCE: metamorworks, "Facial Recognition System Concept," Adobe Stock.

To make comparisons, FR systems typically use a combination of preprocessing, face detection, feature extraction, feature selection, and classification (see Payal and Goyani, 2020; Olszewska, 2016). Preprocessing is used to assess and remediate, if possible, image-quality issues (e.g., illumination, pose) in an original photo to make a more suitable probe photo to submit for an FR search. Face detection isolates the face from other features present in the image. Feature extraction involves extracting important facial features from the image and can be done either locally (certain regions of the face) or globally (the whole face). Feature selection involves methods to select a subset of the most-distinguishing and nonredundant features. Finally, classification involves probabilistic comparison of the vectors created from feature extraction to identify photos in the database with face images most similar to the probe photo. Many of the top-performing FR algorithms use deep convolutional neural networks (Grother, Ngan, and Hanaoka, 2022; Lu, Song, and Xu, 2021) to perform deep-learning tasks of feature extraction, selection, and classification.[8] Therefore, these models train an algorithm to identify and compare facial features contained in images by comparing thousands or millions of facial images with each other or by using image variations of the same face (i.e., data augmentation; see Wang, Wang, and Lian, 2020; Zhuchkov, 2021).[9] It is worth noting that many of the FR algorithms used by law enforcement are black boxes: Exactly how they function across the above components is not usually known.

The output of the comparison process is a candidate list, often ranked using similarity scores generated by the algorithm (see Figure 1.3). Importantly, users can set a threshold, or minimum level of similarity or con-

[8] For instance, the top performing algorithm for a $R = 50$ or when the true positive is in the top 50 results has a false-negative identification rate of 0.0006 (99.94 percent accuracy) on mugshot images. This would mean that the search results would exclude a true match one time out of every 1,667 searches.

[9] Data augmentation can help expand training data and, importantly, the range of data features by transforming images to create more-diverse datasets. This may help improve on issues that FR algorithms currently experience, such as occlusion, expression, illuminance (e.g., darker faces), and image quality.

FIGURE 1.3
Face Recognition Process

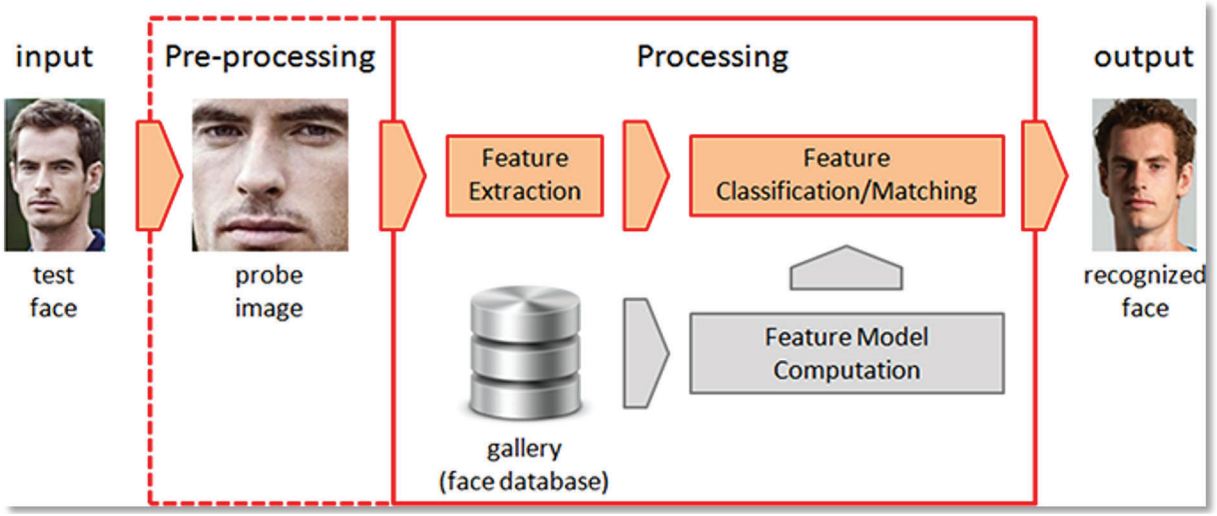

SOURCE: Olszewska, 2016 (CC BY 3.0, https://creativecommons.org/licenses/by/3.0/).

fidence score at which they would like the software to return results. This is set typically to limit the number of false positives. The software can also be set to return results ranked or unranked, to provide a certain number of results, and to display or not display similarity scores. These are all choices that can be made by a user or defined by the organization deploying the technology. Finally, the user has to make a decision about the available options in a candidate list and whether any of the candidates are potential matches. As a result, similar to other forensic identification approaches, there is generally a human in the loop interpreting the results of FR matching. Ideally, the user would be a trained facial examiner, and the process would use blind peer review before the decision to return one or more leads is made. The users may also decide to submit additional images at this point.

Additionally, there are different types of FR searches law enforcement might perform. These include one-to-one or verification searches (e.g., identity confirmation in jail), one-to-many or identification searches (e.g., investigative leads for an unknown suspect), and many-to-many searches, which can occur from stored data (e.g., geofence warrants) or occur in real time (e.g., scanning crowds for individuals with open warrants). We further discuss these search types in Chapter 3.

Current Issues in Use of Facial Recognition by Police: Regulation, Privacy, Civil Rights, Bias, and Public Perceptions

Many communities are considering whether to allow their police departments to use FR. Because it is a relatively new technology, concerns about misuse by law enforcement, differential impacts or bias, and privacy and civil rights are central to the current debate and likely factor into public perceptions. As part of this study, we fielded three survey questions through the RAND American Life Panel (see Appendix A for more detail) to take a snapshot of current public perceptions of FR. Through that survey, we found that 49 percent of American adults supported or strongly supported police use of FR systems, and 43 percent had no concerns about being included in their local police department's database if the agency did so. A 2022 Pew survey (Rainie et al., 2022) found that 46 percent of American adults believed that "widespread" use of FR would be a good idea for society, while 27 percent felt that it would be a bad idea and 27 percent were unsure.

However, general trust in the police was a major factor influencing the survey respondents. Our findings from the survey show that those who trust the police more also support police using FR technologies (see Figure 1.4). When there is lower trust among the community members, there is also lower acceptance of the use of these technologies (dark blue in Figure 1.4). Table 1.1 examines the demographics of these groups. The groups that trust police more tend to be older and have a higher proportion of individuals who reported their race as White.

The Pew survey had similar findings. Young respondents were more likely to consider it a bad idea (42 percent), while Black respondents only differed slightly, with 30 percent responding "bad idea" and 40 percent "good idea" (Rainie et al., 2022). The lack of certainty about the use of FR technology is not surprising given how new it is and the general lack of details about how it is being used by law enforcement. It is also likely that some people are comfortable with some uses and not others. These findings suggest that it is important for law enforcement to be clear about their plans for FR use and seek community input about specific uses. If FR use contributes to increased distrust of the police among communities who already have low trust in police, it could counteract the goals of using the technology in the first place—namely, improving crime clearance and reducing incidence of crime.[10]

Importantly, law enforcement use of FR comes at a time where there is growing use of FR by other entities and acceptance (e.g., commercial systems, public availability, and social media companies). Other government applications, such as identity verification at airports and the proposed use by the Internal Revenue Service (IRS) for tax returns, have drawn differing levels of scrutiny. Notably, the IRS backed away from its plans after public uproar and concerns from the public and political figures (Ebeling, 2022). Despite concerns, regulation of both government and private use of FR technology lags nationally, with few state laws and no federal legislation.[11]

Legislators and policymakers will need to confront these issues and regulate use in a systematic way. FR presents challenges that have some analogues in other areas of law and policy, but its differences from other identification methods also raise unique issues. For instance, the facial image databases being used for FR in law enforcement typically appear to be expanding rather than narrowing, presenting issues around privacy and anonymity, consent, data security, accuracy and bias, data sharing, and so on that will need to be addressed.

Controversy surrounding the use of FR by law enforcement is not unrelated to ongoing discussions about policing in general. Most prominently, this includes debates about the expansion of state power and associated due process and privacy concerns, along with concerns about bias and disparities in policing and how FR could exacerbate both. In the most extreme application, a law enforcement organization, if given access to a large-enough network of cameras equipped with FR, could track everyone's public movements extensively and with much less effort and with questionable legal authority compared with current practices (e.g., wiretaps, physical surveillance). Expansive use of FR technologies in this way would certainly raise privacy issues and would also potentially invite scrutiny about the appropriate balance of power between law enforcement and the general population.

Additional concerns relate to potential bias and disparate impact. These concerns stem from the data sources being searched, the nature of the searches, bias in the algorithms at the heart of FR systems, and disparities in policing in general. Police FR searches commonly include mugshot searches, and, for a variety of

[10] Research on trust and legitimacy of the police suggests that negative views of the police are associated with less crime reporting and higher levels of crime in a community. For examples, see Desmond, Papachristos, and Kirk, 2016; Kirk and Papachristos, 2011; Papachristos, Meares, and Fagan, 2012.

[11] However, the Illinois Biometric Information Privacy Act was enacted in 2008 (Illinois Compiled Statutes, 2008).

FIGURE 1.4

Overview of Survey Responses by Trust in Police

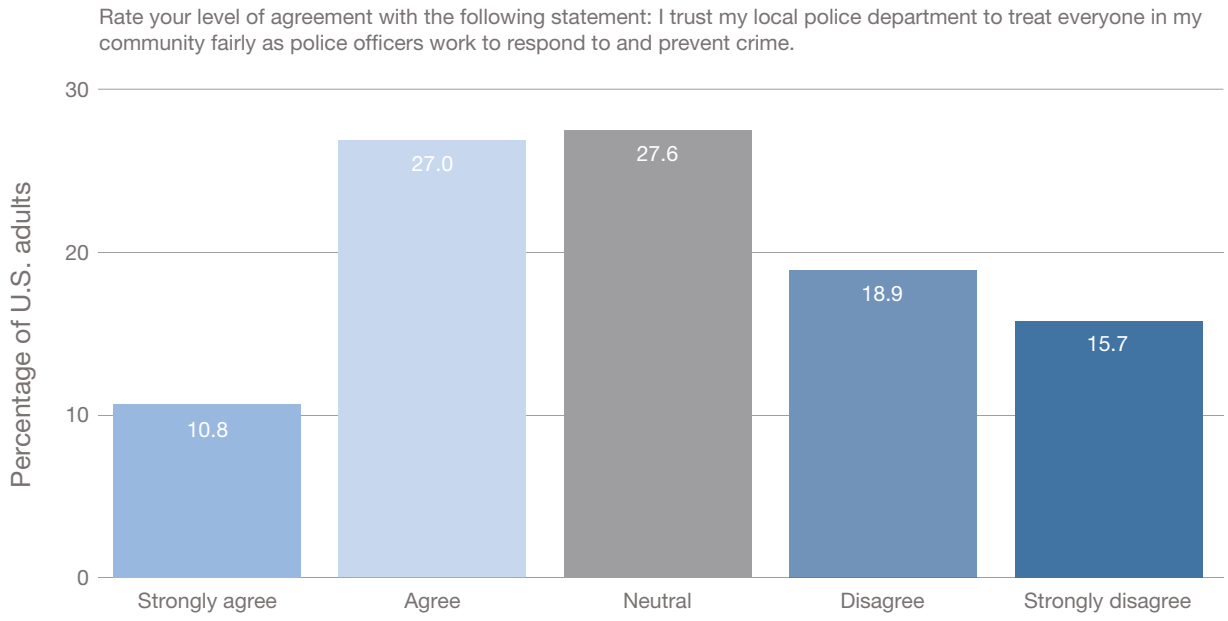

NOTE: Percentages might not add to 100% because a small number of nonresponses were excluded after tabulation.

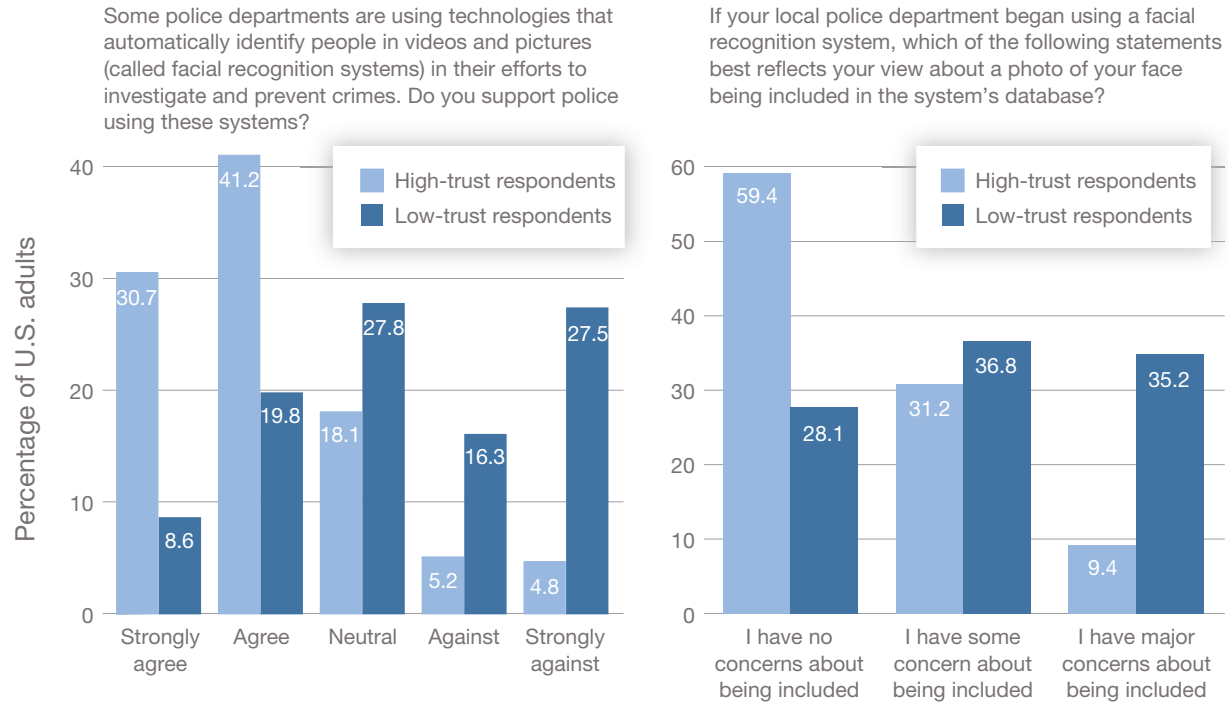

SOURCE: RAND American Life Panel.

NOTE: In the bottom two charts, results are broken out by those who responded in the top panel with "strongly agree" (high-trust respondents) and "strongly disagree" (low-trust respondents).

TABLE 1.1
Overview of Survey Response by Trust in Police

Demographics	Agree or Strongly Agree with Trusting Local Police	Disagree or Strongly Disagree with Trusting Local Police
Male	50%	48%
Female	50%	52%
White	85%	72%
Black	6%	16%
Other race	9%	12%
Average age	52	42

SOURCE: RAND American Life Panel.

historical reasons, there are racial disparities in arrests and therefore an overrepresentation of racial minorities in mugshots (Lofstrom et al., 2019). An agency relying primarily on mugshots as one of the primary sources for FR matches would amplify those historical disparities by increasing the odds of being identified. Indeed, recent research has found an association between deployment of FR by police and racial disparities in arrests (Johnson et al., 2022). More research on this is critical for understanding disparate impact. Moreover, crimes that occur in public view may be more likely to be committed by some sociodemographic groups relative to crimes that occur in private, further enhancing police intervention for some groups and not others (Engel, Smith, and Cullen, 2012; Fagan and Davies, 2000; Gaston, 2019). Disparate impacts may apply not only to criminal justice involvement but to privacy issues as well. Groups that are more likely to have their image captured in public, through CCTV or direct police contact, or are searched in an FR database and have all of this data retained are therefore more subject to privacy invasion than others. One article (Moy, 2021) notes that police technology use can aggravate inequity by replicating the inequities that exist in policing, masking inequities in policing, transferring inequities from elsewhere to policing, exacerbating inequitable policing harms, and compromising oversight of inequity in policing. Police use of FR may create risks across these five aggravating factors, most notably through disparities in the underlying training data (images used to train an algorithm) or gallery data (from a database, such as for mugshots) being used in FR searches (see Chapter 3 for more information). Disparities in training data may reduce accuracy for certain groups, which transfers inequities from elsewhere into policing and compromises oversight because FR algorithms are not public and information on their use by police might not be reported. Disparities in gallery data could replicate inequities that exist in policing because mugshot data, for example, reflect disparate policing impacts and could exacerbate inequitable policing by being more likely to return results for certain groups due to having disparate mugshot data. The nature of this underlying data might not be known or reported, which also compromises oversight.

Lastly, it is well established that FR algorithms are less accurate for darker faces and female faces (Cook et al., 2019; Grother, Ngan, and Hanaoka, 2019b). FR companies have been substantially improving accuracy overall and among demographic subgroups in controlled accuracy testing, leading to questions about the practical implications of small differences in accuracy (e.g., 0.06 percent or lower depending on the application).[12] Questions of how well these tests generalize to real-world conditions also need to be answered.

[12] Testing of demographic differences by NIST shows that false positives are highest for American Indian, Black, and Asian mugshot images (0.0063, 0.001, and 0.0013, respectively, relative to 0.0001 for White males). There is moderate variation across algorithms. Some developers, however, provided algorithms that showed no evidence of false-positive differentials (Grother, Ngan, and Hanaoka, 2019b). Setting higher thresholds, refining training data, more-diverse training datasets, and

Law enforcement documents currently cite racial differences in accuracy as "undetectable," but it should be noted that the tests being cited did not use "wild" or unconstrained images (e.g., images from a surveillance camera of someone in motion), which are more similar to those used in law enforcement investigations and are not tested. NIST notes that error is highly application specific (Grother, Ngan, and Hanaoka, 2019b). Particularly when it comes to discussion of accuracy and bias, it is important to be clear about what tests actually show and how well those tests translate to a law enforcement context.

What Do We Know About Existing Policy and Practice?

In *The Perpetual Lineup*, Georgetown University researchers identified 52 law enforcement agencies across the United States that use or have used FR. However, other reporting suggests this number is much higher (Garvie, Bedoya, and Frankle, 2016; U.S. Government Accountability Office, 2021b). Calls for federal legislation on FR have existed for some time. A bill called the Facial Recognition and Biometric Technology Moratorium Act of 2021 was introduced in June 2021 in both the House and the Senate and was referred to the Committee on the Judiciary and the Committee on Oversight and Reform (Senate Bill 2052, 2021). This bill would enact a moratorium on federal use of FR systems until Congress passes a bill authorizing their use and prohibits receipt of federal law enforcement grants for state and local governments that use FR. Meanwhile, the Federal Bureau of Investigation (FBI) continues to use FR and train local law enforcement on FR use and facial examination techniques. Other moratoria on FR use by law enforcement have been enacted at the state and local levels, but moratoria are often temporary. Indeed, of the two states that had enacted full bans, Virginia has recently enabled use through regulation (the other state being Vermont).

Where FR is allowed, a more standardized approach would be of interest to all parties; many of our law enforcement interviewees expressed concerns about how public perceptions can be negatively affected when an agency elsewhere engages in improper use (see Chapter 2 and Appendix B for details about our interviews with subject-matter experts). A report by the Center for Strategic and International Studies (Lewis and Crumpler, 2021) details some of the current regulatory efforts, and we briefly focus on a few. Two recent efforts to legislate the use of FR in law enforcement come from Massachusetts and Virginia. In Table 1.2, we compare the Massachusetts law and recommendations made in a recent special commission report (Facial Recognition Commission, 2022), which have yet to be enacted, with Virginia's law, which was recently enacted. Importantly, these two statutes cover many of the same policy dimensions but vary in terms of the requirements within those dimensions. For instance, the Massachusetts special commission recommended limiting FR use to felony offenses—and only after probable cause and a judicial warrant are obtained. There are exceptions for emergencies and deceased persons. Virginia requires reasonable suspicion, but FR can be used for any crime, can be used to identify victims, and can be used on detained persons.

In reality, there is no such thing as a perfect FR system that will satisfy the interests of all parties. Inherent in the application and regulation of FR technology use in law enforcement is a series of trade-offs, value judgments, and unknown or unquantifiable risks and benefits. This is true for a variety of general policing practices and police uses of other technologies, but applications of FR in law enforcement are particularly high stakes. Policy decisions could mean the difference between solving a murder and seeking justice or the murderer going free and possibly an innocent person arrested in their place. Values are challenging to reflect adequately in policy, but sufficient ground has been covered in existing policy to guide the key dimensions of law enforcement FR policies. What is missing is an assessment of the trade-offs of different policy dimensions

discriminative feature discovery may help with demographic differences in false positives. Nevertheless, it is not clear how well these tests generalize to a law enforcement setting in which image quality is likely poorer.

TABLE 1.2
Example State Facial Recognition Statutes

Policy Dimension	Massachusetts	Virginia
Authorization	Felony offenses	Crime suspects, crime victims, human or other trafficking, impaired or incapacitated persons, deceased persons, person who is danger to self or others, lawfully detained persons, imminent threats to public safety, life, or national security, vetting undercover law enforcement, driver's license and other identification verification
Authorization	Probable cause and judicial warrant	Reasonable suspicion
Centralization	Massachusetts State Police only	Virginia State Police to draft model policy as minimum standards for local-agency policy
Acquisition	Registry of Motor Vehicles or FBI FR systems, or approval from Executive Office of Technology Services and Security	Division of Purchases and Supply approval; must be NIST tested, with 98% true positives or better, minimal variations across demographics, annual NIST testing, and notification of local governing body first
Exigent circumstances	Emergency involving immediate danger of death or serious injury; deceased person	Not applicable
Scope of search	Ban on surveillance, tracking, and emotion recognition	Ban on tracking and creation of databases of images from live video
Transparency	Notification to the defendant; quarterly reporting on use; documentation of each search; annual report summarizing use	Annual policy updates; meet or exceed State Police Model Facial Recognition Technology Policy; maintain complete records for discovery, postconviction, public reporting, and auditing; annual report with main reporting categories from records
Accountability	Exclusion of information obtained in violation	Penalties for misuse: misdemeanor or termination
Training	None	Training requirements; procedures for confirming findings and secondary examination

SOURCES: RAND analysis of Massachusetts Bill S.2963, 2020; Virginia Senate Bill 741, 2022.

and practices that exist and whether they actually serve to minimize the various risks while maximizing the benefits.[13]

The Current Study

This report explores potentially important features of different types of FR use; the risks, benefits, and trade-offs; and policy guidelines that might address those risks. This description is intended to reflect the state of knowledge and practice at present, in an effort to summarize various perspectives that are currently captured only across disparate sources (e.g., news media, scientific literature, legal scholarship, and existing policies and practices). The main goal of this report is to help inform the debate about law enforcement use of FR technology for policymakers.

[13] Other documents exist that have provided policy, practice, and regulatory guidance for different audiences, such as Garvie, Bedoya, and Frankle, 2016; Major Cities Chiefs Association Facial Recognition Working Group, 2021; the Policing Project, undated; and International Criminal Police Organization, 2021.

Broader Concerns About Pervasive Surveillance and Artificial Intelligence

The technological capability to surveil large areas or track people across several areas using FR or other surveillance technologies already exists today. For instance, in China, government surveillance automatically links to "social credit scores" that dictate access to public services (Kobie, 2019). Other surveillance opportunities are being developed and may be deployed at scale in the near future (e.g., object recognition). Moreover, the combined use of various surveillance technologies (e.g., drones or aerial surveillance, license plate readers, stingray devices) in a network of surveillance has multiplicative effects on privacy and fundamental civil rights that are beyond the scope of this current project. Many of these applications in law enforcement also have unknown or questionable benefits.[14] We do discuss issues related to surveillance uses of FR in this report, but because of the lack of widespread employment by law enforcement in the United States, this is covered lightly in our findings. The increased use of big data and data analytics by police (Ferguson, 2017), networks of surveillance (Friedman, 2022), and frameworks for artificial intelligence regulation generally are critical discussions that FR is a part of but are beyond what we can cover here.

Further, this report is focused on FR use by law enforcement only, and issues concerning pervasive surveillance using other technologies by other entities (e.g., private businesses) is important but beyond our scope. Lastly, this is a rapidly evolving space, and avenues for untoward governmental use of FR or other artificial intelligence or surveillance systems depends on both the capability of the technology and the evolution of legal regimes.[15] Regulation will have to continue to evolve in response.

Approach

We conducted three interrelated research activities: literature and document review, interviews with subject-matter experts, and a nationally representative survey. First, we consulted the various sources mentioned above (news media, scientific literature, legal scholarship, existing policies and practices, technical documents, etc.) to better understand the landscape of FR use in general and FR use in law enforcement in particular. This literature and document review helped us understand the key debates, ongoing technological developments, and policy frameworks that currently exist. We then used this information to frame our research interview protocols and develop a set of notional use (misuse) cases and policy dimensions for discussion in our research interviews. Finally, given the importance of public perception around this subject, we sought to add to the existing knowledge about public perceptions of FR use by law enforcement by asking three questions as part of RAND's ongoing American Life Panel survey, which periodically surveys a nationally representative panel of Americans. In particular, we were interested in how trust in law enforcement might be affected by people learning that their face was part of an FR database.

In Chapter 2, we explain the development of the use cases and policy dimensions in more detail. In Chapter 3, we further describe the findings of the literature review and the interviews by building a framework of the concerns about FR use. In Chapter 4, we propose the development of a road map to help policymakers navigate the concerns. Finally, in Chapter 5, we provide our conclusions. Appendix A describes our survey methodology, and Appendix B presents our interview protocol.

[14] See, for example, a RAND evaluation of Baltimore's Aerial Investigation Research Pilot Program (Smart, Morral, and Schell, 2022), evaluations of predictive policing (Saunders, Hunt, and Hollywood, 2016), or evaluations of live FR (Fussey and Murray, 2019).

[15] For example, researchers have used FR to predict people's sexual orientation (Hawkins, 2017). In light of the 2022 *Dobbs v. Jackson Women's Health Organization* Supreme Court decision, this brings up questions as to whether FR technology or data could also be used to identify pregnant women, although other algorithms could also be used to predict both pregnancy and sexual orientation (Siegel, 2020; Kahn, 2022).

CHAPTER 2

Developing a Spectrum of Use Cases as a Step Toward a Broadly Practical Approach for Regulation

To explore the nature of concerns about FR and to identify a wide variety of potential approaches to addressing those concerns, this research effort began by identifying the spectrum of current or potential applications of the technology. This exercise was designed to develop both a lens for analyzing existing literature related to FR and an elicitation tool in interviews with relevant subject-matter experts to help us explore three key questions:

- What is the accuracy of the technologies and how should that accuracy shape use?
- What differences are there across uses and policy needs, and how should those needs vary by use?
- What are the complex organizational and procedural issues, such as the ownership of FR systems and data (e.g., whether there are differences between police using a commercially developed and maintained system and the police building an FR capability)?

We called these application examples use and misuse cases to explicitly acknowledge the potential for both appropriate *and* inappropriate use of FR by police agencies. In this of this chapter, we will clarify how we categorize the spectrum of cases and the details of our analysis.

What Do We Mean by Use (and Misuse) Cases?

The concept of a *use case* is frequently used in the examination of specific technologies and their applications by individuals and organizations. A use case is defined generally as an example application of a technology or system to perform a specific task or achieve a specific goal (see, e.g., Brush, 2022). Often applied in the development of new information systems or technology products, a well-described use case can also be utilized to examine the policy and other implications of a new technology. Framing where the technology is going to be used and, in the case of FR, whose photos will be included in the database or submitted to the system for identification allows much clearer examination of the potential risks and benefits of the application. Put simply, clear and specific use cases take examination of a new technology from the abstract—where all possible uses of the technology are considered simultaneously and debate may be unduly shaped by the most favorable or unfavorable examples of those uses—to the concrete: a more measured consideration of how and why a new technology tool might be used.

In considering FR employment in law enforcement and those potentially negative uses, we also considered the concept of *misuse cases*: applications of the technology that might be potentially attractive to individual members of law enforcement (e.g., an officer with direct access to an FR system) or to a police department but for which the use of the technology might be inappropriate or potentially damaging to individuals, groups, or society as a whole.

Although there is almost certainly a subset of misuse cases in which broad consensus can be built to disallow the application of FR, the difference between what is considered use and misuse is contested for most potential use cases. Just as use cases can help to set up more-systematic consideration of the implications of a new technology, misuse cases can help enable practical approaches for managing risks if and when the technology is deployed and used. Where consensus exists regarding a misuse case, such actions and applications can be prohibited, and the design of management and oversight regimens around the technology must focus on how to do so in a way that is effective and credible for all concerned stakeholders and groups. In contrast, if the decision is made to deploy a technology when use versus misuse is contested, employment of that technology becomes a target for monitoring, oversight, and ongoing assessment. Only through collecting and providing such information to stakeholders will it be possible for those stakeholders to reassess their initial view, potentially shifting toward a consensus about whether that use should be allowed or prohibited.

Use and Misuse Across the Current and Proposed Implementation of Facial Recognition

Current uses of FR involve a system of inputs and outputs involving both human decisions (e.g., as influenced by policy, training, and experience) and FR software (e.g., data inputs, algorithm quality, software settings). The entire system involved in any given use of FR has implications for the appropriateness, accuracy, results, and actions taken. Most importantly, ensuring that FR is used appropriately (if some uses are deemed appropriate) and protecting against misuse are critical for avoiding some of the concerns we discuss in the next section.

Misuse can occur at any point in the FR process, so we cover some illustrative examples here. For reference, we map out the general FR use process in Figure 2.1. The entire use process has opportunities for misuse, whether in the form of violating policy or in the form of currently allowed practices that create concerns about FR results. For instance, *authorized use* covers policies regarding what FR should be used for and by whom. Letting any officer use FR for any reason would likely not be seen as legitimate. Next, images vary in terms of their quality, and occasionally *probe images* are altered in various ways that lead to questions

FIGURE 2.1
Common Facial Recognition Use Steps

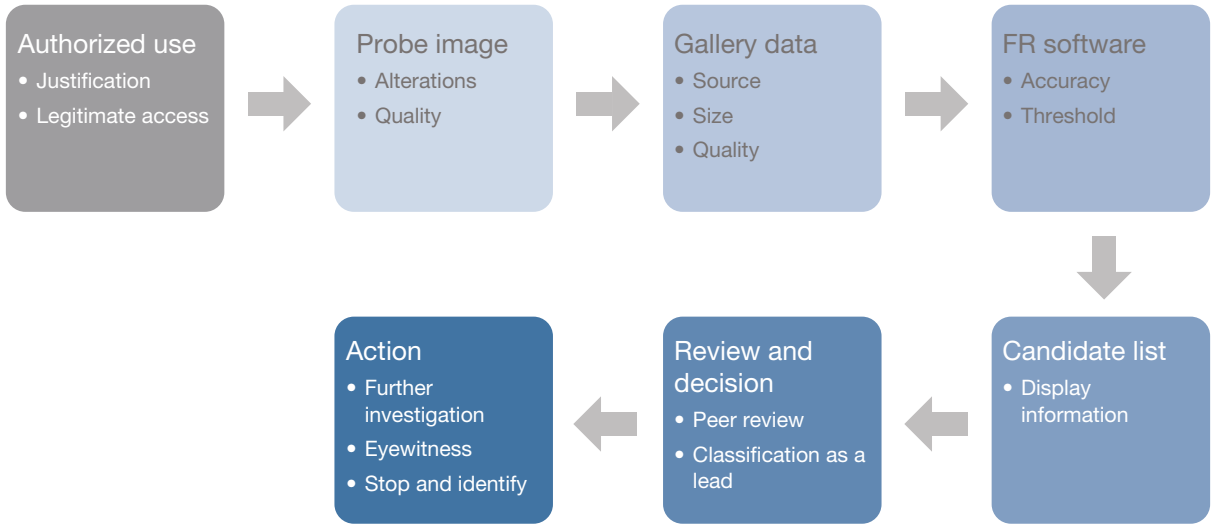

about the accuracy of the results (i.e., garbage in, garbage out). Concerns about bias point to *gallery data* and *FR software* (algorithm) performance as potential sources of bias.

Figure 2.2 summarizes some of the issues that exist in the use of FR by police that can be considered misuse.

The Potential for Clear Misuse

Many of the existing examples of clear misuse of FR technology or other law enforcement technology by police officers involve a lack of authorization for use. This could take the form of someone gaining unauthorized access to the FR system (e.g., stealing a login password) or someone with authorization using the technology without proper justification. The most common example of the latter is police officers' personal use of the technology. For example, as a means of intimidation, 58 officers in the Minneapolis Police Department searched another officer's personal information, consisting of her address, age, height, weight, and other information, without legal justification (Human Rights Watch, 2019). Civil suits in other localities have involved similar allegations about officers accessing law enforcement databases for personal use—for instance, collecting personal information (e.g., home addresses) of romantic interests, ex-partners or their new partners, and celebrities. Officers could use the technology to identify both personal information and personal locations or travel patterns, clear examples of misuse. Finally, another clear misuse, albeit slightly harder to detect, is the uncritical acceptance of FR results that either are directly used for eyewitness identification or, more problematically, noneyewitness identification (as in the case of Robert Williams in Michigan; see Burke, 2021), to directly make an arrest or take some other official action. In addition, some experts consider it a misuse to inform an eyewitness that FR was used for the lineup or showup because it creates an assumption that a true match has been found.[1]

Questionable Facial Recognition Practices

Next, there are a number of questionable practices on which experts differ in terms of their level of risk. These mostly relate to questionable use of probe images—for instance, using very poor-quality images, using probe images of celebrities or artist sketches, or heavily editing original images to be used as probe images (Garvie, 2019). These are concerning practices, because, by their very nature, they obscure features of the face, such as eyes and noses, that are critical for FR similarity determinations. Additionally, the FR results must be compared with the original image rather than the probe image, which will be more difficult if the

FIGURE 2.2
Categories of Potential Misuse

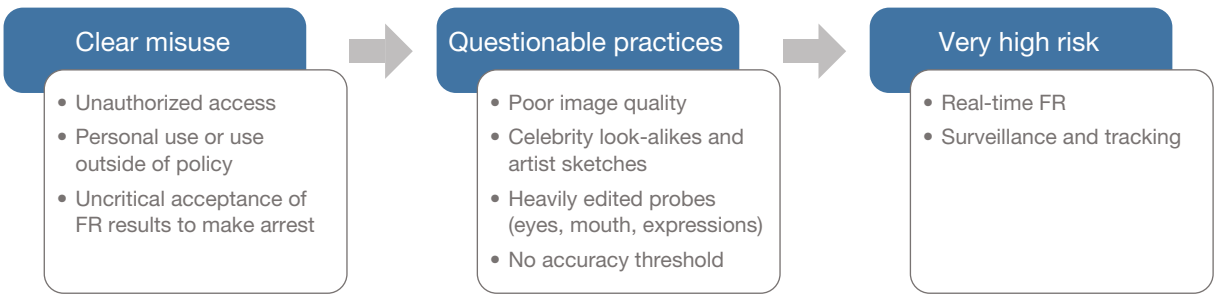

[1] A showup is when an officer brings a witness directly to identify a single suspect in person; this commonly happens at crime scenes (e.g., in the back of a police car), but there are also examples of suspect images being sent to witnesses via text message for identification (see Johnson, 2022).

original image was poor quality or if the creation of the probe significantly modifies the original image (i.e., with features that were not in the original). The outcome would be based on the modified image. Lastly, many of these modifications are done to address existing limitations of FR technology because matching is typically less accurate when there are obstructions, expressions, or other limitations present in the probe image. It is not clear how much such alternative practices affect the FR results. For instance, attempts to improve low-resolution images have found that they potentially correspond to many realistic high-resolution images, meaning that there are multiple potential ways to enhance an image, and it would not be clear which is correct (Menon et al., 2020). Thus, the choices made in image modification could vary with unknown impact on the final results. Other practices, such not using an accuracy threshold (i.e., a minimum similarity score or confidence score), can increase the risk of false positives being included in the candidate list returned in an FR search. There are some situations in which no threshold may be acceptable,[2] but much as with image modifications, the number of false positives that turn into an investigative lead is not known.

Very High-Risk Facial Recognition Practices: Forensic Facial Recognition Versus Surveillance

In the very high-risk category of FR practices,[3] we distinguish between forensic uses of FR (i.e., for investigative purposes after a crime has been committed) and surveillance uses of FR (e.g., real-time FR or facial tracking). It is worth noting that some authors use the term *surveillance* to also describe the source of imagery, such as store surveillance cameras. However, for the purposes of this report, we use *surveillance FR* to mean persistent uses of FR, whether used on people within a certain field of view or focused on a specific individual (i.e., FR tracking). The ongoing use of FR and its employment outside the context of a criminal event are distinguishing features of surveillance.[4]

Several studies have covered the issue of FR for surveillance purposes and typically distinguish it from forensic uses in terms of the legal, societal, and practical differences across the two broad types (see, e.g., Ferguson, 2021; Garvie, Bedoya, and Frankle, 2016; Rudin and Bushway, 2021). The legal considerations stem from concerns about violations of civil rights and due process, including the First, Fourth, Sixth, and Fourteenth Amendments to the U.S. Constitution, as well as other case law that covers issues related to government use of technology. Societal implications of surveillance are connected to the legal implications but include expectations of a certain level of privacy, concerns about bias, and concerns about government overreach (i.e., trust, legitimacy, and consent). Finally, there are practical concerns about surveillance uses of FR. First, the number of searches and potential matches created by an ongoing persistent surveillance application would require much more review capacity than is typically available in police FR units. Next, once a potential lead is generated, there is a need to determine what to do with that information. Should the police send an officer to approach and identify the individual? Should an officer just follow the person? Initiating police-driven stops through FR or committing resources to investigating potential matches both have costs to account for.

---

[2] In this case, the FR algorithm would return the top-ranked results in terms of similarity or confidence score. This might be used in an exploratory way but could return low-quality results.

[3] We borrowed the label for this category from the risk framework created by Center on Privacy and Technology at Georgetown Law, see Garvie, Bedoya, and Frankle, 2016.

[4] To be clear, with tracking FR applications, the individual being tracked is likely suspected of some crime, but the use of FR for tracking is connected to the individual rather than the crime.

Defining Use and Misuse Cases for the Interviews with Subject-Matter Experts

Ensuring that decisionmakers have a thorough and realistic idea of the use and misuse cases is critical for setting policy and practice for law enforcement's employment of FR. Such policy decisions set the basis for the scale of FR use by police, which helps determine the impact, including benefits or risks, of FR.

As of now, there is no universally accepted set of uses by law enforcement; each agency tends to have different accepted uses. A few documents have tried to characterize the range of existing uses and use processes (Law Enforcement Imaging Technology Task Force, 2019; Louradour and Madzou, 2021).

The IJIS Institute and International Association of Chiefs of Police report categorizes uses into three categories (Law Enforcement Imaging Technology Task Force, 2019):

- field use (random field interaction, reasonable suspicion interaction, active incident, deceased identification, lost and missing, identify fraud, interdiction)
- investigative use (active incident, photo array construction, evidence compilation, participant party identification, victim identification, criminal identification, suspect or associate identification)
- custodial or supervisory use (admittance identification, access control and movement, identification for release, identification for program participation, court appearances).

The World Economic Forum provides two sets of use cases (Louradour and Madzou, 2021). One set of use cases walks through more-specific applications or use cases for FR, including identity checking at border control, identifying someone suspected of ATM fraud, identifying a rioter, identifying a museum thief, actively looking for a terrorist in a public place, investigating suspected child abuse, and searching for missing persons. The other set of World Economic Forum use cases identifies both governmental and private uses for "safety and security of public spaces," including identity checking at borders (repeat); tracking persons of interest based on a warrant or terrorism risk (repeat); neighborhood-watch front-door cameras or external vehicle cameras; searching for missing persons (repeat); and providing private security, safety at public events, and safety in public spaces—including automated CCTV and movement tracking, police body cameras, and student-attendance tracking.

Use and Misuse Cases in the Current Study

Many of the existing law enforcement use case lists are not structured to identify potential issues within each application and often do not consider potential future uses of FR and how they might relate to current practice. Our intention for developing our set of use cases was to include those that tap into various policy dimensions and concerns relevant for current uses of FR and to explore how policy might need to adapt for the future.

To start, we reviewed some of the existing use case documents and other research to identify important considerations for current FR use in law enforcement and potential uses with little or no uptake yet. We began with two dimensions that seemed to distinguish many different uses: (1) the importance of the purpose for which FR is being used and (2) the scope of the search. The scope includes how many people are being searched (one or multiple) and how many people the searches are being conducted against (one, thousands, millions, etc.). We also wanted to consider forensic, limited real-time, and surveillance use cases. The use cases we created are shown in Table 2.1. Note that these are not all specific to law enforcement, but they either implicate law enforcement (e.g., law enforcement could be expected to act based on FR use) or represent other potential governmental uses.

TABLE 2.1
Hypothetical Use Cases for Interviews with Subject-Matter Experts

Type	Description
Investigative use for theft	CCTV images of a jewelry theft suspect are captured from multiple private businesses in the same area. Police use multiple still images from different angles in FR search via a mugshot database. *Features: medium/low importance, one-to-many scope, forensic use case, current time frame*
Investigative use for serious crime with image modification	A hit-and-run that injured two pedestrians is captured on video. An image of the driver's facial profile is captured on CCTV video. Officers create a mirrored image and conduct a FR search of state driver's license images. *Features: high importance, one-to-many scope, forensic use case, current time frame*
Private use; police implicated in response	Continuous FR is used at a concert to identify known individuals who have made threats against the performing artist and the concert. *Features: medium importance, several-to-many scope, real-time use case, current time frame*
Video analytics	Continuous video analytics are used to collect high-quality facial images from public and private CCTV video feeds (i.e., creating a face database from various video sources). *Features: low importance, many-to-many scope, real-time use case, future time frame*
Terror attack	After a terror attack, police launch a drone to collect video over the entire area and begin pulling all surveillance camera video from the area. Police submit all facial images to FR system to identify witnesses and suspects. *Features: high importance, many-to-many scope, real-time use case, future time frame*
Body-worn camera; real-time FR	Police body-worn camera video is equipped with live video FR software to scan for open warrants on people who refuse to provide identification. *Features: medium importance, one-to-many scope, real-time use case, current time frame*
Airport real-time FR	Authorities in an airport set up live-video FR to monitor for missing children *Features: high importance, many-to-many scope, real-time use case, future time frame*
Other government use to access public services	FR is used to access public transit. Public transit users must submit facial image and create digital account to set up payment. *Features: low importance, one-to-one scope, real-time use case, future time frame*
Social "credit scores"	FR is used as part of a system that computes social credit scores for various behaviors (e.g., debt, poor driving, buying too many video games, smoking in nonsmoking areas) to assign a value of "trustworthiness." Travel, loans, and education are restricted if an individual's value drops too low. *Features: low importance, many-to-many scope, real-time use case, future time frame*

In parallel, we also reviewed the research literature and constructed an initial set of policy dimensions to include as part of our interviews. After constructing the initial list, we consulted with other RAND researchers with backgrounds in law enforcement (specifically real-time surveillance), police use of technology, and law. We also intentionally asked our subject-matter-expert interviewees for their impressions about whether this list was complete and useful. Where they suggested it was not complete, we typically made adjustments or additions or combined categories. Table 2.2 provides the final list of policy dimensions that were included in our interviews. It does not exactly match the policy dimensions discussed in this report because some categories were added or combined later based on interviewee feedback.

Subject-Matter-Expert Interviews

To help us better understand how these use cases and policy dimensions apply to the real-world use of FR by law enforcement, we conducted 19 interviews with subject-matter experts. The interviewees spanned a range of stakeholders, including law enforcement leadership and officers with FR expertise ($n = 7$), representatives

TABLE 2.2
Policy Dimensions of Facial Recognition Use for Subject-Matter-Expert Interviews

Policy Dimension	Description
Importance of purpose	Does the current purpose justify the use of FR?
Accuracy and image quality	What level of accuracy is necessary across different types of uses?
Potential for error and bias	Are there aspects of the process that contribute to error and bias (e.g., match data, probe image, training, review)?
Potential for abuse	What is the risk of government overreach? What about improper use by individuals?
Underlying data source	If the source is governmental data, is there already a legitimate purpose for an image? If private, does it infringe on protected activity?
Human in the loop; training, standards, and guidelines	What decisions are being made by human users, and how might those enhance or mitigate other concerns?
Privacy and enrollment	Are data coming from public or private sources? Is there a reasonable expectation that behavior would not be shared with agents of the state? Is the facial image enrolled in a database for later use?
Character of observed activity	Is the activity being observed protected?
Level of security concern	Is FR use limited to situations in which the safety of everyone is paramount?
Consent	What form does consent take (e.g., warnings, terms of service, none)?
Scope of search and level of individualized suspicion	Are searches being conducted for a targeted or general purpose? Are searches one to one, one to many, or many to many?
Data sharing	What type of data are being shared? With whom and for what purposes?
Practicality	Would other means of identification be overly burdensome, or is the situation time sensitive? What other biometric identification options are available?

of civil rights groups or legal scholars (n = 5), researchers or oversight officials with expertise on FR use by law enforcement (n = 3), and FR software company representatives (n = 4).

Recruitment

We used a combination of independent outreach, referrals from existing contacts, and referrals from interviewees themselves. Two of the software-company connections were from forms submitted to their websites' "contact us" forms, and two were referrals. All of the law enforcement interviews were from existing contacts or referrals, and nearly all of the civil rights groups, legal scholars, oversight officials, and researchers were contacted via independent outreach (e.g., identified through research, quoted in news articles), with the exception of one referral. We also attempted to contact prosecutors' associations and judges' associations.

Interview Process and Interview Analysis

Interviews followed a general format of asking interviewees about their experiences, perceptions, and knowledge about FR use by law enforcement and associated policies and practices. We have included our interview protocol in Appendix B. Next, we introduced the use cases and asked interviewees to identify whether the use case was within their department's policy or explain whether they had any concerns about that use case. Lastly, we asked them to comment on the policy dimensions and to identify the most-important dimensions

from their perspective. We also asked them to weigh in on whether our list of policy dimensions was relevant and complete.

We did not engage in typical qualitative analysis methods (e.g., thematic analysis) but used the structure of the interviews to identify the most-pertinent information to our discussion of policy dimensions, existing practice, and considerations for regulating FR use.[5] Insights from our interviewees are combined with information obtained from policy and research documents to provide a broad sense of current policy and practice.

[5] Depending on the participant's expertise, we might have explored different items in more detail than others, contributing to slightly different interview questions across interviews. Therefore, we do not quantify interview findings.

CHAPTER 3

Understanding the Landscape of Concerns

A variety of efforts are currently underway to better understand the implications of FR use in law enforcement and throughout society. These span many legal, societal, psychological, practical, and policy considerations. Of note, concerns about FR depend greatly on how FR is used and what alternatives are available in a given situation. Moreover, there appear to be fairly large differences in perceptions of the acceptability of use between *forensic* and *surveillance* use cases, which differ on a variety of important dimensions that we discuss in more detail (Rudin and Bushway, 2021). In this chapter, we lay out the landscape of concerns about FR use that have been raised in news media, research literature, and by our interviewees that will need to be considered when crafting FR regulation, policy, and practice.

First, we describe how ongoing developments in policy and practice are attempting to mitigate concerns and what we need to know to improve policy and practice. Second, we detail a framework for thinking about risks and benefits of FR in the law enforcement setting and across use types. Finally, we map out the landscape of concerns and discuss the potential risks and benefits across a variety of policy and practice dimensions. The findings discussed in this chapter are based on reviews of existing policy, research and technical documentation, discussions among our research team, and interviews with subject-matter experts across the policing, policy, legal, and technological domains.

Thinking About Potential Benefits and Risks of Facial Recognition

The clash between the desire of law enforcement to use FR for identifying individuals and concerns from civil society groups stems partly from different emphases on the risks and benefits of using FR. Legal and policy actions in reaction to FR use or proposed FR use can respond to those risks and benefits, seeking—in different ways—to reduce risk or maximize the potential benefits. Use of FR by law enforcement is also occurring in a broader technology ecosystem in which private-sector entities are using the technology for various purposes and are providing FR services to others, including the general public. Although a private company using the technology for such purposes as customer profiling is qualitatively different from police using it to inform actions that could lead to an individual's arrest, the line between government and nongovernment use is blurry, since law enforcement agencies could purchase services to use FR even if they do not own it, or others with FR could provide information to the police based on FR searches they perform.

Although law enforcement use of FR is often treated as a single activity, FR actually has a wide range of applications and distinct goals with distinct combinations of risks and benefits. Police having FR systems scanning video feeds from across a city, potentially seeking to identify and track thousands of people a day, is qualitatively different from an officer using an FR device to identify someone who refuses to provide identification at a lawful traffic stop. The scale and scope of risks—to individual privacy and individual liberty and the potential for abuse—warrant a very high degree of oversight and public approval. To support robust debate about these types of technologies and public consideration of risks and benefits, pieces of proposed legislation advanced by civil society groups covering police use of technology and surveillance writ large

require public notice, external approval for use, and transparency over the outcomes of the use of new systems and capabilities (American Civil Liberties Union of California, 2014).

The different risks and benefits for different uses of FR accrue to three categories of individuals and groups, with some risks and benefits affecting more than one of them: law enforcement agencies, the subject of a scan, and the broader society (Figure 3.1).

Law Enforcement Agencies

Much of the focus of discussions about FR use by police is legitimately on police departments and the benefits that FR can provide by making those departments more effective. FR could make a variety of identification tasks in policing more efficient (e.g., verifying someone's identity at a traffic stop). Perhaps most importantly, FR could make police efforts to solve crimes more effective, an important consideration given that the percentage of crimes for which a suspect is identified and arrested is quite low (e.g., homicide clearance rates are near 50 percent nationally and decreasing over time; Braga and Dusseault, 2018; Ousey and Lee, 2010; Rodricks, 2021). However, the use of such technologies poses risks to agencies as well.

Misuse or perceptions of misuse can undermine public trust, and willingness to cooperate with police could substantially offset the benefits of FR for investigative purposes. If actions taken based on the use of such technologies are viewed as so inappropriate that they lead to judgments against law enforcement agencies, they could have tangible monetary costs as well (Ouss and Rappaport, 2020). In the longer term, concerns about police misuse of technologies could lead to unwillingness to provide new tools to agencies, hampering their ability to respond to new needs and requirements.

Society

More-effective policing would be a substantial benefit to society. Solving more crimes and holding perpetrators accountable would advance justice and provide a deterrent against additional crime.[1] Efficient law

FIGURE 3.1
Three Groups with Perspectives on Facial Recognition

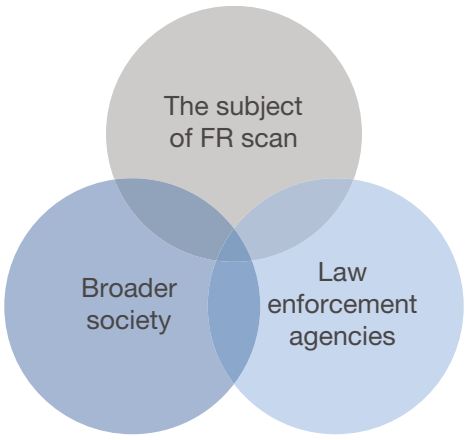

[1] It is worth noting that the public is not broadly convinced of the effectiveness of this technology. In a recent Pew Research Center study on emerging technologies, even though a larger percentage of the respondents thought that broad use of FR by police was a "good idea" than those that thought it was a "bad idea," only a third of the respondents thought that FR would

enforcement can also benefit society in another way: Funds not spent on policing can be spent in pursuit of other priorities.[2]

Counterintuitively, efficient law enforcement could also be viewed as a *risk* of FR technology. Practical constraints are a real check on the power of police and government, and concerns about governments' excessive use of power run deep in American society across the political spectrum. Questions about disparate impact are also important to consider. Police have the capability and authority to identify individuals in many circumstances, but doing so one at a time—e.g., by stopping a person to request identification and then assessing whether the documentation is genuine—requires time and effort. Being able to rapidly identify everyone in a large area with little effort is fundamentally different, stripping away the ability to be both in public and also semiobscure (if not anonymous). As a result, broad use of surveillance technologies, including FR, substantially changes the balance of power between police and the public and removes barriers that have constrained police power in the past.

Subjects of FR Scans

For individuals directly affected by law enforcement use of FR, risks mostly outweigh any benefits. These individuals fall into two groups: individuals whose photo is taken in a particular location and submitted to a FR system for identification and the larger number of people whose photos are included in the FR system for matching and identification purposes. Different types of risks are attached to each.[3]

With the first group, there is a potentially important privacy risk, depending on the circumstances in which the identification is made.[4] With both groups, there is risk associated with technology misidentifying one person as someone else. Although the accuracy of FR systems in matching photos to one another can be very high according to current testing, there is always risk of an innocent person being misidentified as a suspect in a crime they did not commit, with the potential for substantial consequences. Subsequent investigative efforts might be able to rapidly discover and resolve such a misidentification, but that outcome is certainly not assured. History provides many examples of individuals who have been wrongfully convicted of crimes that they did not commit, and a misidentification by FR—if not caught and discarded—is a potential starting point for such a miscarriage of justice.[5]

However, FR could also provide a potential benefit here. Each of the methods police use to identify suspects in criminal investigations has its own limitations and risks of error. Eyewitness identification—

cause crime rates to go down. In addition, 57 percent thought that crime would stay about the same, and 8 percent thought that it would increase as a result (Rainie et al., 2022).

[2] There is some similarity in this argument to elements of the debate in favor of defunding police: the need to use resources for other responders to incidents involving, e.g., mental health concerns versus resources spent on criminal justice responses to those problems.

[3] Note that we have framed this issue regarding individuals who are affected by FR societally—i.e., on members of the public. However, these risks could apply to police officers as well. In some cases, police seek to hide their identities from the public because of concerns about risk when officers identities become publicly known (e.g., doxed). Resistance to broad use of FR by police—put another way, a capability for a law enforcement user to dox any individual at any time—comes from an analogous concern.

[4] For example, law enforcement using FR to identify everyone in the vicinity of a sensitive medical provider where a crime had occurred—e.g., at a substance abuse treatment facility or mental health care provider—could have a substantial privacy impact.

[5] Although some past cases have involved intentional actions by law enforcement officers that resulted in an innocent person blamed for the actions of others, wrongful conviction does not necessarily require malfeasance by investigators. Mental biases can lead individuals to seek out information that confirms or reinforces their existing beliefs and theories, and such biases could shape how investigators collect and use data in criminal investigations.

essentially the human version of FR—is notoriously inaccurate. The use of FR might reduce the occurrence of misidentification by narrowing the group of candidates early in the process.

Mapping the Landscape of Concerns

With this general framing in place, we now turn to a discussion of some of the key concerns identified in news reports, research literature, existing policy, and discussions with subject-matter experts. In this section, we describe the various concerns, the range of how concerns might apply across different uses of FR, and potential implications for risks and benefits to different groups. We organize the concerns into the following categories:

- importance of the purpose for which FR use is being considered (and authorizations for different types of use)
- scope of the search and level of individual suspicion
- sources of error in FR accuracy
- sources of bias in FR accuracy
- data sources used for training and for comparison.

Importance of Purpose: Concerns About Use of Facial Recognition for Minor Incidents

Most policies governing FR use by law enforcement start with defining what FR can and cannot be used for. This is a critical first step because the different uses of FR may have different implications for some of the other concerns, such as accuracy, privacy, or government overreach. To address these concerns, some agencies limit their use of FR to cases in which there is probable cause that a crime has been committed. The seriousness of the suspected crime can also be a factor: Some departments allow FR to be used for all crimes, whereas others limit use to such crimes as violent crimes or felonies.

Limiting FR use to criminal incidents is central to the stated purpose of law enforcement acquisition of these systems, but it is often not the only authorized use. Some agencies may allow their personnel to use FR to identify individuals stopped in the field and on reasonable suspicion; to identify deceased, incapacitated, or missing persons; or to identify victims, witnesses, and associates of a suspect. There may also be extenuating circumstances clauses that allow for the use of FR when there is a legitimate threat of serious harm by an individual or a group to themselves or others.

Authorization by Crime Type or Severity

Authorizing FR to be used for all types of crime is more beneficial to the police and potentially society overall than limiting use to serious crimes, but this increases risks of false positives, expands police power, and increases resources to the criminal justice system. Perceptions of or actual misuse or wrongful arrests could negatively shift societal approval. Some less serious types of crime may disproportionately be committed by poor or marginalized groups. Lastly, using FR for all types of crime might change the way detectives prioritize and investigate cases, causing them to become too reliant on video evidence or preferring to pursue cases with video evidence. In contrast, solving serious crimes provides more benefit per crime solved but ignores most crimes and eliminates the potential to capture prolific offenders for lower-level offenses (e.g., a theft suspect who also is wanted for robbery). Additionally, more-serious crimes may include more pressure to solve the case, potentially increasing the rate of false positives leading to wrongful arrests.

Authorization for Field Identification

Using FR in the field for random interactions (e.g., FR on body camera video in near real time or taking a photo on a smartphone) might benefit police by contributing to officer safety (e.g., verifying an individual's identity) and potentially uncovering existing warrants. However, this creates risks of false positives combined with a short time frame to act, which increases chances of false identification and false arrest, increases the number of subjects scanned, and expands police powers to situations in which no crime is suspected. In reasonable-suspicion scenarios, police officers have stronger legal standing to identify an individual and to benefit from improved safety (e.g., ability to maintain distance) and warrant and case clearance. Society and individuals being scanned may benefit if using FR reduces the length of a stop, but there is also a risk of false identification and a short time frame to act, potentially increasing costly mistakes. Not permitting field identification would potentially increase risks to policing in terms of safety and fewer cleared crimes, as well as to individuals detained while other methods of identification are used (e.g., fingerprints, witness showup).

Authorization for Deceased, Incapacitated, and Missing Persons and Victims, Witnesses, and Associates

In these use cases, there are typically benefits to police, society, and individuals scanned because identifying them usually leads to getting them help and improves the amount of information available in an investigation. In the case of the deceased, early identification provides the benefit of family notification and an improved investigation. There is a risk of misidentification, but other methods of identity verification (dental, DNA) might be available for confirmation.[6] Future capabilities might include facial reconstruction or other methods to help identify the deceased. FR for people who are incapacitated or missing carries similar risks and benefits. One key question for this group is whether FR results can be used against the subject being scanned (e.g., the subject has an open warrant). This also applies to victims, witnesses, and associates, but for this group there is also a question of whether they desire to cooperate with law enforcement and whether using FR might reduce the desire to work with law enforcement by individuals being scanned or society at large.

Authorization for Exigent Circumstances

The typical authorized uses of FR can be fairly well defined, which allows for a clearer understanding of the underlying risks. There may be special cases, however. For example, a clear and eminent threat to public safety (e.g., terror attack, active shooter) might mean that certain components of the authorized use are allowed to be bypassed—whether forgoing reasonable suspicion and scanning everyone in an area (e.g., scanning witnesses, associates, and victims) or engaging in real-time FR surveillance in an area. This may also mandate judicial or civilian oversight review or an independent investigation (e.g., after-action report, sentinel event review). The key benefit in this circumstance would be prevention of further harm and apprehension of identified suspects. The major risks in this situation involve enabling use of technology that, no matter how limited the intended circumstantial use, may threaten fundamental rights or lead to false positives. Some of our interviewees argued that opening the door too far is a slippery slope to abuses in the future.

Authorized Uses for Administrative or Secondary Investigative Purposes

There are also key ways that FR can be used to improve overall investigative quality, particularly as it relates to concerns about data integrity, mistaken identity, or witness identification practices. First, FR can be used

[6] A recent study of automated FR on deceased persons found that the best-performing algorithm had a 1.5 percent error rate, but in the condition most similar to a real-world forensic setting (e.g., predeath images as probes and postmortem images as gallery), the best-performing algorithm had a 13.2 percent error rate. Human FR of deceased persons typically has around a 20 percent error rate (Labati et al., 2021).

to ensure that existing data systems are as accurate as possible, most notably by verifying that all photos associated with an individual are actually that person (e.g., accounting for aliases). Especially if there are multiple databases being used, identity verification in the source data is important. Second, many agencies are currently using FR to create photo arrays used in witness identification. This can assist in ensuring that similar individuals are chosen and that features of the photos themselves are similar (e.g., lighting, pose, size), so as to reduce the chances of a false identification, although there are arguments that selecting lineups by similar facial features using FR could increase false identifications (Jackson, 2019). To our knowledge, no research on this method exists.

The box, and the others in this chapter, summarizes the key points of concern regarding when law enforcement is authorized to use FR.

Concerns about whether FR use by law enforcement is justified across a variety of situations:

- Will police use FR for minor crimes?
- Will FR use for field identification lead to an unwanted expansion of government power (e.g., identify everyone at all times) and to potentially riskier interactions?
- Will FR use in noncriminal situations create issues with consent and error? Will the results be used against the subject of the scan (e.g., if they have an open warrant)?

Scope of the Search and Level of Individualized Suspicion

After considering the concerns of various situations for which law enforcement may be authorized to use FR, it is important to determine the types of searches that are acceptable or the criteria under which the scope of the search should be broader or narrower. The level of individualized suspicion is connected to the scope of search because this implicates the degree to which there is a legitimate law enforcement purpose to identify each individual in the search. This is related to various concerns about the different uses of FR. For instance, police officers can legally request identification if they have reasonable suspicion that an individual has committed a crime or the individual is otherwise legally required to provide identification (e.g., while driving). The level of suspicion here is lower because providing identification is fairly unintrusive. If, however, police officers want to conduct a full search, they will ordinarily need to establish the higher burden of probable cause. This can justify searches of a group as well as an individual. For instance, techniques to track an individual (e.g., wiretapping, GPS tracking) require probable cause and a judicial search warrant.[7] Additionally, probable-cause search warrants have allowed geofencing to identify all cell phones present within an area, drawing on the assumption that the suspect possessed a phone at the location being searched. Here, the warrant serves as a practical check on the police to show that the search has a defined purpose for solving a crime. The ability to identify groups of people in an area through FR is analogous to geofencing, although a face that is publicly viewable is arguably less private than a cell phone. Nevertheless, there are questions about the right to anonymity in public, which we discuss later (Perritt, 2021). For instance, police also contend that a geofence result is not probable cause to arrest on its own, similar to arguments made for FR results. A recent case has called this practice into question after a federal judge ruled that it violated the Fourth Amendment (i.e., protection against unreasonable search and seizure) by collecting information on innocent people without evidence that they might be suspects, which has potential implications for FR (Schuppe, 2022).

[7] Getting to this point often means that less intrusive means have been explored or are determined to be likely to fail.

Table 3.1 depicts the different levels of search that can be used with FR technology, along with examples of uses. Verification is a one-to-one matching process that is used to confirm an individual to a known identity. *One-to-one matching* means the software is matching a single probe image to a single gallery image. This is the same process people use to unlock their iPhone, for example, but can also be used to confirm that someone is who they say they are. In the jail setting, verification may be used at booking, release, or transfer to ensure that the appropriate person is being processed. Verification is the scope of FR search that is least subject to concerns. It has a defined purpose—to confirm whether someone is who they are thought to be—which is in line with other methods of biometric identity verification, such as fingerprinting.

Next, identification is most commonly used in investigations. Here, an acquired image of a potential suspect is compared with multiple gallery images (potentially millions or billions). Often, gallery images are mugshots or state driver's license images, although this can also include images scraped from the internet (e.g., via Clearview AI). With one-to-many scope, note that all enrolled images in the gallery are potentially subject to being included in the results returned by the software, which typically has 20 to 100 candidates (some law enforcement documents recommend returning up to 500 candidates; see Rodriguez, undated). Sometimes, the similarity or confidence scores are displayed, and sometimes they are not. Here, FR technology moves law enforcement far beyond manual searching of mugshots. The speed and accuracy involved in searching millions of images makes it much easier for law enforcement to narrow down the list of potential subjects. Concerns here relate to how the one-to-many search includes a vast number of people who realistically have no chance of being the suspect. Since the FR search often is conducted on an entire database, rather than being filtered down first,[8] people are being repeatedly included in FR searches that, while not intrusive, increases the chances of their erroneous inclusion in the set of candidates.

Finally, there are two types of many-to-many searches that are worth distinguishing. First, with aggregation, law enforcement is using archival or stored data to search many unknown faces and compare them to gallery images. In this case, law enforcement might be going back through video footage to identify suspects, victims, witnesses, or individuals from a watchlist. This might also involve using existing data sources, such as video from a body-worn camera, to create a database of faces that can be searched at a later date. The second type of many-to-many searches is the real-time or automated application that continuously scans faces and compares them with the gallery. Real-time FR applications by law enforcement have occurred. For instance, several trials have been evaluated in the UK, wherein surveillance cameras conducted real-time

TABLE 3.1
Scope of Search and Typical Use

FR Type	Scope of Search	Uses
Verification	One-to-one matching	• Identity confirmation • Jail booking and release
Identification or investigation	One-to-many matching	• Investigative lead • Missing or incapacitated person
Aggregation	Many-to-many matching (stored)	• Investigative lead • Database creation • Geofence warrant
Real-time or automated	Many-to-many matching (real time)	• Watchlist or prohibited person • Monitoring crowd or event • Body-worn camera video • Passive surveillance

[8] Filtering could occur based on location, demographics (gender, age, race), or other relevant information, such as a specific list of potential suspects.

FR outside events or in busy areas to identify and intercept individuals on a watchlist (Fussey, Davies, and Innes, 2021). Real-time body-worn-camera FR also exists and has been pilot tested in the United States (Gershgorn, 2020). Real-time FR raises questions about officer decisionmaking and automation that are not well understood. The use of such technology potentially increases reliance on the results of FR being directly used for official law enforcement action, but it also does not tell the officer how to act, leading some to relabel it *assisted* rather than automated FR (Fussey, Davies, and Innes, 2021). Additionally, although being done in the background, this real-time FR will inevitably be used on people for whom officers do not have reasonable suspicion, increasing the scope of law enforcement's ability to identify anyone at any time for any reason (i.e., arbitrarily; see Ferguson, 2021). Moreover, the quality of the data underlying a real-time system is less clear. If there are lags in clearing prior arrests, warrants, or other pivotal information in the underlying data, this could lead to further errors and improper detentions. These many-to-many applications are the clearest examples of concerns about a police surveillance state.

As the scope of the search moves from one-to-one, to one-to-many, to many-to-many applications, the level of individualized suspicion decreases, and FR is then being used more as a dragnet than a focused tool that maximizes the benefit of its use. Additionally, the increased number of searches is not trivial: More searching means that more errors will be made even with the most accurate of algorithms. This also brings up the question of what to do if, in the course of investigating one crime, information related to another crime or other pertinent law enforcement information is revealed (e.g., if someone has a warrant for traffic tickets). Lastly, a move to real-time or automated FR also has implications for how FR is used and the quality and timeliness of the underlying data.

Concerns about whether there a legitimate law enforcement purpose for scanning someone:

- Since FR is a powerful tool, will it be used to scan people in public for no legitimate law enforcement reason or track their movements?
- Will accuracy be reduced as the number of searches increases?
- Could more-expansive searches violate civil rights?

Accuracy: Sources of Error

In general, the broader the scope, the more likely accuracy concerns and their consequent effects are present. Importantly, concerns about accuracy and bias have been central to discussions about whether FR is appropriate for the law enforcement context. Currently, FR results are not commonly admitted as evidence in court, which is one sign of legal acceptability and reliability (e.g., DNA, fingerprints). However, some of our law enforcement interviewees argued that, as a tip or lead-generating tool, FR does not need to be admitted as evidence or even disclosed to the defense. From this perspective, the primary function of FR is to reduce the list of candidates and investigate from there. Since this proposed approach arguably involves a high degree of human involvement, the FR software and its accuracy is a small part of what happens during an investigation. As one interviewee noted, a potential drawback was that high advertised accuracy can provide FR analysts with too much confidence in the results, which raises concerns about automation bias. Most other interviewees, however, stressed that high accuracy across the various FR applications was perceived as critical for FR to be used in the law enforcement setting. As a source of lead generation and an element of investigative decisionmaking, accurate FR results provide benefits to police and society in the form of improving the efficiency of the investigative process (e.g., police are less reliant on public tips for leads), potentially increas-

ing clearance rates, and by appropriately excluding innocent individuals. As a source of identity verification or field identification, accurate FR benefits police by efficiently establishing subject identity, possibly improving officer safety, or ensuring that the right person is released from custody.

Despite substantial improvements in accuracy since even the late 2010s, there are still areas that need to be pursued further. For instance, although NIST testing is the gold standard for determining accuracy across several applications (e.g., one-to-one matching, one-to-many matching, mugshots, webcams), some of our interviewees expressed concerns about the external validity or generalizability of NIST accuracy tests (e.g., pixelated, off-angle, occluded) to many law enforcement settings. This and the lack of clear standards about image quality—although standards and algorithms to assess image quality are in development (NIST, 2022)—are two of the biggest current question marks for accuracy of FR in a law enforcement setting. Moreover, very few agencies do formal accuracy testing themselves, and NIST testing occurs in a more controlled setting. In fact, FR use in law enforcement is a system of human-driven and software-driven steps with implications for the accuracy of the search and the decision to act on the results. For example, the software could rightfully return results suggesting low similarity to any gallery images, but the officers could still review the images and move forward with a lead. The officers may be correct or incorrect, but the rate at which this happens is not typically known. Existing research suggests that human examination of faces is particularly poor for unfamiliar faces and reduces overall accuracy, although training appears to help (White et al., 2015). Other research has shown that FR errors are compounded in human review due to prior identity labels (e.g., the FR result previously suggested a match; see Howard, Rabbitt, and Sirotin, 2020). There is much more research needed examining the role of human accuracy and human-algorithm interaction in law enforcement's use of FR (for example, see Alon-Barkat and Busuioc, 2022; Peeters, 2020; Saxena et al., 2021).

Eyewitnesses, Facial Examiners, and Super Recognizers

An important question for evaluating the accuracy of FR systems is to compare them with the usual human identification (i.e., is FR better than what we have currently?). Here, we can compare FR system accuracy with three types of human identification: eyewitnesses, facial examiners, and super recognizers. Super recognizers are people who have a unique ability to recognize faces and score highly on FR tests, whereas facial examiners are trained and mentored to evaluate and assess facial features, using a scientific approach to facial comparison. Super recognizers have tested as high as 95 percent on some FR tests (Phillips et al., 2018, but there are concerns about whether this testing translates to real-world performance (Ramon, Bobak, and White, 2019). Some of research in this area notes the importance of confidence judgments in assessing the quality of the positive identification and finds that facial examiners are more consistent as a group, suggesting the importance of training and experience (Hahn et al., 2021). Moreover, there remains evidence of cross-race bias effects for super recognizers, just like the general population (Bate et al., 2019). The main equalizer for trained and untrained humans, as well as for FR software, is low-quality and variable image conditions (Hu et al., 2017; Lee et al., 2009).

Operational Workflow and Accuracy

There are a variety of other factors in the FR process beyond image quality that can affect accuracy. Figure 3.2 shows the different elements of the FR search process that potentially affect the accuracy of FR searches. First, the quality of the image that will be used in the search is paramount. The Facial Identification Scientific Working Group (FISWG) has several guidance documents on factors to consider in terms of image acquisition (i.e., capture) and image quality and notes that there will be vendor-specific impacts of image quality.[9]

[9] FISWG is a subcommittee of the Organization of Scientific Area Committees (OSAC), which is part of an initiative by NIST and the U.S. Department of Justice to improve forensic science. OSAC works to support forensic science standards and guidelines and improve the scientific basis for a variety of forensic science fields. See FISWG, undated.

FIGURE 3.2
Factors in the Facial Recognition Search Process That May Affect Accuracy

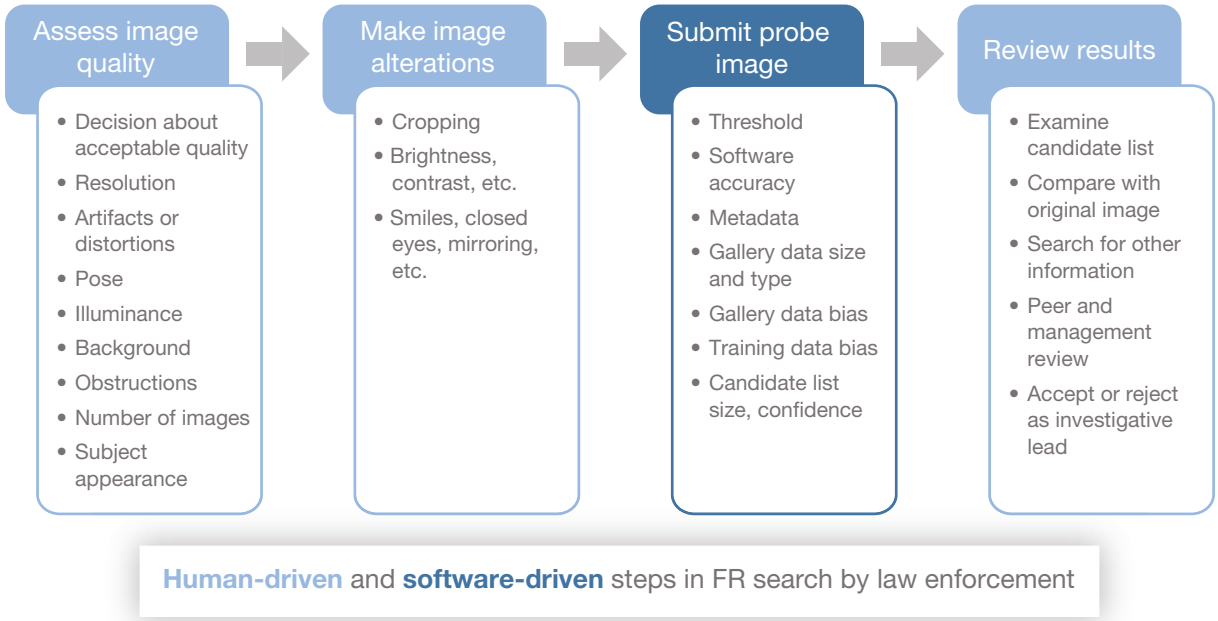

There are a variety of factors that affect whether an image is acceptable for an FR search, including the resolution or pixelation, artifacts or distortions, pose, illuminance or brightness, background, obstructions, whether multiple images will be used, and subject appearance (e.g., expression, makeup, hair, glasses). One especially relevant factor that often implicates multiple image factors is the setting of the image capture—or whether it occurred in a controlled or uncontrolled setting (i.e., "wild" capture). In the law enforcement context, mugshots are from a controlled setting, whereas surveillance footage comes from an uncontrolled setting, meaning that most investigative probe images are uncontrolled and therefore of lower quality. There are efforts underway to improve standardization and automated testing of image quality (Grother, Hom, et al., 2022; Grother, Ngan, and Hanaoka, 2022), as well as efforts to deal with some of these image-quality issues from uncontrolled settings by training the algorithm with data augmentation methods (e.g., Fredj, Bouguezzi, and Souani, 2021). Having more-uncontrolled images may also lead to more attempts at image alterations, which are allowed under some agencies' use policies. Some of our interviewees, however, had reservations about making image alterations beyond cropping and basic image edits, such as brightness, contrast, and sharpness. Making more-drastic changes, such as closing mouths, opening eyes, or mirroring images, was viewed as problematic by some of our interviewees because it can distort features and make machine and human matching more erroneous (e.g., one interviewee said that there is "not much evidence to show it is safe"). Other interviewees thought that altered images could be used in the search but that search results should be compared back with the original unedited image.

Once the image is submitted to the FR software, there several software-driven factors can affect accuracy: the threshold setting, the quality of the algorithm, the size of the candidate list, the gallery size and type, gallery data bias, and training data bias. Setting a threshold is intended to balance the trade-offs between false positives and false negatives such that the software only returns results that meet a certain level of confidence. Setting this threshold depends somewhat on the application and the level of concern about false positives. For instance, in automated FR applications that have limited capacity for review, the threshold should likely be set higher to strictly limit the number of false positives. However, in cases that involve a

clear process of human review, the threshold might be set lower, although there is a trade-off if this requires more human review of images. Some interviewees thought that it was acceptable not to use a threshold at all for applications with human review of a ranked list (i.e., what NIST refers to as the *investigation model*; see Grother, Ngan, and Hanaoka, 2019a), while others advocated for setting the threshold at a level that is either the default for the software or based on some expectation of results, such as the number of expected false positives. No-threshold searches that return ranked lists tend to have a lower miss rate (i.e., the false-negative identification rate) but produce weaker results because a top result can be returned that would not have met the threshold had it been applied (Grother, Ngan, and Quinn, 2017). This is because although the top-ranked candidate may be correct with no threshold, the similarity score might be very low.

Algorithm quality varies based on the application, image conditions, representativeness of the training data, and size of the database being searched. Many of the current top-performing algorithms perform the best for verification (one-to-one) applications but perform somewhat similarly across image conditions (e.g., mugshot photos versus webcam photos) and in one-to-many applications. Higher accuracy algorithms that perform well with challenging image characteristics (e.g., webcam, kiosk, or side-profile images) are more likely to return useful results.

Searches can also be conducted by filtering or binning on metadata or specific characteristics, such as race, age, and gender (Sankaran et al., 2018).[10] If the search metadata are correct and the metadata for gallery images are accurate, the search is more likely to be accurate. If inaccurate, metadata filtering or binning can exclude the true match from being included in the search. Although some law enforcement agencies use metadata, this does not appear to be part of standard search practices. FISWG notes that a key consideration is to ensure metadata consistency (FISWG, 2019). However, using metadata can help reduce low-quality search results, particularly when a threshold is being used.

Concerns about how accurate FR is in the law enforcement setting:

- Will a lack of common practices or standards affect understanding of accuracy at scale or of the impacts of different practices (e.g., image quality, image alterations, threshold settings)?
- Are training and gallery data high quality, and how is accuracy complicated by the true positive not being included in the gallery (testing always includes correct answer)?
- What is the additive or interactive effect of human and algorithm error on overall accuracy?

Next, the size and type of gallery data have implications for accuracy. In general, gallery data are trending toward being more expansive and including multiple types of image sources, such as mugshots and driver's license photos. Having multiple sources of gallery data can affect accuracy as well because images are likely standardized differently (e.g., lighting, resolution). Larger gallery databases also generally produce lower accuracy due to fewer rank-one true positives (Grother, Ngan, and Hanaoka, 2019a).[11]

[10] In the FR setting, metadata can include a number of tagged pieces of information related to the person in the image, such as race, age, or gender, expression, or features of the image itself, such as illuminance, yaw, pitch, and face size.

[11] This means that fewer actual matches are being ranked as the highest-ranked image in terms of similarity score.

Accuracy: Sources of Bias

Concerns about the potential for differential impact or bias across different sociodemographic groups are central to considerations for authorizing and implementing FR systems for law enforcement use. Critically, there is concern that bias is baked in to the various elements of the FR system, most notably into the FR software itself (training data and algorithm performance), in the image data sources being used (gallery and probe), and in the decisionmaking process and actions taken by the person in the loop (confirmation bias, context bias, implicit bias). Thus, given the way in which FR is currently used in law enforcement, the accuracy of the FR system is not solely determined by the accuracy of the FR software; whether the police arrive at the correct decision also matters.

In Table 3.2, we highlight some of the more prominent ways in which bias can seep into the FR system in a law enforcement scenario. First, if the data that the FR software was trained on do not reflect the population being examined, accuracy issues will result. This is more problematic if one particular group is more affected by this issue. Relatedly, the algorithm may perform differently for certain faces or certain types of images such that it affects one group more than others. Next, the data being used by the FR software may have issues that contribute to bias, whether in the gallery data or in the probe images. These sources of bias tend to reflect broader policing or societal biases being incorporated into FR processes, such as whose images are more likely to be included in either gallery or probe images. If more searches are conducted for one demographic group as probes or if there are disproportionate images of one demographic group in the gallery data, then that group will have more opportunities to be affected by the use of FR by law enforcement and might be disproportionately affected by FR error rates. Finally, the FR process always involves a human decisionmaker who may have implicit or explicit biases that affect their choices. We know little about how these human decisions affect overall FR accuracy or policing actions that occur as part of the FR process.

TABLE 3.2
Potential Sources of Bias in Facial Recognition Search

Source	Potential Type of Bias
FR software	Training data • Not reflective of population • Include image features not present or different from gallery data (e.g., more-illuminated images) Algorithm performance • Systematically performs worse for some faces (e.g., darker faces) because of training data and decisions made by developers
Image data	Gallery data • Not reflective of population • Contains preexisting biases (e.g., disparities in arrests) Probe images • Potentially more likely to reflect crimes committed in public • Potential disparities in CCTV placement
Human decisionmaking	Image submission • Requires deciding which images are of acceptable quality to submit • FR results and facial comparison • Known issues with cross-race facial identification that factor into the comparisons made by human reviewers (e.g., unfamiliar faces) • Actions taken • Deciding whether to continue investigation • Deciding whether to examine other metadata to justify further investigation (e.g., criminal history) • Deciding whether to use legal authority to interact with person (e.g., house visit, stop and identify)

Concerns about whether FR systems are biased against certain demographic groups:

- Are training data reflective of the population?
- What are the effects of gallery data representativeness and demographic clustering?
- What are the effects of a higher algorithm error for certain demographic groups, and how do demographic error rates interact (e.g., older Black females)?
- What is the disparate impact due to source of data (e.g., placement of public CCTV cameras) and use of supplemental information (e.g., criminal history) to guide FR search decisions by human operators?
- What are the effects of human bias for other races and unfamiliar faces?

Data Sources: Training Data and Gallery Data

Closely related to concerns about accuracy are concerns about data sources used in FR. The main concerns focus on the training data, the gallery data, and how gallery data are obtained. Concerns about the training data relate to the size of the training database, the representativeness of the training data, and the quality of the training data. With regard to the size of the training database, the key point is that larger training datasets are likely to contain more features on which the algorithm can be trained, which improves accuracy. This is related to the representativeness of the training data because the existence of more features means that the algorithm is more likely to recognize various features in the matching context. For instance, training data that mostly include lighter faces mean less accurate results for darker faces. The same goes for gender, age, and numerous other features. The representativeness of features also applies to image quality and pose, whereby training data that include a variety of resolutions, poses, or occlusions (e.g., masks) are better suited to be accurate when those features are part of probe images being used in matching. For instance, most, if not all, algorithms submitted to NIST testing perform worse for "in the wild" photos, which likely contain features that are less likely to be captured in training data. Lastly, the quality of training data also appears to matter, as cleaned and labeled training data seem to show accuracy improvements (Anwarul and Dahiya, 2020).

Gallery data sources include police-generated images, including mugshot data, other government identification (e.g., driver's license and passport photos), and third-party image data, which might be scraped from the internet and social media sites. Concerns about gallery data include the size of the database, the representativeness of the database, and how gallery data are obtained. The size of the gallery database is directly related to the scope of search and individualized suspicion discussion above, as well as concerns about accuracy (Grother, Ngan, and Hanaoka, 2019a). The larger the gallery, the likelier that more nonsuspects will show up in FR searches and the lower the accuracy. Moreover, larger databases mean that more people will be matched on a multitude of demographic features, potentially contributing to bias due to demographic clustering or combinations of demographic features (e.g., Howard, Sirotin, and Vemury, 2019, found the highest error rates for older white men). Larger databases also raise concerns about consent. When people get driver's licenses or passports, the implication that the government can use their image in subsequent FR searches conflicts with expectations of privacy and anonymity.

Next, concerns about representativeness of gallery data reflect whether the images in the gallery database reflect the general population. As the subjects of an FR search, overrepresented groups are a greater risk of being included in candidate lists than their representation in the general population. This is a particular concern for mugshot gallery data, since these data are likely to have baked in racial disparities that exist in arrests and the criminal justice system. There are also concerns over whether arrests that do not result in charges or conviction should remain in mugshot gallery data or whether they should be purged. Finally, the

use of third-party gallery data increases concern about privacy and database size. The prominent example is Clearview AI, which has scraped billions of images from the internet and social media. Data privacy concerns have landed the company in court, but many law enforcement agencies still use the software (Hatmaker, 2022).[12] If the data source is not legitimately obtained, the implications for law enforcement use are unclear. Moreover, images scraped from the internet may, on their own, be considered illegitimate and still need to be authenticated—whether the image is real and whether the identity can be verified. Questions of the legitimacy of an image create issues for use in investigations. Although people are able to find and remove their images from the internet, when companies collect these data, they remain at a level of permanence that prevents individuals from being in control of their image.

A key question for data sources, particularly when it comes to size, is the exact trade-offs between expansive gallery data (e.g., Clearview AI) and limited gallery data (e.g., mugshots). In both cases, there is a chance that the real suspect's image will not be in the gallery data. This may be less likely in the expansive case, but FR use in law enforcement is always in an open-set context (e.g., probe images may include identities not in the gallery data). More-expansive data reduce the risk that the best result according to similarity scoring is incorrect. However, with more-limited data sources and those generated by law enforcement for its own purposes, there is much less of a "creepy" factor (i.e., that people's images are being used without their knowledge and without clear purpose)—and less expansion of police power.

Concerns about whether the underlying data are flawed:

- Are training data sufficiently representative of not only demographics but variation in image quality (both of which affect accuracy)?
- What are the issues surrounding privacy and consent of being included in gallery data?

What Happens After the Facial Recognition Search?

Of key concern is what happens with the results of FR searches. This is where the less invasive image search connects to actual law enforcement activity, which could be very invasive. All of our law enforcement interviewees and many existing FR policies stressed that an FR search result can be used only as lead generation and is not considered an actual identification. The relevant analogies include manual searching mugshot photos and image releases of "persons of interest" or "wanted" suspects to the media to solicit public tips and information. Many of our law enforcement interviewees considered the use of FR to be more accurate and efficient than these other alternatives, but they are still viable in the event there is no useful FR result. Additionally, in describing known cases of FR misuse or FR leading to wrongful arrest, law enforcement interviewees stressed that it was officer decisions that happened after the FR result, rather than the FR result itself, that was problematic in those cases.

Actions taken after the result of an FR search depend on the FR application being used, whether it is verification, identification, or real-time surveillance. In the identification setting, our interviewees noted that detectives typically should not go directly from receiving the FR results to making direct contact with the

[12] Other countries have banned or fined Clearview AI due to privacy violations, including Australia, Italy, Greece, France, Canada, and the United Kingdom (GRC World Forums, 2022).

subject.[13] In this application, the first step will be to look for any other supportive or disqualifying information, including, but not limited to, whether the subject has a twin, whether the subject was out of town or already in custody at the time of the crime, the subject's prior record, whether the subject lives near the crime, and other forms of background research that would potentially exclude the subject. If this effort points to the same person, the detective would still likely need more information to establish probable cause to arrest. A common way this might occur is through eyewitness identification. A major concern when eyewitness identification is combined with FR is whether the eyewitness is made aware of the FR or any computer matching at all, which would be suggestive and problematic because the eyewitness is likely to assume that FR has found a match. The usual eyewitness concerns and considerations still apply (e.g., showup, lineup, photo-array procedures, and a present defense counsel). Another potential action is to attempt to make contact with the subject and question them. Here, the officer may use FR results to attempt to solicit a confession, although this carries a risk of coercing a false confession.

In the real-time application, there is a much more direct connection between FR results and deciding to act. Here, the predicate for conducting an FR scan might be to look for prohibited persons or known offenders with warrants. In these cases, the officer needs to make a decision to stop and identify the individual or to continue observing or conducting surveillance. Tests of this application have shown fairly low rates of both the perceived credibility of watchlist matches and even lower rates of confirmed identity matches after a stop is conducted (Fussey, Davies, and Innes, 2021),[14] which calls into question the utility of this approach. The need to stop and identify someone to confirm the FR match potentially increases confrontational interactions as well.

Concerns about what happens after police use FR software:

- Will uncritical acceptance of FR results by officers lead to false arrests (i.e., lack of probable cause)?
- How are FR results handled with eyewitnesses, and does FR increase similarity enough to contribute to misidentifications?
- Will FR use ever require or prompt an officer to act immediately on FR results (e.g., a field stop)?

Legal and Ethical Implications of Facial Recognition Use

The extent of FR use, the way it is used and the accuracy of that approach, and how FR is acted on all have implications for various legal and ethical concerns—that is, concerns about privacy, human rights, civil rights, and data security. Part of the challenge in understanding these concerns as they relate to FR is that although there are analogues for how these concerns might play out, the use of this technology is new enough that much of the discussion still falls in the realm of extrapolation. In this section, we discuss these various concerns and how they relate to current and potential future uses of FR.

[13] However, this has happened, and there is no guarantee that it will not happen again. Preventing this from occurring might require several steps, such as legislation, policy, training, and checks and balances (e.g., internal auditing, external auditing).

[14] Ongoing tests in London do not appear to have shown improvement (Rollet, 2022).

Human Rights and Civil Rights Considerations

FR is distinct from other common methods of biometric identification in its ability to identify people from a distance and over large areas—including when they are engaged in political, speech-related, and relatively private activities. These features of FR implicate values protected by the First and Fourth Amendments most directly. Its use in criminal investigations has implications for whether and how police and the courts handle FR results as evidence and dealing with values called out in the Sixth and Fourteenth Amendments, and this use brings up key issues of disclosure (e.g., Brady) and the reliability of scientific evidence or expert testimony (e.g., Daubert and Frye).[15] These issues have been written about in depth elsewhere (e.g., Ferguson, 2021; Perritt, 2021), but we cover them briefly in this section to identify the main concerns. Although some of the concerns discussed here may eventually find their way into the courts, policymakers may approach this subject with an emphasis on preventing the most-likely and most-significant harms to the aforementioned constitutional rights in the first place. Relying on litigation and resulting precedent to set policy risks tolerating harms to constitutional and human rights but also delegitimizes FR and further erodes trust in law enforcement. Like many law enforcement practices, one misuse, false arrest, or rights violation might outweigh the benefits of many legitimate arrests made using FR, given the weight the U.S. justice system affords to the presumption of innocence and related values.

First Amendment Concerns: Freedom of Assembly, Association, and Speech

Potential First Amendment concerns about freedom of assembly, association, and speech are important to civil rights groups in terms of law enforcement use of FR. Notably, many of our interviewees, even those in law enforcement, were concerned about the potential chilling effect on First Amendment–protected activity, particularly protests and other forms of public speech, as individuals may fear being identified and targeted by law enforcement on the basis of their viewpoints. Indeed, some police departments have come under scrutiny for using FR on protestors, leading to concerns about retaliation, due-process violations, and First Amendment violations.[16] This also recalls the previous discussion on the scope of the search or the level of individualized suspicion. Searching multiple people or everyone in a crowd is a many-to-many application of FR, which is not typically defined as an acceptable use in policy. Moreover, such use conflicts with the notion of the right to the privacy of one's associations that is central to rulings in such cases as *NAACP v. Alabama ex rel. Patterson* (1958) and *Shelton v. Tucker* (1960), which prevented disclosure of group affiliations or associations (Nguyen, 2002). The concern of a First Amendment chilling effect is certainly stronger if participation in protests or other group events, albeit in public, comes with the expectation that police will engage in many-to-many FR matching. The ability of existing FR algorithms to quickly scan large databases further contributes to this concern. For instance, some law enforcement agencies currently have access to statewide or multistate driver's license images, and others use Clearview AI's database of billions of public images scraped from the internet. Notably, the strength of this concern likely varies depending on how law enforcement is using FR in this scenario. If used to identify a group of individuals engaged in criminal activity within a large group, concern is lesser than if law enforcement uses FR to identify and contact all people in the area, create a watchlist, or take some other action that is more invasive and includes more people.

[15] Simply put, the First Amendment protects the freedom of speech, the press, assembly, and religion and the right to petition the government. The Fourth Amendment protects people from unreasonable search and seizure by the government. The Sixth Amendment includes guarantees for the right to confront one's accusers. Brady material is evidence that would be favorable for the defense and therefore must be disclosed by the prosecution. Daubert and Frye pertain to the admissibility of scientific testimony and the reliability of scientific or forensic procedures that were used to obtain evidence.

[16] At least three police departments in Florida submitted protestors' images to the state's Face Analysis Comparison & Examination System (FACES), which includes both driver's license photos and mugshot photos (Simpson and Freeman, 2021).

Fourth Amendment Concerns

FR also poses concerns about privacy and the potential for government overreach in breaching privacy expectations. In the U.S. Constitution, privacy is protected primarily through the Fourth Amendment guarantee against "unreasonable searches and seizures." Whether a law enforcement technique constitutes a search, and is therefore subject to the reasonableness standard, depends in part on whether it violates "reasonable expectations of privacy" (*Katz v. United States*, 1967). Fourth Amendment jurisprudence has long held that government's visual public surveillance as such is not considered to be a search and is therefore not subject to reasonableness requirements, such as individualized suspicion or a warrant. FR might not, therefore, implicate the Fourth Amendment's constitutional protections. However, some FR uses go considerably beyond technologically unaided visual surveillance and may present a closer constitutional question under some specific circumstances. For example, as the Supreme Court suggested in *Carpenter v. United States* (2018), prolonged and continuous surveillance accomplished with the use of a new technology (in that case, with the use of cell-site data) may breach reasonable expectations of privacy in ways that more-limited surveillance did not.[17] This suggests that some uses of FR may, in some hypothetical case, be sufficiently different from plain visual surveillance to lead courts to conclude that reasonable expectations of privacy have become breached through. Nonetheless, absent any legislation regulating FR, it is unclear whether constitutional law alone can supply any firm legal limits to its uses (Ferguson, 2021).

Furthermore, if FR does not itself require some level of individualized suspicion under the Fourth Amendment, there is a further question of whether information gathered through FR can itself create reasonable suspicion or probable cause for further searches or seizures. The answer, it would seem, must depend on the specifics of the FR use case.

Sixth Amendment Concerns: The Right to Confront an Accuser

The Sixth Amendment provides that defendants have the right to confront their accusers or witnesses against them, among other rights. Should prosecutors seek to introduce FR evidence in court, criminal defendants could likely raise reasonable arguments that such evidence is at odds with their right to confront "witnesses" against them. Although FR identification or verification is not a witness per se, the introduction of FR results may well raise confrontation rights concerns. In general, constitutional case law bars the introduction into evidence of statements made outside of the courtroom by a witness who does not appear at trial, if these statements are "testimonial" (*Crawford v. Washington*, 2004)—i.e., serve primarily to establish "past events potentially relevant to later criminal prosecution" (*Davis v. Washington*, 2006). The Supreme Court has previously held that forensic laboratory reports, which contain an expert's statement about the results of the lab test, are "testimonial," and would thus violate the defendant's confrontation rights if introduced in court without that expert's appearance (*Melendez-Diaz v. Massachusetts*, 2009; *Bullcoming v. New Mexico*, 2011). Similarly, introducing FR evidence, which could contain statements describing the results of FR analysis by the individual that analyzed it, might also violate defendants' confrontation rights (Celentino, 2016).

Fourteenth Amendment Concerns

The Fourteenth Amendment provides equal protection under the law and protects against violation of due process at the state level. The use of FR might also present due-process concerns—even though, as with the other questions touching on constitutional rights, constitutional law at present does not necessarily provide firm limits on the use of the technology. One way that FR searches can establish probable cause to arrest is through subsequent eyewitness identification. Although the standard practices for eyewitness identifica-

[17] Some justices expressed a similar concern in *United States v. Jones* (2012), although the holding that placing a GPS device on a car qualifies as a search was based on a different rationale.

tion also apply here, a key question is whether additional suggestiveness is presented by the use of FR—in particular, if the eyewitness is informed that FR or computer matching was used at all. Eyewitnesses may be more inclined to believe that sophisticated technology is likely to have made a correct identification. Therefore, they may be more likely to make an identification of a suspect, regardless of FR's actual accuracy. At the same time, the revealed use of FR is unlikely to amount to the kind of "unnecessarily suggestive" procedure arranged by law enforcement, such that it might raise due-process issues at subsequent criminal proceedings.[18] Although legal arguments that challenge eyewitness identification of suspects initially identified through FR may theoretically succeed, at present, constitutional law does not suggest firm limits to the use of FR in this way (see, e.g., Jackson, 2019).

Another way FR technology affects eyewitness identification is through the selection of similar-looking photo array images, controlling for appearance, lighting, or other features that may inappropriately contribute to the suggestiveness of a particular image. It is not clear how a similar-looking lineup might affect an eyewitness's decision to name any subject or how it might affect their confidence. On the one hand, a lineup with suggestive features removed and more-similar faces could lead to more-accurate identifications, as the eyewitnesses are clearly able to select the faces they recognize. On the other hand, an eyewitness who feels compelled to pick someone from the lineup might be adequately satisfied with any of the faces, leading to eyewitness errors that also have higher confidence. Jackson (2019) explores this idea, but more research is needed, as is more research on human-machine interactions and how they affect decisionmaking (Peeters, 2020).

A second due-process concern FR raises is related to exculpatory information. *Brady v. Maryland* (1963) is the landmark case that mandates the disclosure of evidence that is exculpatory and material by the prosecution to the defense. Some attorneys argue that the use of FR could potentially provide *exculpatory* evidence for a defendant, most notably whether there are other individuals of comparable similarity identified by the FR algorithm. A key question is the materiality of the evidence to the defendant's case, which is up to interpretation of the courts (Goldberg, 2020). Here, the court is supposed to consider whether the defendant received a fair trial without the evidence included and whether it is reasonable that a different verdict would have resulted had the evidence been included (*Turner v. United States*, 1919; Goldberg, 2020).

One of the few cases to cover this specific issue, although under state law, *Lynch v. State* (2018), involved a drug-distribution case wherein investigators used FR. Prosecutors never turned over several potential matches to a cell phone image. The probe image was also blurry, taken from a side angle, and the confidence of the search results was low. Thomas Lynch requested that the state disclose other images from the search, but the state refused. The Brady argument was denied because Lynch could not show that the other images were of similar resemblance. Other key elements were that detectives in the case had also not seen the other images; only the crime analyst had. There was also an eyewitness who identified Lynch as the suspect, but it is not clear whether this was the only other evidence connecting Lynch to the crime.

Arguably, the confidence ratings of the FR software would be Brady evidence, similar to how eyewitness confidence is potentially exculpatory. The quality of the probe image, the enrolled database, the software algorithm, and other key parts of the FR process could be potentially important as well. The combination of a FR search leading to an eyewitness identification as the sole basis of probable cause could compound errors in both parts of the identification that led to an arrest.

Another potential legal issue for FR is whether FR evidence shows what it claims to show—for instance, whether the person in the probe image can actually be linked to the crime committed. This means connecting the probe image to the crime scene, the probe image to the set of FR matches, and the matches and the

[18] In *Perry v. New Hampshire* (2012), the Supreme Court decided that when identification is "procured under unnecessarily suggestive circumstances arranged by law enforcement," the Due Process Clause calls for a preliminary judicial inquiry into the reliability of that eyewitness identification.

defendant. Some cases (e.g., *People v. Beckley*, 2010) have rejected photographs from the internet because they lacked authentication or verification that the content from the site was reliable.

Concerns about whether FR use bypasses civil rights protections:

- Could FR be used to stifle free speech, religion, or assembly?
- Will FR use lead to unreasonable searches or violate expectations of anonymity (e.g., pervasive surveillance)?
- Does nondisclosure of FR impede the right to confront witnesses or to disclose potentially exculpatory evidence to the defense?
- What are the effects of the courts not having paid much attention to FR use so far and legislation lagging behind use?

Privacy, Anonymity, Enrollment, Notice, and Consent

The use of FR, both broadly and in the law enforcement context, carries substantial implications for privacy and, more specifically, anonymity, which lie beyond the sphere of rights protected by settled constitutional law. As seen in recent public outrage about the Internal Revenue Service using FR to verify identity on tax returns (Metz, 2022), public expectations of privacy led to pushback on government use. The mass collection of biometric data by the government, in this case facial imagery, is something that many people are not comfortable with, even though many states' FR systems already contain driver's licenses and other forms of government identification. Moreover, concerns about privacy are further exemplified by the use of FR surveillance in other countries. China's social-credit-score system, for instance, can identify people in public and penalize them for doing actions such as smoking in nonsmoking areas. The use of FR during the war in Ukraine highlights how pervasive and accessible the software is to the public. There are even examples of private citizens using FR to identify police officers (K. Hill, 2020). In this section, we discuss concerns about privacy and anonymity and more-specific concerns about enrollment, notice, and consent.

FR in the law enforcement context involves several sources of data (e.g., sources of gallery images and probe images). A key question for how FR affects privacy is the extent to which facial image data already exist and the implications of FR being layered on top of those existing data. That is, FR provides powerful analytic capabilities to identify faces in images and can be used to connect identities to those facial images. FR use by law enforcement may ultimately lead to coercive action against someone (e.g., arrest). Moreover, FR presents a potential motivation to collect and store facial image data for the later purpose of using FR. In a criminal investigation, the probe image can be sourced from private (e.g., ATMs, security cameras) or governmental video sources at or near the scene of the crime. These sources do not exist for the purposes of using FR, but FR can be used on images captured from video. Although FR does not change the existence of these data, it does increase the amount of information that can be extracted from video—and potentially at large scales. Thus, FR increases law enforcement's ability to collect information on faces in public and information on specific individuals and their locations, which makes tracking and surveillance applications more efficient and increases risks of pervasiveness. If these data are kept long term, there would be historical records for law enforcement to access at any time.

As it relates to gallery images, the use of more-expansive data necessarily reduces privacy to some extent, even if those images exist for other legitimate purposes. For instance, combining mugshots with driver's license and other identification images with publicly scraped images from the internet exposes more people

to FR searches, likely without their knowledge that their image data would be used for FR purposes. Frequent searches on large databases implicates more people. Moreover, arbitrary searching of existing video creates more privacy concerns than searching with a more specific purpose, such as a criminal investigation. Some agencies do not require reasonable suspicion to conduct an FR search, which increases this risk. Another concern relates to law enforcement actions to collect or improve FR data, whether by making stops with people to get updated or higher-quality images or collecting images on networks of associates with no clear investigative purpose.

With the expansion of FR into public use, discussions around expectations of privacy are likely to continue. When it comes to law enforcement's use of FR, a question closely related to privacy is what to make of public behavior and expectations of anonymity. Although a person might not have an expectation of privacy if their face is captured on video in public, they do still have an expectation of anonymity, especially to law enforcement. Guidelines around when people are required to identify themselves to police likely do not cover most people's public existence. Thus, searches of broader scope increase concerns about violating expectations of anonymity in public.

Another privacy concern is how enrollment and candidate lists are handled, as well as whether there are notice and consent procedures in place. Since search results may contain as many as 500 images, whether those candidate lists are retained, for how long, and for what purpose presents privacy concerns. Having your image included and retained in a list of potential matches is more invasive than just having your image included in a search. Next, most people are not aware that their images are included in gallery data, and there are questions about whether there should be automatic notice. There are questions about whether people should be notified of how many times their images are included in FR searches and how often their images have shown up in a candidate list. This may coincide with allowing people to review the images associated with their identity and allowing them to request that some images be removed or to opt-out entirely. Retaining more data, including candidate lists, and not having clear protocols for deleting data increases privacy concerns.

Concerns about whether FR use by law enforcement violates people's privacy:

- What are the potential effects of ever-expanding FR databases with images that were not collected for any government purpose?
- What are the potential effects of more people being included in searches and search results being retained?
- What are the potential effects of data being kept indefinitely?

Data Sharing

The ways in which FR data are currently being used point to ever-expanding gallery databases, wherein law enforcement agencies share their images with other agencies under a memorandum of understanding. One such example is Florida's FACES, which contains more than 33 million images of driver's licenses and law enforcement images and is managed by the Pinellas County Sheriff's Office. Agencies across Florida have access to this system and conduct around 8,000 searches per month (Perpetual Line-Up, 2016b). The FBI and the Facial Analysis, Comparison, and Evaluation (FACE) Services Unit maintain a database of 411 million images, including mugshots, visa photos, and driver's licenses and other identification cards (Perpetual Line-Up, 2016a). Shared systems such as these benefit law enforcement by providing more-complete data but the

variation that exists across agencies presents concerns in terms of data access, data quality, and data security and retention.

Data Access

When it comes to data access, some agencies have dedicated facial analysis units, while others allow patrol officers to use FR in the field. The different types of users present different implications for oversight of the entire system. Although agencies are accountable for actions of their own personnel, it is much more challenging to control system misuse with individual agency access. An alternative to individual agency access (i.e., open system) is to centralize FR to a statewide agency (i.e., closed system). In Arizona, the Department of Public Safety conducts all searches for state and local law enforcement agencies, triaging cases by severity and public importance. This has the benefit of having a single source conducting FR searches, meaning that FR use is more easily tracked and authorized users can develop more-specialized expertise.

Data Quality

Next, data-quality variation is a potential concern. There is little known about how image quality, the most important data type, varies across departments. Even variation in how mugshots are taken can affect search results across the system, leading to concerns about "garbage in, garbage out." Alterations to images may contribute to this as well. Moreover, some states require purging from mugshot data those who are not charged or not convicted, while other states do not. In states without such a policy, local departments likely vary on their policies. Lastly, departments may vary in their policies on whether and how long candidate lists are retained. Departments that also shared match or candidate list data may be contributing error into the shared system.

Data Security and Retention

Last, data security can be more of a risk in open than in closed FR systems, particularly open systems with more authorized users. As the number of users increases, the risk of a data breach or misuse of data (e.g., personal use) increases. Additionally, the collection and retention of images into these databases requires scrutiny in terms of what gets added and how long it is retained. Although law enforcement agencies are used to handling sensitive data, the existence of large, shared FR data systems requires an additional set of policies on handling facial data.

Concerns about whether FR data are being handled appropriately by law enforcement:

- What are the policies for oversight and data quality in FR data-sharing networks that include millions of images from multiple states and multiple sources?

Unintended Consequences?

With the implementation of any new technology, there may be unintended consequences that occur internally in relation to the technology use (e.g., practices for solving crimes) or external to the intended outcomes of the technology (e.g., community expectations and trust). The internal unintended consequences may come about due to officers' use of the technology and their assumptions about what it does. First, regular use of FR for investigations might lead to overreliance on FR and its findings in the course of an investigation. That is, law enforcement might come to overvalue video evidence relative to other forms of evidence, letting the presence of video evidence drive the effort put into cases with video or letting the results of FR drive the investigative scope. For instance, if the FR search does not return a match, officers might just assume that

the investigation is unlikely to be successful and move on to another case. Alternatively, an FR match might dictate whom investigators focus on, leading them to overlook other potential suspects or types of information that could aid the investigation. Unintentionally, continued reliance on FR could lead to diminishing investigative skills over time. If a large part of the job is using FR, the skills and experience required for other investigative skills might not be as well developed. This was brought up as a concern in one of our interviews with law enforcement.

Another source of potential unintended consequences relates to community views about FR use by law enforcement. Similar to other forensic technologies, FR may contribute to a "*CSI* effect" whereby the results are viewed as definitive by the public (referring to the television show and others like it that could inflate expectations about forensic evidence). This might also lead people to assume that FR and video evidence are ubiquitous and could lead to them coming forward and cooperating with the police less. This is similar to how the ShotSpotter gunshot-detection system has been associated with reduced calls for shots fired because people assume that the police are aware of the shots or are already responding before they are able to call (Mares and Blackburn, 2012, 2021). Relatedly, community expectations about FR may also provide police with another source of deception in interrogations. If people believe that FR is ubiquitous and accurate, and if the detectives claim to have FR evidence (true or not), this may contribute to a confession. This could be a real confession, but this approach also runs the risk of increasing false confessions.

Lastly, although FR is intended to be used to solve crime, many police uses of technology can create a distance between officers and the community. In the extreme case that identification through FR is highly pervasive, it could discourage officers from getting to know the community because they will not need the community to identify suspects—replacing technology for street knowledge. Alternatively, the community might believe that all interactions with officers are simply attempts to capture their image. Concerns about false confessions, false arrests, and increased distance can hinder community trust, no matter how likely these things are.

Concerns about the unintended consequences that might come from FR use:

- Could FR use change police investigations by overemphasizing video evidence or diminishing other investigative skills to solve crimes?
- Could the community view FR results as definitive or reduce reporting or cooperation because they assume video evidence exists?
- Could FR use reduce community engagement due to officer overreliance on FR or due to community mistrust?

CHAPTER 4

Developing a Policy Road Map to Help Navigate the Landscape of Concerns About Facial Recognition

Given the diversity and variation in levels and types of concern about FR, developing policies that balance the risks and benefits of FR use by law enforcement will likely be a location- and jurisdiction-specific endeavor. In this chapter, we present a road map for developing that path—looking at the different components that could be included in a policy and management structure for FR and how requirements for transparency and oversight that are responsive to different levels of trust in police agencies may be addressed.

Finding a Better Path

By drawing from existing resources, which have often focused on specific concerns or come from particular disciplinary or stakeholder perspectives, the road map brings together the full range of options and explores the particular concerns and risks. The goal is to inform local efforts in shaping legal, policy, and oversight framework so it is focused appropriately on the risks stakeholders want to manage and is balanced to preserve as much of the potential benefit of using the technology as possible.

Doing so will require accepting that concerns originating from all perspectives are legitimate. Police organizations may have concerns about how requirements for public oversight and transparency will affect their work—and might not always trust that groups opposing their use of the technology are operating in good faith. Civil society and other groups might not trust that police agencies will faithfully abide by restrictions on FR use unless transparency forces them to do so. If mutual mistrust dominates, efforts to develop a reasonable and sustainable approach for governing police use of FR will likely fail. In that case, the potential for pendulum swings between permissiveness and restriction will continue—as high-profile events or shifts in public expectations make the arguments of one side seem more persuasive than the other.

Finding a better path requires accepting the reality of trust concerns and building in policy and procedures that seek to address them—even if doing so is much more complex than simpler solutions in which one viewpoint dominates and the concerns of others are ignored. Doing so also may require designing a system—including the technology supporting FR, the policy around it, and oversight—that does not give individual users or their agencies the benefit of the doubt that they will use it appropriately. Individual users would have to accept that as necessary for being given access to powerful technology tools in a high-risk setting. A system that does not trust its users needs to both build in incentives or drivers to encourage appropriate use and facilitate later verification and auditing—even at the cost of some efficiency or potential effectiveness.

The following sections map out the range of options for finding that path between regulation and efficiency, drawing on existing literature and the discussions from our expert interviews. The options fall broadly into the categories shown in Figure 4.1, with community approval being an overarching factor for the other categories. Different options in each category are responsive to different concerns about risks affecting the groups discussed in Chapter 3. In some cases, trade-offs—even within single policy options—are not straightforward and can pull in different directions with respect to managing risk and seeking benefit

FIGURE 4.1
Building a Road Map Through the Landscape of Concerns About Law Enforcement Use of Facial Recognition Systems

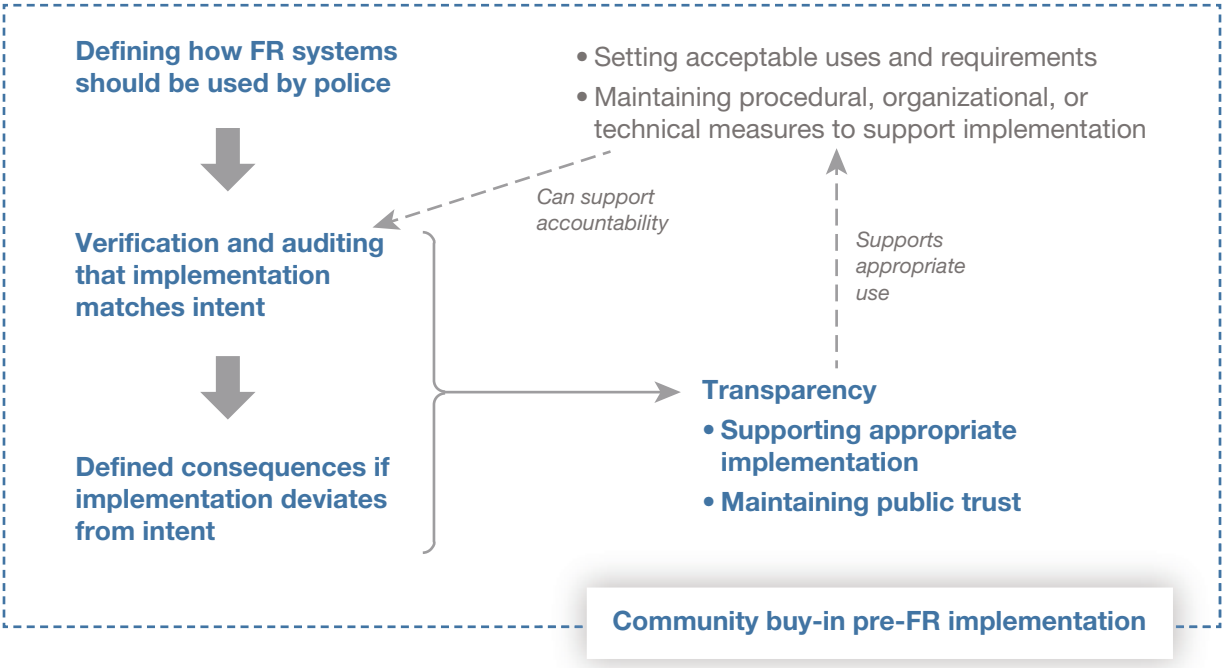

from FR technologies. Throughout the next section, we discuss this balance in terms of ways that policy can either narrow or broaden the "amount" of FR being applied by law enforcement. Where available, we provide examples of existing policy and practice to provide concrete examples and to highlight what is already being done to navigate concerns.

Building a Path Through the Landscape: Defining How Facial Recognition Systems Should Be Used by Police

Where Can You Walk? Setting Acceptable Uses

A foundational policy option for regulating the use of FR by law enforcement is defining the circumstances in which it is allowed—if at all—and when it is prohibited.[1] Between the extremes of free use of FR at any time and full prohibition of law enforcement use lies a wide variety of options for narrower application;[2] the specific choices of how to narrow use address some of the risks and concerns regarding the potential consequences of the usage.[3] Returning to the language of the previous section, this option defines—locally, on

[1] Note that in practice there is the precedent that *exigent circumstances*—a situation in which there is a time-sensitive risk of harm or other damage—can be used to justify law enforcement setting aside restrictions that have been placed on its behavior or use of tactics or technologies. We address this below in the section on consequences for not abiding by restrictions on FR usage.

[2] A total ban on use essentially is declaring that no use is acceptable and forgoing any potential benefits to avoid risks and costs (Conger, Fausset, and Kovaleski, 2019; Richardson, 2021; Spivack and Garvie, 2020).

[3] This is related to the concept of *proportional use* described in International Criminal Police Organization et al., 2021. See also the discussion regarding the U.S. context in U.S. Government Accountability Office, 2021a.

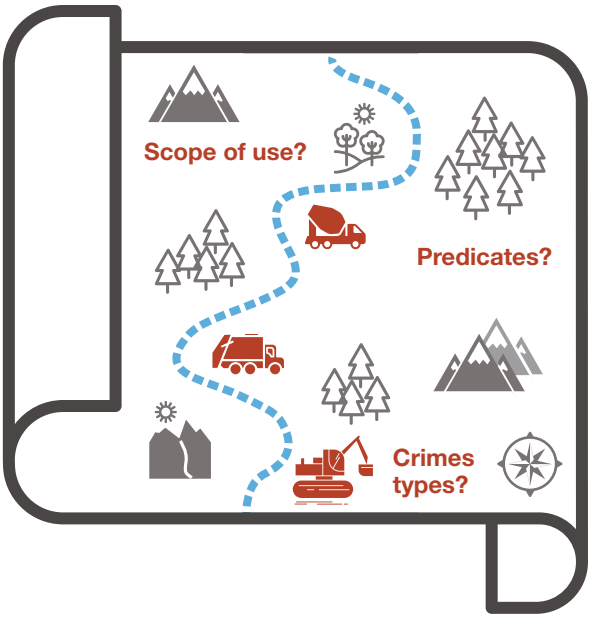

the basis of local concerns and preferences—what constitutes an acceptable use case of FR and what tasks or applications are misuse cases that are then prohibited. Defining the authorized uses, if any, is a key first step in policymaking for FR use by law enforcement.[4]

Facial Recognition for What Crimes?

The crimes that FR is used to prevent, interrupt, or solve define the potential benefit of an individual application of the technology. Stopping a child abduction in progress is much more valuable than stopping an instance of petty vandalism. Solving a murder is more valuable than solving a shoplifting incident. One option for navigating concerns about risks of FR use is to restrict the crimes it can be applied to, approving use for applications in which the potential benefits are high and prohibiting it for cases in which the benefit would be lower (U.S. Department of Homeland Security and Bureau of Justice Assistance, 2017, p. 14; International Criminal Police Organization et al., 2021, p. 14). This approach accepts that there are perceived costs for use, and seeks a path that balances them with the societal benefits of crime prevention and response. Limitations that prohibit some uses of particular concern (e.g., routine identification of participants in First Amendment–protected activity)[5] can be specified to ensure that those concerns are minimized. Currently, some states have limited the use of FR to felony crimes (e.g., Utah, Maine, and Massachusetts are considering this). Many states also allow noncrime exceptions for FR use to identify missing, endangered, incapacitated, or deceased persons.

The minimal baseline for this type of path to address concerns is policy prohibitions of any use of FR systems that is not related to a law enforcement task—that is, instances akin to officers searching police databases for their own reasons and purposes. Beyond preventing these types of uses that have only costs and no agency or societal benefits, where to draw the line for other uses could differ across jurisdictions. Past analyses have sought to determine the value of criminal justice activities by estimating the costs associated with different types of crimes and, therefore, the potential value of preventing them. Cost-of-crime estimates seek to combine costs to society (associated with investigating, prosecuting and punishing the perpetrators of different types of crimes) with the costs borne by victims of crime. Assigning dollar values to the effects of a crime on victims can be difficult. Valuing a stolen car or belongings taken from a home is one thing, but valuing the more intangible effects (both on direct victims and others around them, such as their families) of violent crimes is more difficult. Nonetheless, estimates done by a variety of methods can provide some insight. Figure 4.2 shows approximate values used in a past RAND effort to summarize value associated with different crimes, with the costs of a murder above $8 million.

[4] Several states have passed statutes that require law enforcement agencies to receive local governmental approval before beginning an FR program. Other states centralize FR use by allowing only state police agencies to acquire and use FR.

[5] Washington's law identifies protected characteristics; see Revised Code of Washington, Chapter 43.386.

FIGURE 4.2
Relative Estimates for the Costs of Different Crime Types

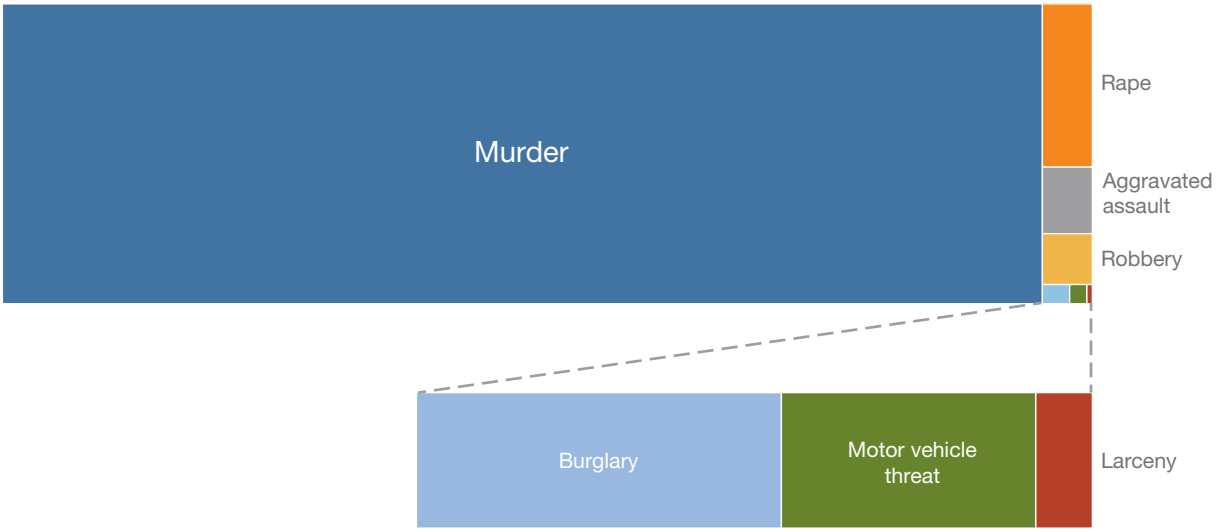

SOURCE: Adapted from Heaton, 2010.
NOTE: Other, more-recent analyses have made estimates of just the governmental costs associated with different types of crimes (Hunt, Anderson, and Saunders, 2017; Hunt, Saunders, and Kilmer, 2019).

With the estimated costs of individual crimes varying over several orders of magnitude, those costs could provide a guide for what applications of FR an area would want to allow and what to prohibit on the grounds that the estimated benefits do not outweigh concerns.

What Predicates Are Needed to Use Facial Recognition?

A second potential narrowing of how FR is used by law enforcement is to define the specifics of the situations in which an officer can use the technology to identify someone. Although a law enforcement officer can certainly ask anyone to voluntarily identify themselves at any time, there are differences in local laws regarding when the officer can compel an individual to identify themselves. Assuming that the officer is in a position to take a photograph of a person without their consent (and assuming broad effectiveness of the FR system being used for the general population), if there are no restrictions on FR use, this could be functionally equivalent to officers being able to compel people to identify themselves at all times.

Existing limitations on officers' ability to compel individuals to identify themselves are anchored in the same concepts that guide other searches and intrusions into individuals' privacy. Compelling someone to identify themselves may require *reasonable suspicion* that the person has committed or is involved in committing a crime. More-invasive searches—i.e., receiving a warrant to search someone's home or vehicle—requires *probable cause* that a crime has been committed or that there is evidence of a crime in the location police want to search.[6]

[6] Reasonable suspicion has a lower standard of proof than probable cause and requires that officers have a specific and articulable basis for suspecting that someone has or is about to commit a crime based on the current facts or circumstances. In person, this means that officers can detain and even frisk someone based on reasonable suspicion but is not by itself grounds for arrest. Probable cause means that there is a reasonable basis based on the totality of circumstances to believe that a crime has been committed or evidence of a crime exists. Probable cause is required for searches or arrests.

An area that treated FR as similar to compelled identification could require that officers have reasonable suspicion before FR use was authorized. Determining what sort of documentation would support whether FR use was appropriate would be an implementation issue, since the level of trust in a department's or individual officer's adherence to such a policy could vary from place to place. If some uses of FR (including broader scope uses, which we discuss next) are considered more akin to an invasive search, probable cause could be necessary—and potentially similar processes around the request and approval of use, such as the issuance of a search warrant—and application would be more restricted and controlled.

A recent Massachusetts special commission has recommended the standard of probable cause with a warrant, and FR is limited to felonies (Facial Recognition Commission, undated). Conversely, Virginia just passed a statute that allows for a reasonable-suspicion standard (Virginia Senate Bill 741, 2022). This is an important choice for policymakers, as a narrower predicate (probable cause) relative to a broader predicate (reasonable suspicion) could mean fewer FR searches and more due-process protections. But it also likely means reduced efficacy of FR for law enforcement, in terms of the number of crimes solved. The actual difference between how the two standards operate in practice is still unclear but could be assessed by comparing states with different policies.

What Scope of Facial Recognition Application Can Be Used?

As discussed in the previous chapter, the application of FR in use cases where one individual is being identified is qualitatively different from broad implementation of FR for large numbers of people at one time. Although such use cases as searching crowds for individuals in watchlists have been explored, and would appear to fall in an area where there is no expectation of privacy, such broad-scope applications raise the most concerns regarding the effect of the technology on privacy and on the relationship between police and the public.

In addition, such real-time usage has many similarities with other situations in which law enforcement has sought to identify large numbers of individuals present in a given area—from mass detentions during protests where there were some violent participants to the use of geofence warrants to identify all of the cell phones present in an area where a crime occurred. In those analogous cases, external oversight entities have pushed back—including judges reviewing geofence warrants and narrowing their scope due to the perceived intrusiveness and civil actions after broad detentions resulting in significant judgments against police departments.

Another facet of scope in the use of FR is how the capability to identify individuals in photos at specific locations and in specific contexts could reveal other things about their actions and activities, potentially impinging on other types of rights or constituting more-significant invasions of privacy than even FR alone.[7] An example of such a combination of data would be photographic information collected at a location viewed as sensitive, including medical facilities, counseling or treatment centers, or religious locations. Deploying FR in such a way means that barriers could be erected to individuals visiting such locations. The combination of FR with other types of video analytics could also be used to gather data in a systematic way, which could affect individual behavior. For example, a persistent challenge in video analytics for security applications is the detection of firearms in video data. A system that combined that capability with FR could allow the

[7] European policy-guidance documents suggest banning these applications: "[FR], and other face analysis technologies, should be used for no purpose other than biometric identification/recognition/verification. The use of [FR] to infer ethnicity, gender, sex, age, emotion, opinion, health status, religion and sexual orientation, and the use of [FR] for predictive analysis, should not be permitted" (International Criminal Police Organization [Interpol] et al., 2021, p. 15). Some examples of proposed U.S. policies are less restrictive: "Entities should consider an additional level of review and approval in order to enhance protection and ensure appropriate use of this technology in sensitive locations or populations" (U.S. Department of Homeland Security and Bureau of Justice Assistance, U.S. Department of Justice, 2017, p. 21). See also discussion in Spivack and Garvie, 2020.

development of a dataset of people who carried firearms in public, even if their doing so was legal under the laws of their local area. This issue of the combination of video recognition information with other location or contextual data is not solely a concern with FR. Similar questions have been raised about the analysis of license-plate-recognition data and, therefore, the ability to track specific vehicles visiting locations that reveal potentially sensitive information about their owners.[8]

Prohibitions or restrictions on the scope of FR use could therefore be elements of a policy approach seeking to limit risks while still allowing some use of the tools (Spivack and Garvie, 2020). The options for this approach fall along a spectrum, ranging from prohibition of all broad-scope applications (e.g., no many-to-many or real-time searches performed as part of surveillance activities)[9] to the combination of scope restrictions with other approaches (e.g., broad-scope searches could be done *only* for particular crime types or in situations in which rapid action was needed to prevent harm). Credibly implementing such a narrowing of acceptable uses could require additional organizational and procedural elements. Unlike the case of geofence warrants for cellular information—when an agency must receive a warrant for access to data held by others and therefore has a built-in practical constraint on the capability—an agency-owner of an FR system could achieve the same end without the external oversight.

Some states have sought to limit scope of FR applications by specifically prohibiting real-time surveillance (e.g., Washington, although it is allowed with a warrant), a many-to-many application. In other states, the predicates (e.g., reasonable suspicion or probable cause) and authorized uses serve to limit the scope of FR application because they must be connected to criminal behavior or another authorized extenuating circumstance. All three of these factors (authorized uses, predicates, scope) are important to consider and set the framework for how broadly FR will be applied within a jurisdiction.

Policy and law should define

- the types of crimes or other situations that FR is authorized for use in investigating
- the predicates required before FR can be used on an individual
- the acceptable scope of FR applications in the jurisdictions and how FR can and cannot be combined with other types of data in the course of investigations.

Designing Guides and Guardrails to Help Stay on the Path: Organizational and Procedural Structures Around the Use of Facial Recognition

Beyond defining acceptable uses for FR, there is a set of policy elements that can support agencies implementing those choices faithfully and in ways that are credible to outside audiences concerned about the impacts of the technology. Returning to the metaphor of policies defining a path through the landscape of potential FR uses and concerns, organizational choices and procedures around FR use can create guardrails to help ensure that individual users and the agency overall stay on the path that has been defined for them.

[8] For example, the DC Circuit Court summarized: "A person who knows all of another's travels can deduce whether he is a weekly church goer, a heavy drinker, a regular at the gym, an unfaithful husband, an outpatient receiving medical treatment, an associate of particular individuals or political groups—and not just one such fact about a person, but all such facts" (*United States v. Maynard*, 2010, quoted in Brooks, 2019).

[9] Reportedly, some states have prohibited the combination of FR with police body-camera footage, which would be a broad-scope application with officers moving among the population in many areas during the course of their daily activities (Slaight and LeCloux, 2020). The IJIS Institute gives an example of body-worn camera video real-time use in their use case document; see Law Enforcement Imaging Technology Task Force, 2019.

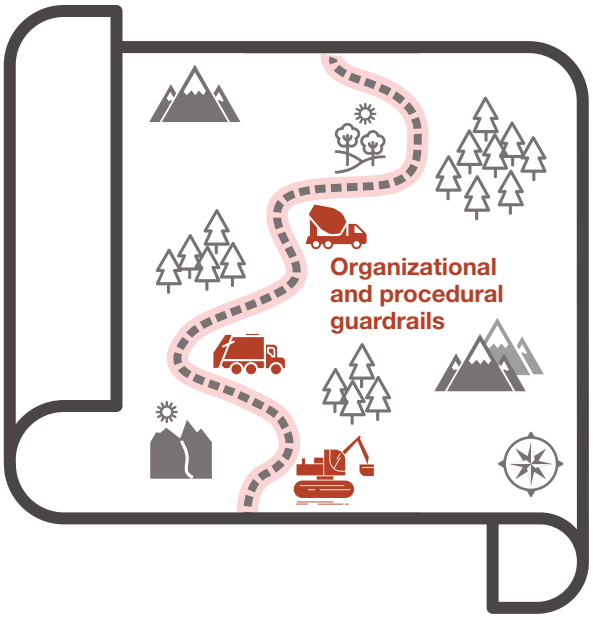

One very basic procedural approach for managing users of FR is having policies for when an officer is authorized to take an individual's photograph at all. Without a photograph, no use of FR is possible and—if a photo is taken under circumstances that are prohibited by policy—the photo itself would provide its own type of audit trail to identify and respond to misuse.

Beyond such baseline policies, however, there are other policy options that seek to build in incentives to abide by defined policies for FR use and to limit incentives to conceal deviations from those policies. The reason why FR is an attractive technology to law enforcement is its potential contribution to identifying individuals to help close cases. Effectiveness in solving crime and closing cases is a key element of how individual officers and their agencies are assessed. Therefore, given any open case in which FR might be useful, there is an incentive to use it. That incentive could reasonably create a temptation to stray from the path, given the potential benefit to the officer or agency for doing so. As a result, choices about who performs two key actions—(1) approves the use of FR for a task and (2) actually performs the FR search—can play an important role in managing how that incentive can weaken an agency's accountability for FR use.

Figure 4.3 shows this dynamic for different options for approval of FR usage and FR search performance. Having individual investigating officers with the authority to use these systems directly without any outside approval for use provides the weakest support to accountability, since there is no external check on whether use matches defined policies and—if the system is used inappropriately—strong incentives to conceal that fact. When authority moves away from the individual directly involved in the investigation (i.e., in Figure 4.3, moving to the right, down, or both), incentives become better aligned—with the individual who is further from the direct benefit of a specific case being closed being better positioned to assess the overall benefits and risks of use.[10]

Centralized decisionmaking about use can also help support external accountability, since patterns of use and misuse could be more easily attached to an individual person. Having staff who are external to a police agency involved in actually performing FR searches provides stronger accountability with respect to whether uses are in accordance with agreed-on policies and also weakens any incentive to conceal misuse—since outsiders would be less susceptible to internal agency cultures or pressures to do so. Options for such external actors include having FR systems resident in multiagency entities or task forces, non–law enforcement organizations,[11] and state-level law enforcement agencies.[12] The area at the far right and bottom of

[10] There is a similar analogy for officers engaged in field uses of FR, whether random interactions or reasonable-suspicion interactions. The officer has a vested interest in maintaining their safety but also should be interested in maintaining the rights of the individual and engaging in constitutional policing. Such field uses also include a limited window of time to verify information and make a decision to act, however, which may increase risks of mistakes.

[11] In a law passed in Massachusetts (Massachusetts Bill S.2963, 2020), "police first must get a judge's permission before running a face recognition search, and then have someone from the state police, the F.B.I. or the Registry of Motor Vehicles perform the search. A local officer can't just download a facial recognition app and do a search" (Hill, 2021).

[12] Note that moving either approval authority or performance (or both) away from the individual-officer level also responds to the potential for clear misuse of these systems (e.g., cases of officers searching for individuals unrelated to police work). See

FIGURE 4.3

Organizational Location of Approval and Performance of Facial Recognition Searches: Effects on Accountability Support

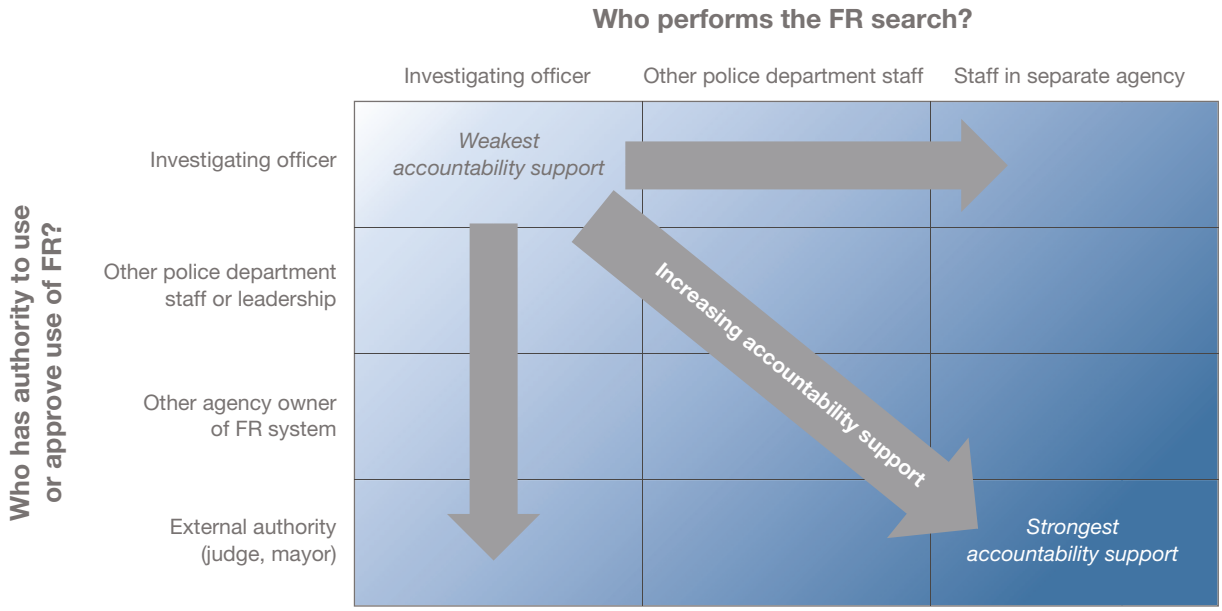

Figure 4.3—approval of FR by an external authority and actual execution of searches by personnel outside the police agency—is analogous to a warrant requirement for access to identification data held by an outside organization in the course of an investigation.[13]

Several states have implemented statutes that call for the centralization of FR searches by law enforcement, with searches being allowed only by a state law enforcement agency or the FBI (e.g., Massachusetts, Maine, Utah),[14] while others are not centralized but have the state agency set minimum policy standards (e.g., Virginia). Additionally, at least one state has set a probable-cause-with-a-warrant standard, while other states have set a reasonable-suspicion standard, and some have not set standards.

Policy and law should define

- when officers are allowed to photograph individuals, therefore limiting use procedurally
- who has authority to approve use of FR by an officer and the requirements associated with that approval
- where the FR system is used and operated organizationally and therefore who has the ability to submit images to the system for analysis.

U.S. Department of Homeland Security and Bureau of Justice Assistance, U.S. Department of Justice (2017, p. 17) for a reflection of some of these issues in a model policy document.

[13] This is the level of external control proposed in Lynch, 2019, and similar to the enacted Massachusetts law regarding FR regulation (Massachusetts Bill S.2963, 2020; Hill, 2021).

[14] Sometimes a state department of motor vehicles has authority to use FR as well.

Treading Carefully: Defining Technical Quality Standards and Supporting Practices

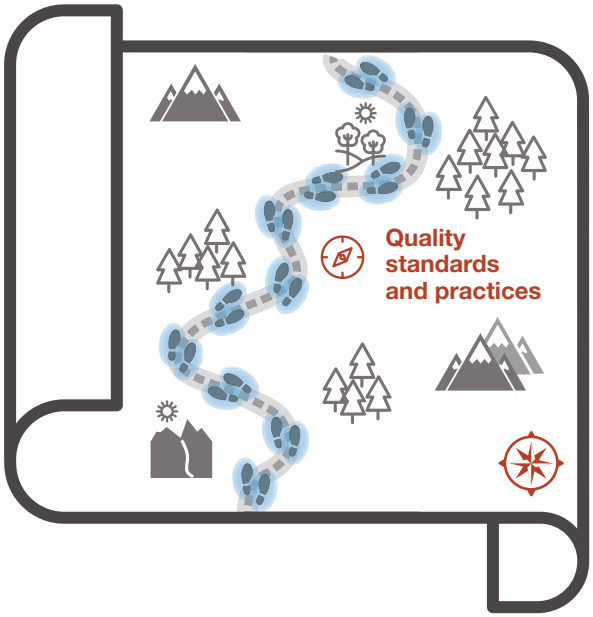

Even when a police agency is scrupulously staying on the path of FR uses that have been assessed as appropriate, if the technology does not perform well—e.g., resulting in an individual misidentified as a suspect in a crime—the net outcome could still be negative. Paradoxically, the higher profile the crime, and therefore the higher value the potential use of FR, the higher the potential costs are of an individual being misidentified by the technology—e.g., being flagged as a suspect in a murder investigation would represent a very different situation from police contacting an individual associated with a less serious crime. As a result, a variety of policy choices could be added to a jurisdiction's approach to police use of FR that force users to walk carefully along their defined path, limiting the potential for misidentification and therefore risks to individuals as a result of FR use. These include limits on the data used by the FR system involved, requirements designed to increase accuracy of matches that FR produces for different law enforcement uses, and limits on what data from FR searches can be stored long term for potential use.[15]

What Data Can Be Used for Facial Recognition? Data Sources and Databases

Concerns about FR misidentifying people in the course of investigations often turn on the quality of images that are submitted for identification. Such concerns derive both from practical constraints, including the conditions in which a photo of a criminal incident was recorded or the quality of the camera doing so, [16] but also from manipulation of photos by users before submission (e.g., mirroring a partial photo, submission of an artist's rendering of a person). Defining explicit limitations on the quality of images that can be submitted—and prohibition of manipulations—can be part of the portfolio of policies intended to reduce the chance of FR misidentifying someone as a suspect in a crime (Richardson, 2021). FISWG has been developing guidance on image-quality assessment, image factors to consider in facial image comparison, morphological analysis, and system testing for accuracy and identity ground truth (i.e., testing and understanding the algorithm before using it operationally). To our knowledge, the extent to which agencies follow these guidelines and standards is not known.

Beyond the images submitted to FR systems, the other core component is the nature of the images in the database used to identify people in those images (i.e., gallery data). As discussed, there are many factors that influence concerns about an FR system's database. The nature of the technology and algorithms used in FR systems shape the ideal design of a database for accuracy and performance in different ways. Just as

[15] These issues can be captured—and effectively enforced—by technical standards. If systems' adherence to specific sets of standards is required by law (e.g., for FR systems purchased by police) or becomes effectively required because of concerns about liability or the credibility of results, these standards may very effectively shape behavior (see the discussion in Kindt, 2021).

[16] High resolution is critical for producing high-quality probe images, but there is no set standard across agencies.

the quality of probe images matters, so does the quality and nature of images in the database (Interpol et al., 2021, p. 19). Database size can matter as well. In some cases, smaller databases result in a lower probability of false matches (Grother, Ngan, and Hanaoka, 2019a), suggesting that departmental policies that limit whose photos are included in their FR system may reduce the chance of an identification error. However, the fewer people that are in the database, the higher the likelihood that any given suspect in a particular crime may not be included, which, depending on how the system functions, may increase the chance of misidentification (e.g., if the algorithm chooses the best-available match in the database even if it is not a particularly definitive match for the submitted photo). Smaller databases—limited to booking photos, for example—could also create bias in the outcome of FR systems, since there may be racial, socioeconomic, and other factors that affect the likelihood of an individual being arrested (and therefore photographed at booking) during an encounter with police (Slaight and LeCloux, 2020).

As a result, the policy lever of specifying the nature of the FR databases that can be used by a local police department can pull in different directions with respect to concerns about the technology. Fewer people in a database may improve performance, depending on the algorithm involved, but decisions about who is and is not in the database shape how the potential costs of false matches are distributed, and that risk may be more equitable and not borne by only part of the population. There is a trend toward larger databases, but some states have set limits on what is allowable (e.g., mugshots only).

How Good Are Results? Increasing Accuracy of Matches

The true accuracy of FR systems in a law enforcement setting, both the FR algorithm and human user, is not currently known. Apart from the images used for FR, there are technical and procedural options that can shape the accuracy of the search results and therefore affect the risk of false matches and how risks associated with errors accrue from use. Measured by independent testing, the quality of FR algorithms has been increasing, but there are still differences in performance depending on how they are used, image quality and other variables (Grother, Ngan, and Hanaoka, 2022). Most importantly, it is not known how well current testing translates to the typical images used in a law enforcement context (e.g., surveillance video). Specific FR systems can also have issues with bias depending on the data they were trained on—e.g., there have been concerns about matching methods performing poorly for many minority groups, leading to higher mismatch rates for some populations.

Design choices for different systems matter as well—for example, when a system returns the top-ten matches, *even when none of those matches is particularly close* and does not provide the user with the information needed to recognize that (e.g., a score rating the level of confidence in the match), risk can be increased. Such a design risks that the user will assume that one of those top-ten matches is correct and act accordingly, but even more important denies the user the information needed to appropriately discount the value of the information in the investigation. In the same way an investigator would discount a vague eyewitness description to a crime, they should also discount a low-quality result from FR—and should reasonably adjust the type of action they can justify taking based on such a result. Moreover, some systems return up to 500 candidates, and some are not ranked by similarity score. Determinations about what threshold to set (quality of the results) and how results should be returned (number and ranked or unranked) are important factors, but different agencies appear to be using different approaches.

With the development of independent testing, quality levels for different algorithms can be established outside law enforcement agencies—and could be imposed in industry standards as well (e.g., industry consensus not to return a result below some level of match certainty). Such results provide ways for potential law enforcement users to distinguish between systems before acquiring them, even if users lack the technical

skills to assess the technologies themselves.[17] Some jurisdictions require pilot testing as part of the FR adoption process, which could be used to examine whether their system is set up to produce quality results.

Although the chance of FR algorithms producing an incorrect result is in part about the systems themselves, it is also about the people who are using them. Blindly trusting the technology and answers produced by FR is no better than assuming any computer necessarily produces the correct answer to a problem in other circumstances. As a result, having defined processes for review and checking of results by human experts can reduce the chances for bad outcomes. Just having a person in the loop is not enough, however. At least one state (Washington) includes a requirement for "meaningful human review" in its statute and also includes provisions for training (Interpol et al., 2021; U.S. Government Accountability Office, 2021b). A person in this role must be trained to both assess the results of FR searches and approach that assessment with a skeptical eye. This can be part of the rationale for limiting individual officers' direct access to FR systems or even managing those systems and staffing their use separately from the police department. Doing so can allow concentrated training and building of expertise among the smaller group of individuals authorized to do searches, rather than trying to build that knowledge broadly in an entire department of officers or investigators. Still, some agencies currently allow for officer field use of FR, which is the most decentralized approach. Although agency-to-agency training varies, FISWG has proposed a set of minimum training criteria and other guides to support training. More research in this area is needed to determine training effectiveness and key features.

The details of the person-in-the-loop process can limit the potential for other biases to affect judgment about whether two photos are indeed the same individual. For example, concerns have been raised about examiners having access to the individual's criminal history data when they are reviewing FR matches. The fear is that an examiner knowing that some of the individuals returned by the system committed crimes previously may create a confirmation bias and skew the examiner's assessment of whether those individuals match the photo associated with the new criminal incident. Another related component of meaningful human review and bias reduction that some agencies use (and our more technical interviewees recommended)[18] is the use of secondary peer review and management review.[19]

How Perishable Should Facial Recognition Results Be? Considering Data Storage Limitations

As storage capacity for electronic information has gotten cheaper, civil liberties organizations have raised concerns about the potential for large video datasets to be collected and stored for extended periods (Lynch, 2019). Their concern is greatest for broad-scope implementations of FR (discussed previously) in which data from many cameras cover broader geographic areas (see Amnesty International, 2021). Such datasets, when analyzed longitudinally using FR, could provide a powerful way to track individuals over an extended period to determine locations they have visited. The utility of such a capability for criminal investigation is obvious—for example, a suspect could be followed backward in time to determine other locations that might yield clues to a crime committed elsewhere.

Rather than storing video data, truly broad implementations of FR could instead retain only the results of FR analysis, building logs over time of the matched identities of everyone passing through a monitored loca-

[17] See the description of *ex ante evaluation* for both quality and bias in Interpol et al., 2021, p. 17.

[18] The Major City Chiefs Association Facial Recognition Working Group (2021) recommends secondary and management review as well.

[19] Peer review usually involves having another officer authorized and trained to use FR review the FR results and come to an independent decision as to whether the results could be used to move the investigation forward. Management review might not be independent but does involve additional scrutiny. Whether and how often peer review catches errors and changes the final decision regarding FR results are unknown and in need of further research.

tion. This would require less storage than retaining the actual images, potentially making some surveillance applications more attractive to some FR users. Although this approach would again be of greatest impact if implemented broadly, it is possible to envision narrower implementations, including in central public areas of a city center.

For policymakers considering these types of applications, the trade is between the investigative utility and the potential invasiveness of such data being collected—as well as how the potential for retrospective tracking of one's movements and activities could chill important activities and degrade individual rights. Retention of just FR matches is potentially problematic from a risk perspective, since not retaining the associated images could propagate errors of identification—and confound the ability of any quality-improvement approaches for detecting that error if the data were used at a later time.[20] Policy decisions around retention policies, therefore, would appear significant, whether that retention is of video or FR results alone.

Policy and law should define

- quality standards for photos that can be submitted to an FR system in the jurisdiction
- the nature of the database used for identification in the jurisdiction and its level of representativeness for all relevant groups or subpopulations
- minimum levels of accuracy for any FR algorithm used in the jurisdiction and the algorithm's performance across all different age, racial/ethnic, or other subpopulations relevant to the community
- requirements for specific technical features supporting accuracy, such as the system *not* returning any match below a specified level of quality or certainty
- standards of the minimum effectiveness for algorithms used in FR systems locally and the provision of data necessary to assess that effectiveness before purchase (i.e., jurisdiction-specific operational testing)
- requirements for system outputs that highlight rather than hide uncertainties about FR matches and processes for trained experts to skeptically review FR outputs before use in investigations
- requirements for FR systems to provide users with measures of match quality and departmental standards for the types of action that can be taken at different levels of match quality or certainty
- limits on how long individuals' photos can be stored in FR databases used in the jurisdiction and how long data on matches from past searches can be retained.

How Do Others Know You Stayed on the Path? Verifying and Auditing Whether Implementation Matches Intent

Although policies that define clearly what can and cannot be done with FR and set guardrails around implementation are critical, they may not be enough to address concerns about police use of the technology. If there are concerns about how police will use FR, communities may not fully trust that law enforcement agencies will actually do what they say they will—and even those who trust police may positively view mechanisms to verify that trust is well placed. That is, the level of perceived invasiveness or concerns about invasiveness might be balanced by the level of transparency and confirmation that the system is working as it should. At a basic level, without some form of auditing, misuse of FR systems would be unlikely to ever be discovered, requiring concerned communities to take users at their word that they are following policies as they should.

[20] Existing policy development templates for FR in the U.S. context recommend against storing match data (see U.S. Department of Homeland Security and Bureau of Justice Assistance, U.S. Department of Justice, 2017, p. 33).

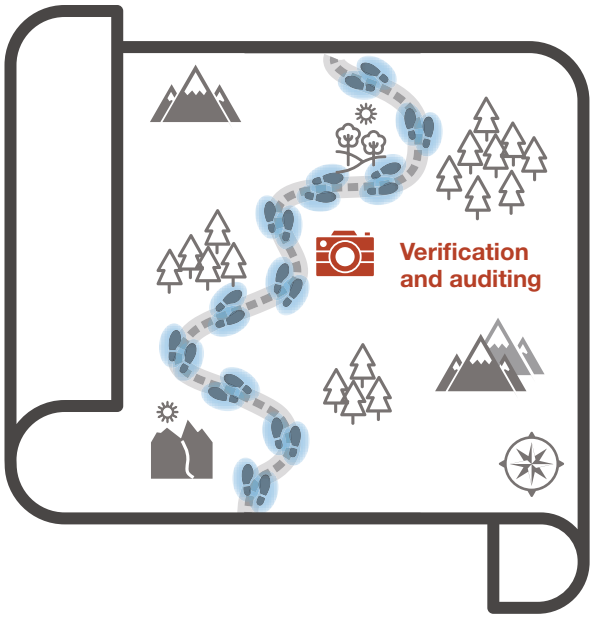

If this is the case, the next class of policy design options should center on the verification and auditing structure that will log police agencies' steps through the FR landscape to make it possible to see whether they stayed on the desired path.

For any framework that will include these types of verification approaches, the first step is prohibition of any nonaudited system. An early source of concern regarding FR was individual officers seeking access to commercial FR capabilities on their own, sometimes without the knowledge of their own departments and—therefore—without any accountability structure. When verification and auditing are part of the desired policy framework around FR, unauthorized access must be prohibited—or any auditing of an official system for access to FR would be rendered meaningless.

The real value of auditing and verification approaches depends strongly on how they are designed and the strength of that design for discovering inappropriate behavior. For example, a system may require that a user input a case number when a search is done, associating that search with an investigation for a specific criminal act. Although that is a first step, if there is no mechanism to check whether that case would be considered an acceptable use of FR or whether it is reasonable that FR is being used for the case, then the user entering a number when they upload a photo does relatively little to ensure accountability. Put another way, to ensure accountability, an audit must be suspicious, looking at system use the way a critical outsider might rather than starting from an assumption that users will use the system appropriately. Organizational separation (i.e., audit by an independent agency or entity) can also strengthen oversight (Lynch, 2019). Other ways that auditing could assess use would be tracking how frequently individual users submit photos to the system, patterns in the types of cases they use it for, and so on—all of which could provide ways to detect anomalies that suggest misbehavior. Standards for FR systems might require that the software tracks every aspect of user access, actions, and results within the system, much as some body-worn-camera data systems are structured currently (e.g., a system automatically logs when a user accesses a file, views the video, or makes any changes).

Other types of audits could focus on key concerns of FR use, such as privacy, accuracy, and bias. Privacy audits might explore how data are collected, stored (e.g., security), managed, shared, retained, and purged. This would provide critical feedback as to how well this biometric information is being handled and uncover potential risks in the data life cycle. Accuracy audits might select a random subsample of prior FR searches and submit them to independent testing to determine whether the results could be replicated. Given the technology-human interaction that is a part of every FR search, this would include a review of how well the process and decisions are documented. Bias audits could explore how the actual use of FR corresponds to differential performance across subgroups (Lee, Resnick, and Barton, 2019). We noted in the previous chapter that the images used in FR demographic tests, and even accuracy tests more generally, do not account for the same real-world conditions that exist in law enforcement use cases. The first step would involve assessing how often different outcomes occur across groups (e.g., the probe image was rejected for poor quality, the algorithm returned the correct person in the top-ranked position, a candidate is selected and referred as

a potential lead, and peer review results in different determinations). In particular, there should be assessments of both algorithm-generated and human-generated outcomes. FR software companies have improved and continue working to improve accuracy across demographic groups, but the software is only one piece of where bias can affect the system (i.e., FR use in law enforcement is not automated).

Although auditing is often viewed as checking up on or tracking the behavior of individual users, the existence of audit capabilities also serves to support the credibility of policies as a way to address the concerns of skeptical audiences. Tracking data on use also is critical for any post hoc assessment of the value of FR systems, since without information on when FR was used and how, no objective assessment of its contribution to criminal justice outcomes can be successful.

Some states have implemented requirements for annual reporting in FR legislation, which are not necessarily audits but could enable auditing (Lewis and Crumpler, 2021). These cover basic features of FR use, such as the number of requests (and types of requests), the number of searches, the number of probable matches, the gallery data source (e.g., mugshots, driver's license), and the race and gender of the probe-image subject. Some states also require reporting on the scope of the deployment and measured performance of the system (e.g., Washington Senate Bill 6280, 2021).

Policy and law should define

- what information must be recorded with each FR use by authorized users of the system, supporting what types of retrospective review of the appropriateness of that use
- frequency of audits of system use and what sorts of analyses are done during the audits to examine and assess how and why specific users utilized FR
- how long use and other audit data must be kept to enable retrospective analysis of system usage and contribution to criminal justice outcomes
- the accessibility of audits, including what and how to release publicly.

What Happens When You Stray from the Path? Defining Consequences If Implementation Deviates from Intent

For there to be true accountability to whatever path a jurisdiction has set for use of FR, there should be consequences for straying from that path. Without some consequence associated with policy violation, there is real risk that the perceived benefits to individual officers or the agency for efficiency or improved performance will weaken the ability of policies or law to manage FR risks.

Furthermore, in some communities, there has been considerable frustration about the ability to hold law enforcement—both at the individual officer and agency levels—accountable for actions that cause harm or deviate from policy or procedure. Much of the public debate around this issue has centered on police use of force and consequences for patterns or serious individual incidents of police violence.

Given concerns that officers and police agencies often do not face real consequences for actions viewed as extremely serious, skepticism that law enforcement will follow the rules regarding FR usage is understandable. As a result, in designing a jurisdictional approach to FR, what sorts of consequences exist for deviations from policy or desired practice is likely an important component. These include the following:[21]

[21] Examples of both types of sanctions are included in U.S. Department of Homeland Security and Bureau of Justice Assistance, U.S. Department of Justice, 2017, p. 35.

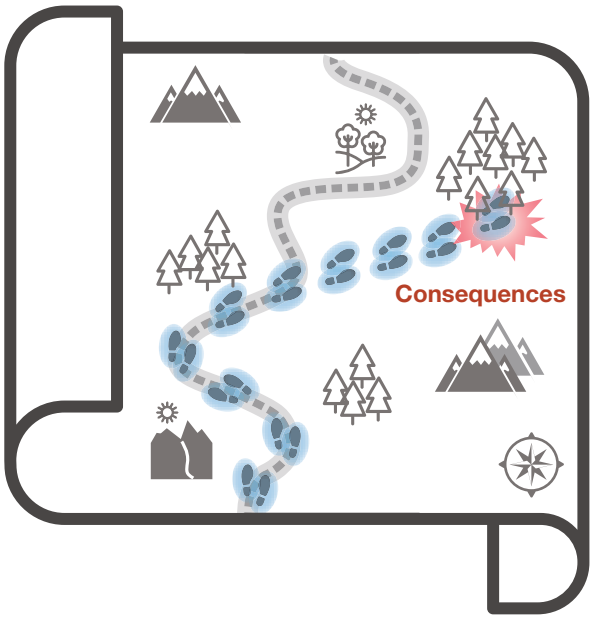

- *Individual sanctions, the penalties that an officer or analyst user faces for misuse:* These could vary from loss of access to the system, personnel sanctions, or even legal consequences for serious transgressions.[22]
- *Departmental sanctions, the penalties faced by the department for misuse of the capabilities:* For systems that are managed outside the police department, these could include severing the agency's access to the system, but other options could include budgetary or personnel actions taken by political overseers of the department (e.g., firing of agency leadership).

Whether a sanctions regime will address outside audiences' concerns may depend on whether results are made public—relating to the broader issue of transparency we will delve into in the next section. Although a strong sanctions regime that is only internally administered may effectively regulate officer use of FR, it may not have significant external trust value if skeptical audiences cannot see its results.[23]

Any discussion of the regulation of law enforcement use of technological or other capabilities raises the question of *exigent circumstances*[24]—situations so dangerous or fast moving that freeing police of restraints would be viewed by a reasonable person as necessary to protect the public, and so their "breaking the rules" is in fact the correct course of action.[25] For example, situations that might be viewed as justifying relaxation of restrictions on FR could include a major and ongoing terrorist attack or a situation in which an abducted child was being removed from an area. In both cases, relaxation of limits on broad-scope uses of FR might allow action that stopped significant future harm. The question of exigent circumstances and FR links closely with consequences for system misuse, since the police's sole authority to decide that circumstances mean that they can ignore restrictions weakens the ability of consequences to support accountability. One option for managing that risk is the requirement for others to be part of decisionmaking during events, although doing so could slow the decision enough to reduce the potential utility of relaxing restrictions.[26] Another option is a mechanism of strong after-incident review, with individual- or agency-level consequences in cases in which that review determines that officers or agency leadership made the wrong call.

[22] See, for example, proposals that legal penalties be increased in Florida for misuse of law enforcement information systems (WCJB Staff, 2021).

[23] Other analyses of the application of surveillance technologies by police emphasize the need for both administrative provisions for enforcement and the creation of a "private cause of action" if administrative enforcement is assessed as failing (Fuson, 2019).

[24] Or "situations of emergency," as they are categorized by Interpol et al., 2021.

[25] An example of where this term is used is cases justifying police entering an area immediately—even though a search warrant would normally be required.

[26] "As a consequence, acknowledging the risks involved in this exceptional emergency situation, an independent authority should be in charge of authorizing this application and, if there is not enough time, it should be authorized by the chain of command. In this case, the chain of command should inform and justify the decision to the independent authority as soon as possible and not later than 24 hours" (Interpol et al., 2021, p. 15).

Policy and law should define

- what the penalties are at both the individual and agency levels when use deviates from approved use cases
- what types of exigent circumstances justify relaxation of jurisdiction-level restrictions on FR use, the process for making the decision to do so if such an incident occurs, or sanctions for cases in which the decision is assessed as improper after the fact.

Why It Matters That Others Know How and Where You Walk: Transparency and Community Approval

Given that a core driver of questions about how and when police should use FR is public concern that the risks of such use may outweigh its benefits, what the public knows about law enforcement utilization of FR and whether the public authorizes or approves of that use matter. Members of the public cannot actively agree with and support something that they do not know anything about and, although it is possible for law enforcement to keep some capabilities secret from the public, such a strategy can be a fragile one. If an agency secretly implements a new technology, it should not be surprised if affected communities—when they find out—reach the conclusion that there was some blameworthy reason why the agency was hiding what it was doing. Secrecy broken also can prompt strong reaction and reinforces the potential for pendulum swings from permissiveness to restriction in police use of technology. As a result, transparency and the related goal of obtaining community authorization and ongoing support can be critical components of a legislative or regulatory approach for FR in a jurisdiction (Spivack and Garvie, 2020; U.S. Government Accountability Office, 2021b).

Transparency

The central policy question regarding transparency is what information should be made public to inform local debate and ongoing public oversight of the use of FR. Transparency in this area is a focus of debate: Policymaking organizations in the law enforcement sector raise it as something that could be done,[27] while some civil society organizations are seeking to have requirements for transparency enacted into state and local law. The most basic level of transparency is disclosing that a police agency is using FR, information about the policies around its use (if those policies are internal to the department rather than externally imposed), and data about the frequency of its use. Although this provides some information to the public about FR utilization, this level of transparency has only limited potential to contribute to accountability and reinforce appropriate

[27] Existing policy development templates for the United States include some transparency recommendations, but they are relatively limited (e.g., policies are publicly available), and accountability elements are focused internally within the agency (U.S. Department of Homeland Security and Bureau of Justice Assistance, U.S. Department of Justice, 2017, pp. 34, 35). A more recent, nongovernmental resource document is more tentative with respect to transparency: "Law enforcement should also *consider* conducting and publishing a Privacy Impact Assessment (PIA) or a Data Protection Impact Assessment (DPIA). Once implemented, agencies should *consider* publicly publishing their policies on facial recognition technology to ensure transparency with the community" (International Association of Chiefs of Police, Law Enforcement Policy Center, 2022, p. 7, emphasis added).

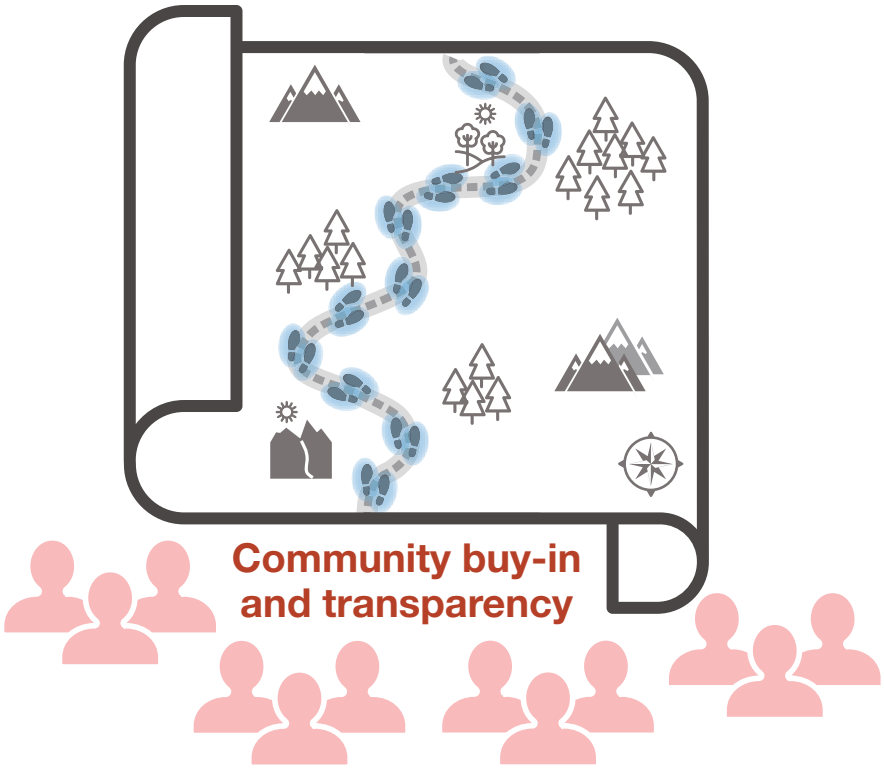

FR use. Other more-substantial types of transparency go beyond just informing the public and can have a more direct effect on incentives and likely FR use behavior.[28] They include the following:[29]

- **Detailed statistics on the use rates for case types over time, and drivers for different levels of use within case types:** Depending on what classes of cases are viewed as appropriate for application of FR in policy, more-detailed case-type data can support public assessment that police are abiding by those policies.
- **Case-level disclosure during discovery that FR was used and the quality of the FR match that led to identification of the suspect:**[30] Rather than viewing FR simply as a lead-generation tool that might not even be disclosed during discovery, its use could be treated more like other forensic science methods in which questions about quality and process are core to assessing the value of the evidence they produce. For example, in cases in which a very poor-quality FR match led to a suspect, supporting evidence that the correct individual was chosen could be examined more skeptically, to hedge against the risk that a bad lead provided by FR then skewed the remainder of the investigation.

[28] Others have noted that whether transparency efforts will shape incentives "depend[s] on whether the parts of government with oversight authority are cooperative and supportive of the law, and whether there is a robust advocacy community to provide external pressure or accountability" (Richardson, 2021).

[29] European policy recommendations for FR include most of these broader forms of transparency; see Interpol et al., 2021, pp. 15–16.

[30] Some law enforcement policy documents argue for this type of transparency, specifically Interpol et al., 2021: "For anyone identified using an [FR] system, that person must be informed that he/she was subject to such a search/that an [FR] system was used to identify them, if they are subsequently taken into custody, brought in as a witness, or have any other official role in a law enforcement process based on their face via the [FR] system" (p. 16).

- **Disclosing FR use to the defense counsel would also support adherence with policies about what cases should and should not utilize FR:** Then, inappropriate use would become a substantial liability for such cases, balancing out any perceived value that using FR could let investigators close the case more quickly. Disclosing the quality of matches could also manage incentives for pursuing very weak matches (and imposing costs on erroneously identified suspects), since the quality of the initial match could support a line of attack for the defense counsel at trial or in plea negotiations.
- **Criminal justice outcome data for cases in which FR was and was not used:** Since a central driver of FR use is the perceived benefit for solving cases, providing outcome data to the public could support judgment that the benefits of use justify the risks. Ideally, such analysis would compare similar cases in which the core difference was application of FR to assess both short-term (e.g., clearance rates, conviction rates, plea decisions) and long-term outcomes (e.g., whether the odds of wrongful conviction are higher or lower when FR is used).
- **Results of analyses examining whether effects of FR contribute to bias and differential outcomes in different populations:**[31] Given documented issues with some FR algorithms performing differently when matching faces from different racial/ethnic groups, building in assessment to determine whether that is the case in an agency's in-practice implementation of FR could be important for maintaining support for the system. Monitoring for such outcomes could build in incentives to be vigilant for ways that application of the technology could create bias and to look for ways that its use could make the justice system more fair and equitable overall.

Some police departments include disclosure that FR was used to defendants as part of policy, while others do not. Those that do not likely do not disclose such information to the prosecution either. These departments consider the FR search results only as a lead and not evidence per se. So far, Washington and Massachusetts appear to be the only states to require notice to defendants. Massachusetts specifies that the results of the search, other matches, the technology's accuracy, and the process of selecting the defendant should be included. Other states are providing notice to the public about situations in which the government is capturing images that could be used in FR searches (e.g., driver's licenses in Utah and Massachusetts). Certain types of FR use might also require more-specific notice—for instance, where ongoing or real-time surveillance is being used, information about the duration, location, number of people affected, and the number of misidentifications could be reported.

Policy and law should define

- what types of data and analysis must be made public regarding FR system use and outcomes and how frequently the data and analysis must be updated.

Community Approval and Ongoing Support

Transparency can help reduce the potential for the public to misunderstand police use of FR and the risk that an absence of information drives mistrust and opposition to the technology. However, more-substantial community *involvement* in authorizing FR use, as well as development of the ground rules for its use, could

[31] Although not framed in the context of transparency or making their results public, these types of ongoing assessments are recommended in U.S. Department of Homeland Security and Bureau of Justice Assistance, U.S. Department of Justice, 2017.

provide the foundation for more-substantial community support. Given that the potential effects of police FR use on the broader public and society are the central drivers of opposition to these technologies, direct involvement of the public in deciding what uses are and are not acceptable—and in the construction of the policy, management, legal, and incentive frameworks to support those decisions—seem critical to building a sustainable path forward (see Hill, O'Conner, and Slane, 2022). As much as this report is intended to guide policymakers, the concepts just listed and others covered in this report can be used by community members to ensure that their local FR policy is in line with their expectations. Like transparency more broadly, the need for substantial public involvement is a matter of current debate, with some law enforcement policy organizations recommending that departments "consider" doing so,[32] while some civil society organizations are seeking to have requirements for involvement enacted into law. Given the known concerns listed in Chapter 3, adopting FR is likely not akin to adopting other forensic practices (e.g., DNA) due to the potential for FR use by law enforcement to become invasive, due to concerns about disparate impact, and the current variation in practices due to a lack of widely accepted standards.

Notably, it is not clear whether equity impact assessments are a common feature of seeking approval and developing policy for FR use. These assessments are used to ensure that the full range of potential issues and impacts are considered, to increase education about these impacts, and to increase community engagement and collaboration. In particular, racial equity impact assessments identify how policies or practices might contribute to inequity even when the policy or practice appears neutral. Moy (2021) identifies options for constructing an equity impact assessment and highlights key questions related to five components of potential inequity that might be brought about by police use of technology: replicating inequity (e.g., does FR use of mugshot data incorporate inequity?), masking inequity (e.g., is equity being measured in FR use?), transferring inequity (e.g., do the data being used have embedded inequity?), exacerbating inequitable harms (e.g., is FR used in a way that enhances current inequities?), or compromising inequity oversight (e.g., are FR vendors monitoring inequity or allowing independent testing?). Moy (2021) recommends that the equity impact assessment should be completed prior to adoption of the technology to mitigate potential harms as best as possible.

Some states and localities require public engagement prior to deploying FR systems (e.g., Washington state and Seattle) or to give notice before operation and provide an opportunity for written public comment (Utah and New York City) (Lewis and Crumpler, 2021).

Policy and law should define

- what levels of community involvement and through what mechanisms that involvement must occur in the decisions to adopt FR and choices about management and oversight of use of the technology.

[32] "Prior to implementing facial recognition technology, agencies should *consider* community engagement initiatives that allow community members to voice concerns while providing law enforcement a platform to understand such concerns. These initiatives would allow agencies to appropriately address community concerns and educate the public about the use of facial recognition technology" (International Association of Chiefs of Police, Law Enforcement Policy Center, 2022, p. 7, emphasis added).

CHAPTER 5

Conclusions

Although there are ongoing calls to ban or pause FR use by law enforcement at the federal, state, and local levels, some federal agencies, states, and municipalities are moving forward with its use, and some have been using it for years. Critically, federal legislation and many state legislatures still lag behind FR use in setting regulations for law enforcement. Addressing the concerns that exist with FR use requires an understanding of how FR is used by law enforcement, the different potential applications of use, and the risks and benefits of those applications. Notably, regulations and policies need to be specified clearly and in relation to actual use so that concerns are adequately addressed. A variety of journalistic, legal, scientific, and technical documents cover various aspects of FR use generally and specifically by law enforcement. We sought to incorporate these various perspectives into a single document that summarizes the existing state of knowledge and provides discussion of policy considerations.

The Potential Ubiquity of Facial Recognition and Options for Narrowing

Combining FR with pervasive surveillance is possible with current technology and represents the most concerning application of this technology because it is overly broad and pervasive. Much of our discussion about regulating FR use by law enforcement involves narrowing the use of FR in ways that maximize benefits and minimize risks, along with careful review to ensure that these regulations were upheld. This narrowing could consist of limiting FR to investigating serious crimes, increasing the level of authorization required, reducing the number of people being searched or doing the searching, restricting the types of images or the quality of images that can be used, setting a threshold for search results, and deleting data.

Having a requirement to seek warrants from an outside source represents both an oversight and practical check on the power of law enforcement to collect and use information, and technological advancements often lead to or even require new investigative tools and techniques. Although other tools and resources also facilitate investigations, they are generally limited in their ability to identify anyone, at any time. However, FR can potentially change that, and that change is one source of concern from civil society: that the technology has the potential to fundamentally shift the balance between the police and the public.

Constraints on how the technology is used—that it cannot and will not be applied anytime, anywhere, and for any purpose—seek to shift that balance back again, partially, although this may not alleviate the concerns of those most worried about the introduction of this technology. While still allowing FR use, such constraints seek to make FR use more rare than commonplace and, by doing so, maintain some of that practical check on police capability at the price of some of the potential efficiencies the technology could provide.

Transparency about when FR is used is also a critical check. This includes public awareness of preapproved-use cases and FR use in individual cases (e.g., notification to defendants). If the system is flawed or used inappropriately, full disclosure can potentially identify issues and reduce or prevent the high cost of an innocent person being convicted.

Mapping the Policy Road Map to Concerns

Although the general idea of narrowing is helpful conceptually, the two main parts of this report focus on concerns and policy dimensions. In Table 5.1, we map the policy road map from Chapter 4 onto the concerns discussed in Chapter 3. As shown in this table, each policy dimension tends to deal with more than one category of concern, and most categories of concern are addressed by several policy dimensions. The first section (gray section), which focuses on setting acceptable uses, outlines the situations in which FR can be used and sets limits to address concerns about pervasive surveillance associated with this technology. The next section (blue section) focuses on the FR processes that can be used to reduce concerns about expansive use but also to limit susceptibility to improper interpretation or application of FR results. The next section (green section) addresses some of the features of FR systems as they are applied in law enforcement settings, which always include a combination of computer and human steps. These policy dimensions may help prevent situations in which computer and human errors are compounded rather than discovered and addressed before leading to a negative outcome. The next two sections of policy dimensions cover concerns that police officials might misapply the technology (red section) or that the risks of FR technology are not adequately being addressed as currently applied (purple section). Finally, the last set of policy dimensions (orange section) covers concerns about public and defendant awareness of FR use and implications for trust in the police, as well as the outcomes of FR use and whether the intended benefits are actually realized or are sufficiently large to outweigh any actual negative consequences of use.

Evaluating Facial Recognition Use by Law Enforcement

Although concerns about FR use are well documented, potential benefits are still mostly assumed. With FR use, as with much other technology use by law enforcement, a key question is whether the technology makes law enforcement practice better (e.g., more efficient, less biased, improved outcomes) than standard practice. This is critical for making assessments about whether benefits outweigh risks. To start, in-domain or field-accuracy testing and human-factors testing needs to be done to address concerns in this critical area.[1] Descriptive information about FR use by law enforcement is a key first step, which might examine the whole chain of use, starting from image acquisition: the proportion of images that are rejected or accepted, the proportion of images that require adjustment, the proportion of candidates who are excluded due to disconfirming external information (e.g., the subject is already in jail), the proportion of candidates that are rejected by peer or supervisor review, the proportion of candidates who are accepted as a lead, how many leads to arrest, and how many are convicted. From there, implementation and outcome evaluations might support or reject the use of FR as an improvement over and above standard practice (e.g., more efficient, more cleared cases, higher conviction rates).

Developing Clear Standards

Related to balancing benefits and concerns, a more fully developed set of standards will be important for the future. Some of our interviewees expressed that much of the policy and practice variations that currently exist in law enforcement use of FR are due to the lack of recognized standards. Since there are no clear standards, different software developers develop different systems that may lead to unintended variability in how law enforcement agencies use FR. Moreover, were there accepted standards, software companies could create software, interfaces, and settings that adhere to those standards and could not be manipulated by individual

[1] Stanford University's Institute for Human-Centered Artificial Intelligence has drafted a protocol for performance assessment that outlines what could be implemented to assess law enforcement applications of FR. See Ho et al., 2020, p. 753.

TABLE 5.1

Policy Road Map and Related Concerns

Policy Road Map	Accuracy: Bias	Accuracy: Error	Actions Taken	Data Sharing and Storage	Human Rights and Civil Rights	Officer Misuse	Overuse or Unjustified Use	Privacy, Enrollment, Consent	Public Perception	Surveillance (scope of search)
FR for what crimes					✓	✓	✓		✓	
Definition of misuse						✓			✓	
Predicates needed					✓		✓			✓
Scope of application		✓			✓			✓		✓
Authorization for taking photos							✓			
Who approves use							✓			
Who performs search	✓	✓	✓			✓				
Limits on data used	✓	✓		✓				✓		
Requirements to improve accuracy	✓	✓								
Human review	✓	✓	✓							
Limits on storage				✓				✓		
Limits on using outside technology						✓				
Technical standards		✓	✓				✓			
Audits and auditors					✓	✓	✓		✓	✓
Specialized audits (bias, privacy)	✓				✓			✓	✓	
Definition of consequences						✓	✓			
Exigent circumstances							✓			✓
Transparency									✓	
Disclosure of use, policies, data reports			✓				✓		✓	
Detailed statistics on case types over time		✓	✓							
Disclosure to defense counsel					✓				✓	
Community approval									✓	

users. This could involve requiring a valid (i.e., linked) case number to be entered for every search, automated processing of probe images (e.g., image editing or rejection of low-quality images), defined accuracy thresholds, or other features that meet agreed-on standards. Variation across jurisdictions subjects the public to different FR procedures, which come with different risks and public perceptions of legitimacy. It is well-known that public perceptions of policing are subject to social contagion, whereby critical incidents of police misconduct that receive substantial public attention will affect perceptions of police elsewhere (Weitzer, 2002, 2015; White, Weisburd, and Wire, 2018). Thus, any misuse of FR anywhere could affect perceptions of FR use by law enforcement more broadly.

Preparing for the Future of Facial Recognition and Other Remote Biometric or Object Identification Technologies

Although the primary focus of this report is FR use by law enforcement, other remote biometric or object identification technologies could potentially involve similar policy dimensions and concerns, accounting for differences across the distinctive uses, use processes, and implications of using those technologies. For instance, current concerns about the accuracy of gait recognition are sufficiently large enough that this technology is not being used by law enforcement, but it would theoretically be a valuable alternative method of remote biometric identification in cases in which there is no usable image of a suspect's face or could be combined with FR to strengthen identification. Object recognition, such as weapon recognition, is being developed and marketed by many of the same software companies that currently have FR algorithms, some of which supply FR technology to law enforcement. Most notably, real-time FR is currently being used in a limited sense in law enforcement settings, but this application may become more desirable in the future, particularly the integrated application of body-worn cameras. Policy considerations and standards for these future uses of remote biometric or object identification should begin to be developed now while they are still being tested or being used in a limited sense.

APPENDIX A

Survey Methodology

To better understand the current perspective of the American public with respect to the use of FR technologies, we fielded three survey questions via RAND's American Life Panel, "a nationally representative, probability-based panel of more than 5,000 participants who are regularly interviewed over the internet" (RAND Corporation, undated).

During the summer of 2022, we fielded the following three questions via the American Life Panel.

Question A:
Rate your level of agreement with the following statement: I trust my local police department to treat everyone in my community fairly as police officers work to respond to and prevent crime.

- ☐ Strongly Disagree
- ☐ Disagree
- ☐ Neutral
- ☐ Agree
- ☐ Strongly Agree

Question B:
Some police departments are using technologies that automatically identify people in videos and pictures (called facial recognition systems) in their efforts to investigate and prevent crimes.

- ☐ Strongly Against
- ☐ Against
- ☐ Neutral
- ☐ Support
- ☐ Strongly Support

Question C:
If your local police department began using a facial recognition system, which of the following statements best reflects your view about a photo of your face being included in the system's database?

- ☐ I have no concerns about being included
- ☐ I have some concern about being included
- ☐ I have major concerns about being included

When fielding this survey, we divided the recipients into two groups to test the impact of presenting the questions in different orders. This was an experiment modeled after one conducted by Merola, Lum, and Murphy (2019) on another policing technology, automated license-plate-recognition systems. The purpose of the experiment was to determine whether being primed with questions about police use of FR technology led respondents to assess their trust in police differently. The first group received the questions above

in the order that they are listed above (A, B, C). In this order, the question about FR systems were presented before questions about trust in police. The second group received the questions about police using FR systems first, followed by the question on trust in police (specifically: B, C, then A). Given random assignment of respondents to the two groups, differences between how the group that was asked about trust in police last and the group that was asked that question first could be inferred to reflect the effect of law enforcement FR use on trust in police. In their experience, Merola, Lum, and Murphy saw statistically significant differences in responses to their question regarding trust in police "to do what is right" based on question order, with individuals asked about license-plate recognition first reporting lower trust. Ultimately, we did not observe the same effect that they observed.

APPENDIX B

Interview Protocol

Background

Experience with facial recognition (or other biometric identification) technologies:

- What experience have you had with facial recognition or other biometric identification technologies? (Probe for different areas of use, particularly criminal justice–related.)

Fundamental differences associated from other methods of identification:

- In your view, how is facial recognition distinct from other methods of identification? What about other biometric methods?

Broad Benefits and Risks of Use: How Do You Think About Balancing the Risks and Benefits of Using FR Technology in Criminal Justice?

Examples of where these technologies served the public interest (e.g., justice):

- In your opinion, in what ways has FR technology benefited society? What about in the future?

Examples of where these technologies did not serve the public interest (e.g., injustice, bias, mistakes) **(highest priority for this section):**

- What concerns do you have about the potential uses of this technology?
- How would you weigh the relative risks and benefits involved in the use of this technology?
- How can biases be identified and/or mitigated?
- What role does public trust play in calculating risks/benefits?

Policy/Regulatory Framework: What Policies Need to Be in Place to Ensure Risks Are Minimized?

Examples of good government policies with these technologies:

- Can concerns be minimized/mitigated through sound policy?
- What policies are most critical for regulating the use of this technology in criminal justice?

- E.g., World Economic Forum policy framework: principles and self-assessment

Examples of undesirable government policies with these technologies:

- Are you aware of any examples of problematic government policies or interventions? What are the main issues you see in those examples?

Regulatory, access, audit, privacy practices (**highest priority of this section):**

- What are the necessary regulatory or oversight elements that should govern the use of this technology?
  - Would certification be a potentially useful route to ensure ongoing regulation?
  - What about limits to access?
  - What about auditing?
  - What about data privacy practices?
  - What about consent options?
    - Are there some uses where consent or lack of consent is not an issue?
    - What about opt-out policies?

Tech Provider

- Description of your software and differentiation from other software
- Perspectives on accuracy, accuracy testing (e.g., NIST), and limits of software
- Potential future uses of software in law enforcement (e.g., live video)
- Guidance for law enforcement
- Guidance for legislators

Law Enforcement

- What software do you use? What was the main reason for choosing that company? Accuracy known?
- What is/was your policy?
- What training is required before use?
- How do you control access to FR software?
- What guidelines are in place for use?
  - Suspicion, probable cause to use?
  - Independent FR analysts, peer review, etc.?

Use Cases

Reactions to researcher provided use/misuse cases:

- Is accuracy a factor here? How so?
  - Is the probe image likely to be problematic in any way?
- Is the source data an issue here?
- Do you have any concerns about the process used here?
- Would this use be appropriate in other circumstances or for other uses?

- What about other methods of identification? Pros and cons in this scenario?

Trends Associated with These Technologies

- Areas of use
- Areas of improvement
- Source data
- Data labeling
- Algorithmic accuracy
- Accuracy measures
- Standards
- Current federal/state/local policies and how they're evolving

Abbreviations

CCTV	closed-circuit television
FACES	Face Analysis Comparison & Examination System
FBI	Federal Bureau of Investigation
FISWG	Facial Identification Subcommittee Working Group
FR	facial recognition
Interpol	International Criminal Police Organization
NIST	National Institute of Standards and Technology

References

Alon-Barkat, Saar, and Madalina Busuioc, "Human-AI Interactions in Public Sector Decision-Making: 'Automation Bias' and 'Selective Adherence' to Algorithmic Advice," *Journal of Public Administration Research and Theory*, February 8, 2022.

American Civil Liberties Union of California, *Making Smart Decisions About Surveillance: A Guide for Communities*, November 2014.

American Civil Liberties Union of Massachusetts, "An Act to Regulate Face Surveillance," undated. As of August 15, 2022:
https://www.aclum.org/sites/default/files/field_documents/regulate_face_surveillance_fact_sheet_final.pdf

Amnesty International, "Surveillance City: NYPD Can Use More Than 15,000 Cameras to Track People Using Facial Recognition in Manhattan, Bronx and Brooklyn," June 3, 2021. As of September 8, 2022:
https://www.amnesty.org/en/latest/news/2021/06/scale-new-york-police-facial-recognition-revealed/

angus mcdiarmid, "Joseph C Ilacqua," October 28, 2012. As of January 20, 2023:
https://www.flickr.com/photos/angusmcdiarmid/8130830700

Anwarul, Shahina, and Susheela Dahiya, "A Comprehensive Review on Face Recognition Methods and Factors Affecting Facial Recognition Accuracy," in *Proceedings of ICRIC 2019*, Springer, 2020.

Bate, Sarah, Rachel Bennetts, Nabil Hasshim, Emma Portch, Ebony Murray, Edwin Burns, and Gavin Dudfield, "The Limits of Super Recognition: An Other-Ethnicity Effect in Individuals With Extraordinary Face Recognition Skills," *Journal of Experimental Psychology: Human Perception and Performance*, Vol. 45, No. 3, 2019.

~BostonBill~, "the usual suspects," March 26, 1959. As of January 20, 2023:
https://www.flickr.com/photos/8533266@N04/4457182603/

Brady v. Maryland, 373 U.S. 83 (1963).

Braga, Anthony A., and Desiree Dusseault, "Can Homicide Detectives Improve Homicide Clearance Rates?" *Crime and Delinquency*, Vol. 64, No. 3, 2018.

Brooks, Julia M., "Drawing the Lines: Regulation of Automatic License Plate Readers in Virginia," *Richmond Journal of Law and Technology*, Vol. 25, No. 3, 2019.

Brush, Kate, "Use Case," TechTarget, last updated November 2022. As of November 29, 2022:
https://www.techtarget.com/searchsoftwarequality/definition/use-case

Bullcoming v. New Mexico, 564 U.S. 647 (2011).

Burke, Melissa Nann, "Michigan Man Wrongfully Accused with Facial Recognition Urges Congress to Act," *Detroit News*, July 13, 2021.

Carpenter v. United States, 138 S. Ct. 2206 (2018).

Celentino, Joseph Clarke, "Face-to-Face with Facial Recognition Evidence: Admissibility Under the Post-Crawford Confrontation Clause," *Michigan Law Review*, Vol. 114, No. 7, 2016.

Conger, Kate, Richard Fausset, and Serge F. Kovaleski, "San Francisco Bans Facial Recognition Technology," *New York Times*, May 14, 2019.

Cook, Cynthia M., John J. Howard, Yevgeniy B. Sirotin, Jerry L. Tipton, and Arun R. Vemury, "Demographic Effects in Facial Recognition and Their Dependence on Image Acquisition: An Evaluation of Eleven Commercial Systems," *IEEE Transactions on Biometrics, Behavior, and Identity Science*, Vol. 1, No. 1, 2019.

Crawford v. Washington, 541 U.S. 36 (2004).

Davis v. Washington, 547 U.S. 813, 822 (2006).

Desmond, Matthew, Andrew V. Papachristos, and David S. Kirk, "Police Violence and Citizen Crime Reporting in the Black Community," *American Sociological Review*, Vol. 81, No. 5, 2016.

Dobbs v. Jackson Women's Health Organization, 597 U.S. ___ (2022).

Ebeling, Ashlea, "IRS to Drop Facial Scan ID.me Verification for Online Accounts," *Forbes*, February 7, 2022.

Engel, Robin S., Michael R. Smith, and Francis T. Cullen, "Race, Place, and Drug Enforcement: Reconsidering the Impact of Citizen Complaints and Crime Rates on Drug Arrests," *Criminology and Public Policy*, Vol. 11, No. 4, November 2012.

Facial Identification Scientific Working Group, "Published Documents," webpage, undated. As of September 4, 2022:
https://fiswg.org/documents.html

Facial Identification Scientific Working Group, "Facial Recognition System: Metadata Usage," Version 1.1, October 25, 2019.

Facial Recognition Commission, homepage, undated. As of September 8, 2022:
https://frcommissionma.com/

Facial Recognition Commission, *Special Commission to Evaluate Government Use of Facial Recognition Technology in the Commonwealth: Final Report*, March 2022.

Fagan, Jeffrey, and Garth Davies, "Street Stops and Broken Windows: Terry, Race, and Disorder in New York City," *Fordham Urban Law Journal*, Vol. 28, No. 2, 2000.

Ferguson, Andrew Guthrie, *The Rise of Big Data Policing: Surveillance, Race, and the Future of Law Enforcement*, New York University Press, 2017.

Ferguson, Andrew Guthrie, "Facial Recognition and the Fourth Amendment," *Minnesota Law Review*, Vol. 105, 2021.

FISWG—*See* Facial Identification Scientific Working Group.

Fredj, Hana Ben, Safa Bouguezzi, and Chokri Souani, "Face Recognition in Unconstrained Environment with CNN," *Visual Computer: International Journal of Computer Graphics*, Vol. 37, No. 2, 2021.

Friedman, Barry, "Lawless Surveillance," *New York University Law Review*, Vol. 97, 2022.

Fuson, Hannah, "Fourth Amendment Searches in First Amendment Spaces: Balancing Free Association with Law and Order in the Age of the Surveillance State," *University of Memphis Law Review*, Vol. 50, 2019.

Fussey, Pete, Bethan Davies, and Martin Innes, "'Assisted' Facial Recognition and the Reinvention of Suspicion and Discretion in Digital Policing," *British Journal of Criminology*, Vol. 61, No. 2, March 2021.

Fussey, Pete, and Daragh Murray, *Independent Report on the London Metropolitan Police Service's Trial of Live Facial Recognition Technology*, Human Rights, Big Data and Technology Project, Human Rights Centre, University of Essex, 2019.

Garvie, Clare, "Garbage in, Garbage out: Face Recognition on Flawed Data," Flawedfacedata.com, May 16, 2019. As of September 4, 2022:
https://www.flawedfacedata.com/

Garvie, Clare, Alvaro Bedoya, and Jonathan Frankle, "The Perpetual Line-Up: Unregulated Police Face Recognition in America," Center on Privacy and Technology, Georgetown Law, October 18, 2016. As of September 4, 2022:
https://www.perpetuallineup.org/risk-framework

Gaston, Shytierra, "Producing Race Disparities: A Study of Drug Arrests Across Place and Race," *Criminology*, Vol. 57, No. 3, 2019.

Gershgorn, Dave, "Exclusive: Live Facial Recognition Is Coming to U.S. Police Body Cameras," *OneZero*, March 5, 2020. As of September 2, 2022:
https://onezero.medium.com/exclusive-live-facial-recognition-is-coming-to-u-s-police-body-cameras-bc9036918ae0

Goldberg, Rebecca Darin, "You Can See My Face, Why Can't I? Facial Recognition and *Brady*," *HRLR Online*, Vol. 5, 2020.

GRC World Forums, "Can European Regulators Stop Clearview AI?" May 27, 2022. As of January 10, 2023:
https://www.grcworldforums.com/can-european-regulators-stop-clearview-ai/5617.article

Grother, Patrick, Austin Hom, Mei Ngan, and Kayee Hanaoka, *Ongoing Face Recognition Vendor Test (FRVT): Part 5; Face Image Quality Assessment*, draft, National Institute of Standards and Technology, September 27, 2022.

Grother, Patrick J., Mei L. Ngan, and Kayee K. Hanaoka, *Ongoing Face Recognition Vendor Test (FRVT): Part 2; Identification*, National Institute of Standards and Technology, NISTIR 8238, November 2018.

Grother, Patrick, Mei Ngan, and Kayee Hanaoka, *Face Recognition Vendor Test (FRVT): Part 2; Identification*, National Institute of Standards and Technology, NISTIR 8271, September 2019a.

Grother, Patrick, Mei Ngan, and Kayee Hanaoka, *Face Recognition Vendor Test (FRVT): Part 3; Demographic Effects*, National Institute of Standards and Technology, NISTIR 8280, December 2019b.

Grother, Patrick, Mei Ngan, and Kayee Hanaoka, *Face Recognition Vendor Test (FRVT): Part 2; Identification*, National Institute of Standards and Technology, NISTIR 8271 Draft Supplement, November 9, 2022.

Grother, Patrick J., Mei L. Ngan, and George W. Quinn, *Face in Video Evaluation (FIVE) Face Recognition of Non-Cooperative Subjects*, National Institute of Standards and Technology, NISTIR 8173, March 2017.

Hahn, Carina A., Liansheng Larry Tang, Amy N. Yates, and P. Jonathon Phillips, "Forensic Facial Examiners vs. Super-Recognizers: Evaluating Behavior Beyond Accuracy," *PsyArXiv Preprints*, August 24, 2021.

Hatmaker, Taylor, "Clearview AI Banned from Selling Its Facial Recognition Software to Most US Companies," *TechCrunch*, May 9, 2022. As of January 10, 2023:
https://techcrunch.com/2022/05/09/clearview-settlement-bipa/

Hawkins, Derek, "Researchers Use Facial Recognition Tools to Predict Sexual Orientation. LGBT Groups Aren't Happy," *Washington Post*, September 12, 2017.

Heaton, Paul, "Hidden in Plain Sight: What Cost-of-Crime Research Can Tell Us About Investing in Police," RAND Corporation, OP-279, 2010. As of December 15, 2022:
https://www.rand.org/pubs/occasional_papers/OP279.html

Hill, Dallas, Christopher D. O'Conner, and Andrea Slane, "Police Use of Facial Recognition Technology: The Potential for Engaging the Public Through Co-Constructed Policy-Making," *International Journal of Police Science and Management*, Vol. 24, No. 3, 2022.

Hill, Kashmir, "Activists Turn Facial Recognition Tools Against the Police," *New York Times*, October 21, 2020.

Hill, Kashmir, "How One State Managed to Actually Write Rules on Facial Recognition," *New York Times*, March 5, 2021.

Ho, Daniel E., Emily Black, Maneesh Agrawala, and Li Fei-Fei, "Evaluating Facial Recognition Technology: A Protocol for Performance Assessment in New Domains," *Denver Law Review*, Vol. 98, 2020.

Howard, John J., Laura R. Rabbitt, and Yevgeniy B. Sirotin, "Human-Algorithm Teaming in Face Recognition: How Algorithm Outcomes Cognitively Bias Human Decision-Making," *PloS One*, Vol. 15, No. 8, 2020.

Howard, John J., Yevgeniy B. Sirotin, and Arun R. Vemury, "The Effect of Broad and Specific Demographic Homogeneity on the Imposter Distributions and False Match Rates in Face Recognition Algorithm Performance," in *2019 IEEE 10th International Conference on Biometrics Theory, Applications and Systems*, IEEE, 2019.

Hu, Ying, Kelsey Jackson, Amy Yates, David White, P. Jonathon Phillips, and Alice J. O'Toole, "Person Recognition: Qualitative Differences in How Forensic Face Examiners and Untrained People Rely on the Face Versus the Body for Identification," *Visual Cognition*, Vol. 25, Nos. 4–6, 2017.

Human Rights Watch, "US: Police Found to Violate Fellow Officer's Privacy: Minnesota Case Shows Need for Stronger Data Protection Laws," June 20, 2019. As of September 2, 2022:
https://www.hrw.org/news/2019/06/20/us-police-found-violate-fellow-officers-privacy

Hunt, Priscillia, James Anderson, and Jessica Saunders, "The Price of Justice: New National and State-Level Estimates of the Judicial and Legal Costs of Crime to Taxpayers," *Journal of Criminal Justice*, Vol. 42, No. 2, 2017.

Hunt, Priscillia Evelyn, Jessica Saunders, and Beau Kilmer, "Estimates of Law Enforcement Costs by Crime Type for Benefit-Cost Analyses," *Journal of Benefit Cost Analysis*, Vol. 10, No. 1, 2019.

Illinois Compiled Statutes, Biometric Information Privacy Act, 740 ILCS 14/, October 3, 2008. As of September 2, 2022:
https://www.ilga.gov/legislation/ilcs/ilcs3.asp?ActID=3004&ChapterID=57

International Association of Chiefs of Police, Law Enforcement Policy Center, *Facial Recognition*, February 2022.

International Criminal Police Organization, United Nations Interregional Crime and Justice Research Institute, Police of the Netherlands, and World Economic Forum, "Policy Framework for Responsible Limits on Facial Recognition, Use Case: Law Enforcement Investigations," white paper, World Economic Forum, September 2021.

Interpol—*See* International Criminal Police Organization.

Jackson, Kaitlin, "Challenging Facial Recognition Software in Criminal Court," *The Champion* (National Association of Criminal Defense Lawyers), July 2019.

Johnson, Khari, "The Hidden Role of Facial Recognition Tech in Many Arrests," *Wired*, March 7, 2022.

Johnson, Thaddeus L., Natasha N. Johnson, Denise McCurdy, and Michael S. Olajide, "Facial Recognition Systems in Policing and Racial Disparities in Arrests," *Government Information Quarterly*, Vol. 39, No. 4, 2022.

Kahn, Jeremy, "After Roe, Fears Mount About A.I.'s Ability to Identify Those Seeking Abortions," *Fortune*, June 28, 2022.

Katz v. United States, 389 U.S. 347 (1967).

Kindt, Els J., "Transparency and Accountability Mechanisms for Facial Recognition," German Marshall Fund, February 3, 2021. As of September 8, 2022:
https://www.gmfus.org/news/transparency-and-accountability-mechanisms-facial-recognition

Kirk, David S., and Andrew V. Papachristos, "Cultural Mechanisms and the Persistence of Neighborhood Violence," *American Journal of Sociology*, Vol. 116, No. 4, 2011.

Kobie, Nicole, "The Complicated Truth About China's Social Credit System," *Wired*, July 6, 2019.

Labati, Ruggero Donida, Danilo De Angelis, Barbara Bertoglio, Cristina Cattaneo, Fabio Scotti, and Vincenzo Piuri, "Automatic Face Recognition for Forensic Identification of Persons Deceased in Humanitarian Emergencies," *IEEE 2021 International Conference on Computational Intelligence and Virtual Environments for Measurement Systems and Applications (CIVEMSA)*, IEEE, 2021.

Law Enforcement Imaging Technology Task Force, *Law Enforcement Facial Recognition Use Case Catalog*, IJIS Institute and the International Association of Chiefs of Police, March 2019.

Lee, Nicol Turner, Paul Resnick, and Genie Barton, *Algorithmic Bias Detection and Mitigation: Best Practices and Policies to Reduce Consumer Harms*, Brookings Institute, May 22, 2019.

Lee, Won-Joon, Caroline Wilkinson, Amina Memon, and Kate Houston, "Matching Unfamiliar Faces from Poor Quality Closed-Circuit Television (CCTV) Footage: An Evaluation of the Effect of Training on Facial Identification Ability," *Axis*, Vol. 1, No. 1, Summer 2009.

Lewis, James A., and William Crumpler, *Facial Recognition Technology: Responsible Use Principles and the Legislative Landscape*, Center for Strategic and International Studies, September 2021.

Lofstrom, Magnus, Justin Goss, Joseph Hayes, and Brandon Martin, "Racial Disparities in California Arrests," fact sheet, Public Policy Institute of California, 2019.

Louradour, Sebastien, and Lofred Madzou, "A Policy Framework for Responsible Limits on Facial Recognition: Use Case; Law Enforcement Investigations," white paper, World Economic Forum, October 2021.

Lu, Peng, Baoye Song, and Lin Xu, "Human Face Recognition Based on Convolutional Neural Network and Augmented Dataset," *Systems Science and Control Engineering*, Vol. 9, No. 2, 2021.

Lynch, Jennifer, "Face Off: Law Enforcement Use of Face Recognition Technology," Electronic Frontier Foundation, May 2019.

Lynch v. State, 260 So. 3d 1166 (Fla. Dist. Ct. App. 2018).

Major Cities Chiefs Association Facial Recognition Working Group, "Facial Recognition Technology in Modern Policing: Recommendations and Considerations," Major Cities Chiefs Association, 2021.

Mares, Dennis, and Emily Blackburn, "Evaluating the Effectiveness of an Acoustic Gunshot Location System in St. Louis, MO," *Policing: A Journal of Policy and Practice*, Vol. 6, No. 1, 2012.

Mares, Dennis, and Emily Blackburn, "Acoustic Gunshot Detection Systems: A Quasi-Experimental Evaluation in St. Louis, MO," *Journal of Experimental Criminology*, Vol. 17, No. 2, 2021.

Massachusetts Bill S.2963, An Act Relative to Justice, Equity and Accountability in Law Enforcement in the Commonwealth, December 31, 2020.

Melendez-Diaz v. Massachusetts, 557 U.S. 305, 311 (2009).

Menon, Sachit, Alexandru Damian, Shijia Hu, Nikhil Ravi, and Cynthia Rudin, "Pulse: Self-Supervised Photo Upsampling via Latent Space Exploration of Generative Models," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, 2020.

Merola, Linda M., Cynthia Lum, and Ryan P. Murphy, "The Impact of License Plate Recognition Technology (LPR) on Trust in Law Enforcement: A Survey-Experiment," *Journal of Experimental Criminology*, Vol. 15, No. 1, March 2019.

Metz, Rachel, "After Face-Recognition Backlash, ID.me Says Government Agencies Will Get More Verification Options," CNN, February 9, 2022.

Miikka H, "21042007," April 21, 2007. As of January 20, 2023:
https://www.flickr.com/photos/miikkah/467234390/

Moy, Laura M., "A Taxonomy of Police Technology's Racial Inequity Problems," *University of Illinois Law Review*, Vol. 139, No. 1, 2021.

NAACP v. Alabama ex rel. Patterson, 357 U.S. 449 (1958).

National Institute of Standards and Technology, "FRVT Quality Assessment—Specific Image Defect Detection: An Ongoing Evaluation of Software That Checks for Quality Problems in Face Images," August 19, 2022. As of September 8, 2022:
https://pages.nist.gov/frvt/api/FRVT_ongoing_quality_sidd_api.pdf

Ngan, Mei, Patrick Grother, Kayee Hanaoka, and Jason Kuo, *Face Recognition Vendor Test (FRVT): Part 4; Morph-Performance of Automated Face Morph Detection*, National Institute of Standards and Technology, NISTIR 8292, March 6, 2020.

Nguyen, Alexander T., "Here's Looking at You, Kid: Has Face-Recognition Technology Completely Outflanked the Fourth Amendment?" *Virginia Journal of Law and Technology*, Vol. 7, No. 2, Spring 2002.

NIST—*See* National Institute of Standards and Technology.

Olszewska, Joanna Isabelle, "Automated Face Recognition: Challenges and Solutions," in Srinivasan Ramakrishnan, ed., *Pattern Recognition-Analysis and Applications*, Intechopen, 2016.

Ousey, Graham C., and Matthew R. Lee, "To Know the Unknown: The Decline in Homicide Clearance Rates, 1980–2000," *Criminal Justice Review*, Vol. 35, No. 2, 2010.

Ouss, Aurélie, and John Rappaport, "Is Police Behavior Getting Worse? Data Selection and the Measurement of Policing Harms," *Journal of Legal Studies*, Vol. 49, No. 1, 2020.

Papachristos, Andrew V., Tracey L. Meares, and Jeffrey Fagan, "Why Do Criminals Obey the Law? The Influence of Legitimacy and Social Networks on Active Gun Offenders," *Journal of Criminal Law and Criminology*, Vol. 102, No. 2, Spring 2012.

Payal, Parekh, and Mahesh M. Goyani, "A Comprehensive Study on Face Recognition: Methods and Challenges," *Imaging Science Journal*, Vol. 68, No. 2, 2020.

Peeters, Rik, "The Agency of Algorithms: Understanding Human-Algorithm Interaction in Administrative Decision-Making," *Information Polity*, Vol. 25, No. 4, 2020.

People v. Beckley, 185 Cal.App.4th 509 (Cal. Ct. App. 2010).

Perpetual Line-Up, "Jurisdiction: Federal Bureau of Investigation," webpage, last updated September 2016a. As of September 8, 2022:
https://www.perpetuallineup.org/jurisdiction/federal-bureau-investigation

Perpetual Line-Up, "Jurisdiction: Florida," webpage, last updated September 2016b. As of September 8, 2022: https://www.perpetuallineup.org/jurisdiction/florida

Perritt, Henry H., Jr., "Defending Face-Recognition Technology (and Defending Against It)," *Journal of Technology Law and Policy*, Vol. 25, No. 1, 2021.

Perry v. New Hampshire, 565 U.S. 228 (2012).

Phillips, P. Jonathon, Amy N. Yates, Ying Hu, Carina A. Hahn, Eilidh Noyes, Kelsey Jackson, Jacqueline G. Cavazos, Géraldine Jeckeln, Rajeev Ranjan, et al., "Face Recognition Accuracy of Forensic Examiners, Superrecognizers, and Face Recognition Algorithms," *Proceedings of the National Academy of Sciences*, Vol. 115, No. 24, 2018.

Policing Project, "Regulating Police Use of Face Recognition Technology—Resources for Legislators," webpage, undated. As of January 10, 2023:
https://www.policingproject.org/regulating-police-use-of-face-recognition-technology

Rainie, Lee, Cary Funk, Monica Anderson, and Alec Tyson, "Public More Likely to See Facial Recognition Use by Police as Good, Rather Than Bad for Society," in *AI and Human Enhancement, Americans' Openness Is Tempered by a Range of Concerns*, Pew Research Center, March 2022.

Ramon, Meike, Anna K. Bobak, and David White, "Super-Recognizers: From the Lab to the World and Back Again," *British Journal of Psychology*, Vol. 110, No. 3, 2019.

RAND Corporation, "RAND American Life Panel," webpage, undated. As of September 8, 2022:
https://www.rand.org/research/data/alp.html

Reuters, "U.S. Cities Are Backing Off Banning Facial Recognition as Crime Rises," *U.S. News*, May 12, 2022.

Revised Code of Washington, Chapter 43.386, Facial Recognition. As of December 15, 2022:
https://app.leg.wa.gov/RCW/default.aspx?cite=43.386

Richardson, Rashida, "Facial Recognition in the Public Sector: The Policy Landscape," German Marshall Fund, February 3, 2021. As of September 8, 2022:
https://www.gmfus.org/news/facial-recognition-public-sector-policy-landscape

Rodricks, Dan, "A Data Dive on Baltimore Homicides Shows the Need to Stop Retaliatory Violence," *Baltimore Sun*, March 9, 2021.

Rodriguez, Roger, "Facial Recognition: Art or Science?" Vigilant Solutions, undated. As of September 2, 2022:
https://www.sheriffs.org/sites/default/files/Whitepaper%20Facial%20Recognition.pdf

Rollet, Charles, "London Police Facial Recognition Struggles to Find Suspects," IPVM, July 28, 2022. As of December 4, 2022:
https://ipvm.com/reports/lfr-london

Rudin, Cynthia Diane, and Shawn D. Bushway, "A Truth Serum for Your Personal Perspective on Facial Recognition Software in Law Enforcement," *Translational Criminology Magazine*, Fall 2021.

Sankaran, Nishant, Sergey Tulyakov, Srirangaraj Setlur, and Venu Govindaraju, "Metadata-Based Feature Aggregation Network for Face Recognition," *IEEE 2018 International Conference on Biometrics (ICB)*, IEEE, 2018.

Saunders, Jessica, Priscillia Hunt, and John S. Hollywood, "Predictions Put into Practice: A Quasi-Experimental Evaluation of Chicago's Predictive Policing Pilot," *Journal of Experimental Criminology*, Vol. 12, No. 3, 2016.

Saxena, Devansh, Karla Badillo-Urquiola, Pamela J. Wisniewski, and Shion Guha, "A Framework of High-Stakes Algorithmic Decision-Making for the Public Sector Developed Through a Case Study of Child-Welfare," *Proceedings of the ACM on Human-Computer Interaction*, Association for Computing Machinery, 2021.

Schuppe, Jon, "Cellphone Dragnet Used to Find Bank Robbery Suspect Was Unconstitutional, Judge Says," NBC News, March 7, 2022.

Senate Bill 2052, Facial Recognition and Biometric Technology Moratorium Act of 2021, referred to the Committee on the Judiciary, June 15, 2021.

Sheard, Nathan, and Adam Schwartz, "The Movement to Ban Government Use of Face Recognition," Electronic Frontier Foundation, May 5, 2022. As of September 2, 2022:
https://www.eff.org/deeplinks/2022/05/movement-ban-government-use-face-recognition

Shelton v. Tucker, 364 U.S. 479 (1960).

Siegel, Eric, "When Does Predictive Technology Become Unethical?" *Harvard Business Review*, October 23, 2020.

Simpson, Joanne Cavanaugh, and Marc Freeman, "South Florida Police Quietly Ran Facial Recognition Scans to Identify Peaceful Protestors. Is That Legal?" *Sun Sentinel*, June 26, 2021.

Slaight, Jillian and Ryan LeCloux, "Facial Recognition Technology: Balancing Safety and Privacy," *Wisconsin Policy Project*, Vol. 4, No. 3, March 2020.

Smart, Rosanna, Andrew R. Morral, and Terry L. Schell, "Evaluating Baltimore's Aerial Investigation Research Pilot Program," RAND Corporation, WR-A1131-1, 2022. As of November 29, 2022:
https://www.rand.org/pubs/working_papers/WRA1131-1.html

Spivack, Jameson, and Clare Garvie, "A Taxonomy of Legislative Approaches to Face Recognition in the United States," in Amba Kak, ed., *Regulating Biometrics: Global Approaches and Urgent Questions*, AI Now Institute, September 2020.

Tolosana, Ruben, Sergio Romero-Tapiador, Ruben Vera-Rodriguez, Ester Gonzalez-Sosa, and Julian Fierrez, "DeepFakes Detection Across Generations: Analysis of Facial Regions, Fusion, and Performance Evaluation," *Engineering Applications of Artificial Intelligence*, Vol. 110, April 2022.

Turner v. United States, 248 U.S. 354 (1919).

United States v. Jones, 565 U.S. 400 (2012).

United States v. Maynard, 615 F.3d 544, 562 (D.C. Cir. 2010).

U.S. Department of Homeland Security and Bureau of Justice Assistance, U.S. Department of Justice, *Face Recognition: Policy Development Template for Use in Criminal Intelligence and Investigation Activities*, December 2017.

U.S. Government Accountability Office, *Forensic Technology: Algorithms Strengthen Forensic Analysis, but Several Factors Can Affect Outcomes*, GAO-21-4335SP, July 6, 2021a.

U.S. Government Accountability Office, *Facial Recognition Technology: Federal Law Enforcement Agencies Should Have Better Awareness of Systems Used by Employees*, GAO-21-105309, July 13, 2021b.

Vincent, James, "Liveness Tests Used by Banks to Verify ID Are 'Extremely Vulnerable' to Deepfake Attacks," *The Verge*, May 18, 2022. As of September 2, 2022:
https://www.theverge.com/2022/5/18/23092964/
deepfake-attack-facial-recognition-liveness-test-banks-sensity-report

Virginia Senate Bill 741, Facial Recognition Technology; Authorized Uses, April 11, 2022.

Wang, Xiang, Kai Wang, and Shiguo Lian, "A Survey on Face Data Augmentation for the Training of Deep Neural Networks," *Neural Computing and Applications*, Vol. 32, No. 19, 2020.

Washington Senate Bill 6280, Concerning the Use of Facial Recognition Services, July 1, 2021.

WCJB Staff, "Florida Lawmakers Look to Crack Down on Misuse of DAVID System," WCJB TV20 News, March 18, 2021.

Weitzer, Ronald, "Incidents of Police Misconduct and Public Opinion," *Journal of Criminal Justice*, Vol. 30, No. 5, 2002.

Weitzer, Ronald, "American Policing Under Fire: Misconduct and Reform," *Society*, Vol. 52, No. 5, 2015.

White, Clair, David Weisburd, and Sean Wire, "Examining the Impact of the Freddie Gray Unrest on Perceptions of the Police," *Criminology and Public Policy*, Vol. 17, No. 4, 2018.

White, David, James D. Dunn, Alexandra C. Schmid, and Richard I. Kemp, "Error Rates in Users of Automatic Face Recognition Software," *PloS One*, Vol. 10, No. 10, 2015.

Zhuchkov, Aleksei, "Analyzing the Effectiveness of Image Augmentations for Face Recognition from Limited Data," *IEEE 2021 International Conference "Nonlinearity, Information and Robotics" (NIR)*, IEEE, 2021.

About the Authors

Samuel Peterson is a policy researcher at the RAND Corporation. He has broad interests in topics within and outside criminal justice, and he has research experience in gun violence prevention, sexual assault at American colleges and universities, school-based victimization among adolescents, and substance use within adolescent social networks. Peterson completed his Ph.D. in criminal justice at the University of Cincinnati.

Brian A. Jackson is a senior physical scientist at the RAND Corporation. His research focuses on homeland security, criminal justice, and emergency preparedness. Jackson received his Ph.D. in bioinorganic chemistry from the California Institute of Technology.

Dulani Woods is a data scientist at the RAND Corporation and codirector of the RAND Center for Scalable Computing and Analysis. He is adept at data acquisition, transformation, visualization, modeling, simulation, optimization, and other methods. He holds an M.S. in agricultural economics (applied economics) from Purdue University.

Elina Treyger is a political scientist at the RAND Corporation. Her work at RAND spans topics in three areas of research: homeland security, national security and defense, and justice policy. She received a J.D. from Harvard Law School and a Ph.D. in political science from Harvard University.