

Inside the Black Box: Algorithmic Policing and the Right to Know

Madisyn Schackow*

Abstract

Artificial intelligence is increasingly being used in the American criminal justice system. These predictive algorithms shape police decision-making, identification, and evidentiary uses. These programs claim efficiency and objectivity, but they operate without bias correction and reinforce stereotyped policing. They also operate without any transparency framework, leaving citizens unaware both of the use of AI and of how it affects them. The current framework was designed to regulate human decision-making, not the predictive algorithms that are currently being utilized across the United States.

This article examines how the use of artificial intelligence in criminal justice implicates core constitutional protections, including Procedural Due Process, Equal Protection, and evidentiary reliability standards. The United States' current framework fails to provide adequate safeguards for AI usage in policing. Transparency, disclosure and reliability requirements grounded in existing constitutional principles would ensure that the integration of AI systems would enhance, rather than undermine, the legitimacy of the criminal justice system.

Keywords: predictive policing, transparency, algorithmic due process, equal protection, artificial intelligence in criminal justice

A Introduction

On a cold winter morning in Detroit, Michigan, Robert Williams was awoken by the sound of police officers banging on his front door. Upon opening his door, Williams was arrested for a crime he did not commit. Police used an artificial-intelligence (AI) based facial-recognition system that incorrectly identified Williams as the suspect in blurry surveillance footage. Williams was forced to spend thirty hours in custody before detectives admitted that the algorithm had been wrong.¹ Williams' experience did not end upon his release from custody, as he later sued the City of Detroit.² The litigation concluded in 2024, with a settlement that imposed restrictions on the Detroit Police

* Madisyn Schackow, J.D. expected May 2027, Indiana University Robert H. McKinney School of Law, Indianapolis.

1 American Civil Liberties Union, 'Williams v. City of Detroit: Settlement Summary', 28 June 2024 (describing the wrongful arrest of Robert Williams on the basis of a facial recognition 'match').

2 Ibid.

Department's use of facial-recognition technology.³ These restrictions no longer allow officers to make arrests based solely on the results acquired from facial-recognition technology.⁴ Officers are now required to obtain independent evidence connecting the identified individual to the crime.⁵ Further, the settlement imposed additional auditing and training requirements governing the use of the AI technology.⁶ These reforms demonstrate that facial-recognition technology can be subjected to meaningful procedural safeguards without requiring a complete prohibition of its use. Williams' case provides a meaningful example of the harms that arise from relying solely on algorithmic identification and the availability of imposing enforceable limitations.

AI algorithms have moved from Silicon Valley laboratories into the daily lives of almost everyone with access to technology. Police departments across the United States are using AI systems that claim to predict crime before it occurs. These systems use criminal records, service requests, and social network activity to predict crime rates. Police then use this data to determine "hot spots" for crime, create personalized risk assessments, and direct the deployment of officers. Predictive technologies have gradually changed policing strategies, but the public remains largely in the dark about how the systems operate, what they encode, and how they influence encounters between the public and law enforcement. Due to historic institutional secrecy, trade-secret claims, and law-enforcement exemptions in public-information regulations, these algorithms frequently evade investigation. Because of this, the populations most impacted by policing are often the ones least able to challenge the technology.

This is not simply a technical oversight; secrecy is doing work of its own. Predictive systems rely on historical data that has repeatedly been shown to be incomplete, racially biased, or produced in violation of the Constitution. These errors go unchecked and are incorporated into automated decision-making because the public cannot view the data or models. The consequences fall particularly heavily on persons of color. In areas that are already overpoliced, algorithmically driven patrol routes increase monitoring, and AI feedback loops increase racial disparities that are already present. People can be classified as "high risk" without any way to find out or deny the claim. Courts struggle to confront these developments under Fourth Amendment and due-process frameworks.

Without an enforceable right to know, predictive policing is incompatible with constitutional protections. Withholding information about these systems directly hinders the ability of communities and courts to evaluate the constitutionality of their use. The technology shapes policing without any means of identifying bias or assessing accuracy. This affects procedural due process, equal protection, and evidentiary claims. The idea that government authority is based on knowable, contestable, and accountable foundations is undermined by algorithmic secrecy.

Part I of this article outlines the use of predictive-policing technologies and data-driven law enforcement. Part II explains the tools that currently limit notice

3 Ibid.

4 Ibid.

5 Ibid.

6 Ibid.

requirements for AI systems. Part III explains the right to know found in constitutional doctrine, including the Fourth Amendment, procedural due process, equal protection, and the right to challenge evidence. Part IV describes potential reforms that promote transparency, like statutory changes, contracting requirements, and judicial standards. Transparency is a central condition for the use of AI algorithms in law enforcement, because without it accountability measures fail and constitutional protections are hollowed out. Without a meaningful right to know, predictive policing functions as an unreviewable black box that threatens to reproduce and intensify inequalities within the justice system.

B Predictive Policing and Algorithm Enforcement

Predictive policing refers to the use of algorithmic tools trained on historical crime data to anticipate future criminal activity, thereby shaping police officers' decisions. Traditional crime analysis relied on officer experience and training and used publicly reported crime statistics. Modern policing has increasingly begun to use AI tools to anticipate future crime, inform patrol strategies, shape investigative efforts, and assess risk. These systems are now embedded in routine policing across the United States, and often operate with minimal public awareness.⁷

I *Types of Predictive Policing*

Place-based prediction and person-based prediction are two different predictive policing technologies. There is a key distinction between place and person-based prediction, because they affect policing in different ways. Place-based systems operate structurally, by influencing the frequency and locations of police encounters. They identify geographic areas where crime is statistically likely to occur during specific times, producing "hot spots" that direct police patrols.⁸ This functions as a form of generalized suspicion: a stop ordinarily requires individualized suspicion of unlawful conduct, not mere presence in a location.

Person-based systems present the greater danger. Person-based systems attempt to assess the likelihood that specific individuals will engage in criminal activity, often by assigning risk scores created from the individual's arrest history, social network activity, and prior police contact.⁹ These technologies use data that often reflects historical enforcement disparities, rather than clear criminal behavior. As a result, these systems use a person's criminal history to justify future surveillance. Vendors and police departments describe these systems as neutral decision-support tools, but their outputs materially influence how officers allocate time and make decisions while working.¹⁰ Unlike traditional tools, these systems

7 Ferguson, 'Policing Predictive Policing', 94 *Wash. U. L. Rev.* (2017) 1109, at 1165.

8 *Ibid.*, at 1120-1122.

9 Ferguson, 'Predictive Policing and Reasonable Suspicion', 62 *Emory L.J.* (2012) 259, at 266.

10 Joh, 'The New Surveillance Discretion: Automated Suspicion, Big Data, and Policing', 10 *Harv. L. & Pol'y Rev.* (2016) 15, at 31.

focus on managing risk rather than responding to wrongdoing. This shifts policing from punishment to prevention.

Although policing used other data-driven models before AI, the systems were far less capable in terms of scale and function. Predictive systems now apply statistical and machine-learning techniques to large datasets, by methods often opaque even to the operator. These systems claim to generate objective forecasts rather than descriptive summaries.¹¹ As a result, the algorithmic outputs are increasingly substituted for individual judgment of officers, changing how suspicion and risk are used within policing.¹² For example, when officers are working under time constraints or internal pressure, they are more likely to defer to these algorithmic predictions, so that reliance on these tools becomes routine over time.

The use of these systems without transparency to those policed will change policing fundamentally. AI-generated predictions will begin to replace human judgment, little will be knowable about how decisions are made, and remedies will be hard to obtain without clear evidence of harm. The constitutional requirement of individualized suspicion becomes unworkable if police decisions rest on, and are shielded by, opaque algorithms.

II Data Inputs and Feedback Loops

The predictions made by AI systems are based almost entirely on historical policing data, including arrest records, calls for service, stop-and-frisk reports, and incident logs. This data reflects criminal activity, but also enforcement choices made by officers, which could have been shaped by biases about race, neighborhood, and gender.¹³ When predictive systems rely on these datasets without bias correction or contextualization, they reproduce existing biases.

Where algorithmically identified areas are more heavily monitored, more incidents are recorded. That teaches the system that those areas require increased surveillance, whether or not it is true.¹⁴ Predictive policing frequently reinforces past enforcement biases under the pretense of data neutrality rather than revealing new criminal tendencies.¹⁵ The automation of this process entrenches the subjective, and often biased, judgments that the models were trained on. Therefore, the new data created reinforces the stereotypes, making it more difficult to spot or contest discriminatory results.¹⁶

III Algorithmic Influence on Policing Decisions

Law-enforcement agencies claim that predictive algorithms do not make final decisions, but the recommendations they provide significantly shape the way officers

11 Ferguson, *supra* note 2, at 1144.

12 Ferguson, *supra* note 4, at 288, 298.

13 Richardson, 'Jason M. Schultz & Kate Crawford, Dirty Data, Bad Predictions: How Civil Rights Violations Impact Police Data, Predictive Policing Systems, and Justice', 94 *N.Y.U. L. Rev.* (2019) 15, at 29.

14 Ferguson, *supra* note 2, at 1149.

15 Richardson, *supra* note 8, at 47.

16 Joh, *supra* note 5, at 36.

act.¹⁷ The predictive outputs contribute to patrol allocation, stop decisions, and investigative focus, even if they should be based on officer discretion.¹⁸ The Fourth Amendment has traditionally required that police have individualized, observable facts to justify suspicion before stopping or searching civilians.¹⁹ Predictive policing complicates this framework because the assessments are not directly observable or understandable. Officers themselves may be unable to explain why a location or individual was flagged without relying on the system's output entirely. This raises questions about whether an AI's suspicion satisfies the requirement of specific and objective grounds for intrusion by police.

The influence of predictive systems extends beyond direct interactions between police and citizens. The risk assessments generated during investigations may follow individuals through charging decisions, pretrial determinations, and sentencing, thus affecting liberty during multiple stages of the criminal process.²⁰ For example, a "high risk" individual may face heightened scrutiny during the police investigation, a reduction the likelihood of pretrial release, an increase bond amounts, or face harsher sentencing outcomes. While these are all individual decisions, they create a compounding effect on the accused, magnifying their consequences without regard for bias or error. Despite the significant impact AI systems have, their design, validation, and operation are rarely disclosed to defendants, courts, or the public. This leaves defendants unable to examine or fight against systems that are making decisions about their rights. The disconnect between usage and transparency raises both legal and constitutional concerns by allowing police officers to evade transparency and contestability.

C The Transparency Problem: Secrecy, The Right to Know, and Public Oversight

Despite the growing use of predictive policing systems in law enforcement, the technologies are institutionally protected from public scrutiny. This is purposeful and reinforced by doctrinal protections, institutional practices, and statutory exemptions. As a result, affected individuals and the broader public cannot reliably determine when predictive systems are in use, what data they rely on, or how their outputs influence policing.

I Algorithms and the Trade-Secret Doctrine

A central source of ambiguity in predictive policing is the reliance on private vendors to design, maintain, and update the algorithmic systems. Police departments frequently use third-party companies that claim ownership over the models, methodologies, and data-processing techniques used in the predictive tools.²¹ These ven-

17 Ferguson, *supra* note 2, at 1115.

18 *Ibid.*, at 1138.

19 *Terry v. Ohio*, 392 U.S. 1, 21 (1968).

20 Garrett and Rudin, 'The Right to a Glass Box: Rethinking the Use of Artificial Intelligence in Criminal Justice', 109 *Cornell L. Rev.* (2024) 561, at 564-565.

21 Ferguson, *supra* note 2, at 1131.

dors typically invoke trade-secret protections to deny disclosure of their methods. These vendors argue that transparency would compromise their interests and give their competitors advantages and access to their technology.²²

The trade-secret doctrine was initially designed to protect commercial innovation in private markets and is unfit to govern technologies that perform core governmental functions.²³ When applied to predictive policing, trade-secret protections allow private actors to dictate the scope of public access to information about how police power is exercised.²⁴ This gives vendors the ability to deny disclosure of critical information, even when it may infringe upon constitutional rights. The outsourcing of predictive analytics creates secrecy and removes accountability in favor of protecting market confidentiality.

II Public Records Law and Law-Enforcement Secrecy

Public-records statutes, including the federal Freedom of Information Act (FOIA) and state equivalents, are commonly used for transparency in government decision-making.²⁵ In theory, these laws provide the public with access to records about government programs, policies, and the technology used.²⁶ Predictive policing systems often invoke law enforcement exemptions, allowing agencies to withhold documents related to investigative techniques, internal practices, or public safety.²⁷

Historically, courts have given significant leniency to law enforcement organizations when assessing whether disclosure will jeopardize investigations or facilitate evasion of policing techniques.²⁸ This deference has been extended to algorithmic tools, even when agencies cannot demonstrate any harm that would result from disclosure.²⁹ The result is a form of law-enforcement privilege that places predictive systems beyond the reach of meaningful public oversight. The laws governing public records were not created to address predictive algorithms, and the limitations are clear when agencies invoke both investigatory and proprietary exemptions.³⁰

Aside from legal protections, predictive policing is further safeguarded by institutional norms that favor secrecy over transparency. Without litigation or legislative requirements, police departments rarely disclose the use of predictive systems.³¹ When agencies do acknowledge the use of predictive tools, the disclosures

22 Citron and Pasquale, 'The Scored Society: Due Process for Automated Predictions', 89 *Wash. L. Rev.* (2014) 1, at 25.

23 *Ibid.*, at 20-21.

24 *Ibid.*

25 Freedom of Information Act, 5 U.S.C. § 552 (2018).

26 *Ibid.*

27 *Milner v. Department of the Navy*, 562 US 562 (2011), at 567 (explaining that Exemption 7(E) protects against disclosure that could enable law evasion).

28 *Ibid.*

29 Citron and Pasquale, *supra* note 17, at 26.

30 Wexler, 'Life, Liberty, and Trade Secrets: Intellectual Property in the Criminal Justice System', 70 *Stan. L. Rev.* (2018) 1343, at 1367.

31 Friedman and Jánosky, 'Policing's Information Problem', 99 *Tex. L. Rev.* (2020) 1, at 43.

are often high-level descriptions that exclude details about data inputs, validation studies, or error rates.³² This partial disclosure gives the appearance of accountability but does not actually provide meaningful information.

Without access to details about the predictive systems, the ability to assess claims of disparate impact, specifically targeting neighborhoods or demographic groups, is impossible.³³ Defendants are unable to argue how the algorithmic outputs influenced officers' decisions to stop, arrest, or surveil them, essentially leaving them defenseless.³⁴ Similarly, the court system lacks modern precedent needed to evaluate constitutional challenges to the predictive policing practices.³⁵ In this environment, secrecy functions as a structural barrier to legal accountability.

III Transparency as a Precondition to Legitimacy

The effect of trade-secret doctrine, public-records exemptions, and institutional secrecy is a transparency gap that undermines public accountability and governance. This transparency gap exacerbates existing power imbalances and secrecy between law-enforcement agencies and the communities they intend to protect.³⁶ Transparency cannot remain optional if predictive policing is to be balanced with the Constitution. Access to information regarding algorithmic policing should become a legal requirement for both effective oversight and constitutional legitimacy.

D Constitutional Foundations

Constitutional jurisprudence is often underscored by the belief that people have access to enough information to comprehend, challenge, and restrict governmental activity when it affects their liberty. The Constitution itself does not explicitly mention a "right to know," but transparency is a precondition from the Fourth Amendment, equal protection, procedural due process, and the evaluation of evidence. The transparency gap undermines procedural due process claims, because they require access to information about the effect of the system's use in practice. Similarly, it affects equal protection claims because proof of discriminatory use or intent is required and could only be proven by the algorithm itself. Finally, when dealing with evidentiary matters, the methods and tools used to obtain evidence are necessary to determine admissibility and weight. Because predictive policing systems are shielded from disclosure, these legal safeguards become functionally dormant.

I The Fourth Amendment and Privacy Protections

Recent Fourth Amendment decisions involving the use of digital information illustrate why technological algorithms cannot place police decision-making outside of

32 Ferguson, *supra* note 2, at 1153.

33 Citron and Pasquale, *supra* note 17, at 27-28, 31.

34 Garrett and Rudin, *supra* note 15, at 564-565.

35 Bloch-Wehba, 'Visible Policing: Technology, Transparency, and Democratic Control', 109 *Calif. L. Rev.* (2021) 917, at 936-937.

36 Joh, *supra* note 5, at 38.

the ordinary constitutional review. In *Carpenter v. United States*, the Supreme Court held that individuals maintain a reasonable expectation of privacy regarding their historical cell-site location data, even though that information is maintained by third-party service providers.³⁷ The Court explained that although the information is held by a third party, it cannot overcome the user's claim to Fourth Amendment protections.³⁸ The Court emphasized that a person "does not surrender all Fourth Amendment protection by venturing in the public sphere."³⁹ Further, the Court explained that society believes that law enforcement officers would and could not secretly monitor every move that an individual takes, but "allowing [the] government access to cell-site records contravenes that expectation."⁴⁰ Because modern location data can reveal an extremely detailed record of a person's historic movements, its usage implicates Fourth Amendment concerns about government power and expectations of privacy.⁴¹

The Fifth Circuit applied those principles to geofence warrants in *United States v. Smith*. There, law enforcement officers knew the location of where the crime occurred, but they did not know the identities of the individuals involved.⁴² Traditionally, a warrant is directed towards a known suspect. However, geofence warrants are used to identify a geographic area and time period, allowing technology companies to use their location databases to determine which users were present at that location during an identified timeframe.⁴³ In *Smith*, officers requested location data from Google.⁴⁴ Google's process required searching its "Location History" database, which tracks user's locations across every app and every device associated with the account once enabled.⁴⁵ These searches provide a wide range of information, including anonymized device data, specific locations, and identifying information regarding users.⁴⁶ The Fifth Circuit held that the use of the Defendants' location history constituted a search subject to Fourth Amendment protections, because the users had a reasonable expectation of privacy regarding that information and, under *Carpenter*, the third-party use did not eliminate that privacy interest.⁴⁷ The Court further concluded that the geofence warrants, as used in this case, were unconstitutional.⁴⁸ The geofence searches used by law enforcement were overly broad, as it was unknown whose information would be used in the search.⁴⁹ The Fifth Circuit reasoned that this procedure resembled the general searches that the Fourth Amendment was designed to prevent.⁵⁰ Although the

37 *Carpenter v. United States*, 585 U.S. 296, 309-319 (2018).

38 *Ibid.* at 309.

39 *Ibid.*, at 310.

40 *Ibid.*, at 310-311.

41 *Ibid.*, at 312.

42 *United States v. Smith*, 110 F.4th 817, 820-821 (5th Cir. 2024).

43 *Ibid.*, at 822.

44 *Ibid.*, at 821.

45 *Ibid.*, at 822-823.

46 *Ibid.*, at 825.

47 *Ibid.*, at 831-833, 836.

48 *Ibid.*, at 836-838.

49 *Ibid.*

50 *Ibid.*, at 838.

Court ultimately affirmed the denial of suppression, because the officers acted in good faith, it held that the geofence warrants themselves violated the Fourth Amendment.⁵¹

The Supreme Court subsequently addressed whether geofence warrants constitute a Fourth Amendment search in *Chatrie v. United States*. There, police officers investigating a robbery used a geofence warrant directed to Google, like the warrant used in *Smith*.⁵² The process began with anonymized location information concerning devices near the crime scene and resulted in the disclosure of identifying information for selected users during the relevant timeframe.⁵³ The Supreme Court held that the police conducted a Fourth Amendment search when they obtained Chatrie's Google location history, because there was a reasonable expectation of privacy regarding cell-phone location information.⁵⁴ The Court rejected both the argument that the search was narrow because only two hours of information were involved and the argument that Google's possession of the data eliminated the defendant's privacy interest.⁵⁵ Because the Fourth Circuit did not believe that a search had occurred, it did not address whether the search was validly executed.⁵⁶ The Supreme Court left that determination to the Fourth Circuit Court of Appeals.⁵⁷

Taken together, *Carpenter*, *Smith*, and *Chatrie* demonstrate how data-driven investigative technologies can infringe upon traditional Fourth Amendment protections. *Carpenter* established that detailed location information does not lose its constitutional protection merely because it is stored by a third party. *Smith* confronted the problem of using technology to generate a pool of people who will become suspects. *Chatrie* then confirmed that governmental use of geofence warrants constitute a Fourth Amendment Search. These cases do not establish that predictive-policing scores or recommendations independently constitute a search. However, predictive systems, like geofence searches, can operate before officers have individualized suspicion through conventional investigation. An algorithm may designate a person as high risk, identify a location for heightened enforcement, or generate a lead based on dataset patterns. This can contribute to the determination of potential suspects, rather than supplying individualized evidence about a person already suspected of criminal activity. This implicates a meaningful Fourth Amendment review of predictive policing. If officers cannot explain what information led to or was used in their investigation, then defendants cannot adequately defend themselves. Further, Courts must be able to evaluate whether those actions rested upon constitutionally sufficient, individualized grounds.

51 Ibid., at 839-840.

52 *Chatrie v. United States*, 225 L. Ed. 2d 793, 801 (2026).

53 Ibid., at 804.

54 Ibid., at 801.

55 Ibid., at 807-808.

56 Ibid., at 819.

57 Ibid.

Without that information, judicial review risks examining only officer's final decisions, while ignoring the process that contributed to it.

II Procedural Due Process and the Right to Meaningfully Challenge

Procedural due process requires that individuals have notice and a meaningful opportunity to be heard before the government deprives them of protected liberty interests.⁵⁸ The Supreme Court has repeatedly emphasized that due process is not satisfied by formal procedures alone; it requires a meaningful reduction in the risk of erroneous deprivation.⁵⁹ In *Mathews v. Eldridge*, the Court created a balancing framework that considers the private interest affected, the risk of error, and the government's interest, including the burdens on enforcement.⁶⁰

Predictive policing exacerbates the very risk of error *Mathews* seeks to reduce. Algorithmic systems rely on probabilistic inferences created from data, but the logic behind the decisions and the error rates are undisclosed.⁶¹ Individuals are subsequently subjected to increased surveillance, stops, or investigative scrutiny based on algorithmic assessments they cannot see or understand. Their opportunity to challenge state action becomes nonexistent. Notice becomes useless if the basis for government decision-making is hidden.⁶² As the Court recognized in *Goldberg v. Kelly*, due process requires disclosure of the grounds for governmental decisions so affected individuals can meaningfully respond.⁶³

The individualized character of the suspicion that predictive policing tools generate further supports the need for transparency. In *Londoner v. City & County of Denver*, the Court held that when government action is directed at specific individuals in an adjudicative manner, due process requires notice and an opportunity to be heard, unlike generally applicable legislative action.⁶⁴ Person-based predictive systems that assign risk scores or designate individuals as enforcement priorities operate in precisely this individualized manner. However, they do so without providing the procedural transparency traditionally required when states target individuals.

Courts have taken differing approaches to transparency when algorithmic systems influence criminal proceedings. For example, in *State v. Loomis* the Wisconsin Supreme Court permitted the use of a risk assessment created by an algorithm at a sentencing hearing, yet conceded that the defendant was unable to examine the algorithm's methodology or data inputs.⁶⁵ Although the court expressed concern about nontransparency and bias, it relied on cautionary instructions rather than disclosure requirements to address due process risks.⁶⁶

58 U.S. Const. amend. XIV, § 1.

59 *Mathews v. Eldridge*, 424 U.S. 319, 344 (1976).

60 *Ibid.*, at 335.

61 Garrett and Rudin, *supra* note 15, at 586-589.

62 Bloch-Wehba, *supra* note 30, at 939.

63 *Goldberg v. Kelly*, 397 U.S. 254, 267-268 (1970).

64 *Londoner v. City & Cnty. of Denver*, 210 U.S. 373, 385-386 (1908).

65 *State v. Loomis*, 371 Wis. 2d 235, 881 N.W.2d 749, 761-762 (2016).

66 *Ibid.*, at 767-769.

Loomis illustrates that due process concerns are recognized, yet transparency remains discretionary rather than mandatory.

New Jersey Courts, however, have taken a different approach. For example, in *State v. Arteaga*, police officers used facial-recognition technology while investigating a robbery.⁶⁷ They used these tools to obtain possible facial matches to the robber, eventually leading them to finding the defendant.⁶⁸ After learning that the technology had contributed to this identification, the defendant sought discovery pertaining to the system and search that produced the result.⁶⁹ His request included 13 categories of information, including: the name and manufacturer of the software, error rates, the original probe photograph, any edited versions of the photograph, candidate-match information, characteristics of the database, reports produced by the analyst, and information concerning the analyst's qualifications.⁷⁰ Initially, this request was denied, however the New Jersey Appellate Division reversed the denial.⁷¹ The Court concluded that the requested information had a direct connection to testing the reliability of the facial-recognition technology and bore on the defendant's guilt or innocence.⁷² The defendant had identified specific categories of information and supported his request with expert assistance, thus demonstrating a particularized need for the discovery.⁷³ This decision recognizes that a defendant cannot meaningfully evaluate an algorithmic identification by examining only the final candidate selected by investigators. Testing the identification requires access to information about the process that generated it, including the system used, the image submitted, modifications to the image, the system's reliability, and the alternative candidates suggested.

The New Jersey Supreme Court further refined their approach in *State v. Miles*. Rather than directly applying all 13 categories used in *Arteaga* when facial-recognition technology is used, the Court held that the appropriate scope of discovery depends on the individualized circumstances presented in each case.⁷⁴ The Court still required the State to disclose non-proprietary information identifying the facial-recognition tools, the materials used, and information explaining how law enforcement used those tools during their investigation.⁷⁵ The Court concluded that such information could bear on the reliability of the technology, the defendant's ability to challenge the investigation, impeachment of identifications, and the possibility of third-party guilt.⁷⁶ The State's limited discovery showed that one facial-recognition search ranked Miles eighth among ten possible matches, while another search returned several additional potential matches, which highlights the importance of access to the

67 *State v. Arteaga*, 476 N.J. Super. 36, 42-43 (Super. Ct. App. Div. 2023).

68 *Ibid.*

69 *Ibid.*, at 43.

70 *Ibid.*, at 43-44.

71 *Ibid.*, at 50, 57.

72 *Ibid.*, at 57.

73 *Ibid.*

74 *State v. Miles*, 264 N.J. 209, 229-230 (2026).

75 *Ibid.*, at 232-234.

76 *Ibid.*

results.⁷⁷ Discovery results gave the defense access to information pertaining to how and why police focused on Miles, whether investigators gave adequate consideration to other possible matches, and how the facial recognition influenced the investigation. The Court in *Miles* still recognized limits to disclosure, emphasizing that *Arteaga* did not render propriety information automatically discoverable in all future cases.⁷⁸ This led the Court to reverse the portion of discovery compelling the propriety information, including the source code.⁷⁹ That distinction shows that transparency requires enough information to understand how reliance on an algorithm affected investigations and to test the reliability of the resulting decision. Transparency does not require unrestricted access to lines of source code within the program. Therefore, courts can protect legitimate proprietary interests while still requiring disclosure necessary for meaningful contestability.

III Equal Protection and Proving Discriminatory Enforcement

The equal protection doctrine similarly depends on access to information about how laws and policies are administered. The Supreme Court has held that facially neutral practices violate equal protection when they are applied in a discriminatory manner.⁸⁰ In *Yick Wo v. Hopkins*, the Court invalidated a neutral licensing law that was enforced discriminatorily against Chinese laundry owners, emphasizing that discriminatory administration is constitutionally intolerable.⁸¹

The Court's decision in *Village of Arlington Heights v. Metropolitan Housing Development Corp.* further emphasizes this point. In *Village of Arlington Heights*, the Court explained that discriminatory intent may be inferred from circumstantial evidence, including historical background, procedural departures, and disparate impact.⁸² Algorithmic secrecy removes access to this exact type of evidence. This problem is worsened by the Court's decision in *McCleskey v. Kemp*, in which it rejected an equal protection claim based on statistical disparities due to a lack of proof regarding intent.⁸³ When predictive systems are unclear, establishing intent becomes nearly impossible, making equal protection claims unprovable in practice.

Predictive policing systems claim to be neutral and data-driven, but they rely on data that is often biased, from enforcement patterns to surveillance.⁸⁴ When algorithms rely on this data, the risk of increasing discriminatory outcomes is clear. Transparency is essential to exposing these unjust outcomes. Without access

77 Ibid., at 220.

78 Ibid at 232-233.

79 Ibid.

80 *Yick Wo v. Hopkins*, 118 U.S. 356, 368-369 (1886).

81 Ibid.

82 *Vill. of Arlington Heights v. Metro. Hous. Dev. Corp.*, 429 U.S. 252, 264 (1977).

83 *McCleskey v. Kemp*, 481 U.S. 279, 292-295 (1987).

84 Richardson, *supra* note 8, at 47.

to information about the data, design, and use, plaintiffs cannot meet the requirements for their equal protection claim.

IV *Evidentiary Reliability*

Predictive policing also undermines principles governing the reliability and contestability of evidence. Algorithmic outputs are typically used as internal tools for policing, but they have increasingly been used as evidence that shapes investigations, charging decisions, and sentencing outcomes.⁸⁵ The Constitution has long required that expert evidence be subjected to scrutiny for reliability and validity. In *Daubert v. Merrell Dow Pharmaceuticals*, the Court held that expert evidence must be based on reliable methods and subject to replication by others in the field.⁸⁶ AI outputs should be subjected to this same questioning when used as evidence. Because the workings of algorithmic systems are kept secret, courts and experts cannot evaluate the reliability of their outputs. For example, error rates, validation studies, and methodology are crucial for assessing evidence that is hidden when AI systems are used. This leaves defendants helpless against the AI's conclusions. This mirrors the concerns raised in *Melendez-Diaz v. Massachusetts*, where the Court emphasized that forensic evidence must be subjected to confrontation rather than be accepted as truth, because .⁸⁷ There, the prosecution introduced sworn laboratory certificates stating that the substance seized from the defendant was cocaine, but it did not call the analyst who prepared those certificates to testify.⁸⁸ The Supreme Court held that the certificates were testimonial statements, because they were created for use as evidence at trial and performed the same function as live testimony.⁸⁹ The analysts were therefore “witnesses” for Sixth Amendment purposes and the defendant was entitled to confront them through cross-examination.⁹⁰ The Court rejected the argument that confrontation was unnecessary because the evidence resulted from scientific testing.⁹¹ Instead, the Court emphasized that the Confrontation Clause is a procedural guarantee that reliability is tested through the adversarial process, rather than presumed from the scientific nature of the evidence.⁹²

Bullcoming v. New Mexico further demonstrates why reliability and confrontation remain concerning when technology contributes to forensic evidence. In *Bullcoming*, the prosecution introduced a laboratory report certifying the defendant's blood-alcohol concentration, but they did not call the analyst who performed the test and signed the certification to testify.⁹³ Instead, the State called another analyst to testify who was familiar with the procedures and equipment used, but

85 Garrett and Rudin, *supra* note 15, at 569.

86 *Daubert v. Merrell Dow Pharms., Inc.*, 509 U.S. 579, 593-595 (1993).

87 *Melendez-Diaz v. Massachusetts*, 557 U.S. 305, 310-312 (2009).

88 *Ibid.*, at 308-309.

89 *Ibid.*, at 310-311.

90 *Ibid.*

91 *Ibid.*, at 317-318.

92 *Ibid.*, at 318-320.

93 *Bullcoming v. New Mexico*, 564 U.S. 647, 650-651 (2011).

did not perform or observe the test.⁹⁴ The Supreme Court held that this testimony did not satisfy the Confrontation Clause, because they could not be cross-examined about the original analysts' procedures, observations, or conclusions.⁹⁵ The analyst represented that the sample had been received intact, identifying information corresponding to the defendant, the prescribed testing protocol had been followed, and that no conditions affected the integrity of the sample or validity of the results.⁹⁶ Those statements include human conduct that could be challenged through cross-examination.

Algorithmic evidence raises a similar problem. An AI system can automatically generate a score, report, candidate list, prediction, or recommendation, but human decisions are intermingled within that process. For example, humans decide which information is supplied to the system, whether an image is altered before submission, what parameters govern the search, how outputs are interpreted, and what conclusions are drawn. *Bullcoming* demonstrates why transparency is necessary before a court can determine whether automated evidence contains testimonial human assertions. Without disclosure, a defendant will be unable to identify whether an algorithmic result reflects automated processing or human judgments capable of triggering confrontation rights.

Transparency also intersects with the government's disclosure obligations under *Brady v. Maryland*, where the Court said withholding of evidence that is material to the determination of guilt of a criminal defendant violates the defendant's constitutional right to due process.⁹⁷ Not disclosing algorithmic systems that generate information that is exculpatory, mitigating, or valuable for impeachment undermines *Brady's* command of fairness. The right to know how predictive tools were used or how they influenced decisions protects defendants' access to evidence that may be necessary for a favorable verdict.

V Transparency as a Constitutional Precondition

Together, the above doctrines promote the idea that state power must be exercised through knowable and contestable processes. The Supreme Court's disapproval of arbitrary enforcement in *Papachristou v. City of Jacksonville* rested on the fear that unchecked discretion would invite abuse of the system and erode the rule of law.⁹⁸ When using unclear algorithms rather than clear choice and decision-making, predictive policing replicates this fear in a digital form.

Transparency is not an external policy imposed on constitutional law, but rather a precondition embedded within due process, equal protection, and evidentiary doctrines themselves. Constitutional safeguards become formalities rather than protections when predictive technologies operate as a black box. Acknowledging a public right to know when AI is used in the justice system is an

94 Ibid.

95 Ibid., at 661-662.

96 Ibid., at 653.

97 *Brady v. Maryland*, 373 U.S. 83, 87 (1963).

98 *Papachristou v. City of Jacksonville*, 405 U.S. 156, 168-171 (1972).

essential change that maintains fundamental rights in the digital era without creating a new constitutional protection.

E Reforming Policing Through a Right to Know

Transparency alone does not satisfy the public's right to know regarding predictive policing. The right to know requires frameworks that protect constitutional principles without endangering the legitimate objectives of law enforcement. Statutory reforms and disclosure rules promote an effective right to know. When set up correctly, these systems can safeguard public safety and advance accountability for algorithms.

I Required Disclosure

A meaningful right to know does not require full public access to source code or proprietary algorithms. What constitutional accountability requires is access to information sufficient to understand how predictive systems influence government decision-making. First, accountability requires disclosure of the existence and purpose of any predictive policing system. Second, the categories of data used as inputs should be public. Third, sharing the validation studies, error rates, and known limitations is required. Finally, policies governing how officers may rely on algorithmic outputs should be clear. This would include mandatory training programs for officers at all levels to ensure that data is used correctly and that guidelines are followed.

Advocates claim transparency would increase democratic legitimacy, prevent misconduct, remedy injustice, create better policies, and build community trust.⁹⁹ This information would enable courts to assess the reliability of the programs and their predictions. This approach is similar to existing regulations in other areas of law. For example, expert testimony is not admissible because jurors understand the science in full, but because the methodology and reliability are subjected to adversarial questioning and peer review.¹⁰⁰ That exact admissibility requirement should apply to predictive policing tools because they influence decisions that directly impinge on freedom. Recent case law demonstrates that this calibrated form of disclosure is workable. As explained above, in *State v. Miles*, the New Jersey Supreme Court required disclosure of information identifying the non-proprietary information for facial-recognition tools, the materials used during the investigation, and information explaining how law enforcement used those tools.¹⁰¹ The Court treated this information as relevant, because it could allow the defendant to evaluate the technology's reliability, challenge the investigation, impeach identifications, and develop a potential third-party's guilt.¹⁰² The court in *Miles*

99 Bakke, 'Predictive Policing: The Argument for Public Transparency', 74 *Ann. Surv. Am. L.* (2018) 131, at 144.

100 *Daubert*, supra note 47, at 593-595.

101 *Miles*, supra note 74, at 232-234.

102 *Ibid.*

drew substantially the same requirements proposed here. Disclosure of how a predictive system was used, the categories of data relied upon, known error rates and limitations, and the policies governing officer use can provide meaningful accountability. If more sensitive proprietary information becomes necessary in a particular case, Courts can use protective orders or require a heightened showing of need. Confidentiality concerns therefore justify tailoring disclosure, rather than eliminating transparency.

Another approach the United States should follow is found in the European Union's Artificial Intelligence (AI) Act, which was enacted in 2024.¹⁰³ This Act is the first comprehensive statutory framework to govern the development and use of AI systems. The EU AI Act created a risk-based framework in which systems more likely to infringe fundamental rights attract more regulation. It categorizes AI systems into four categories: minimal risk, limited risk, high risk, and unacceptable risk.¹⁰⁴ For high-risk systems, the act imposes affirmative obligations including mandatory pre-use risk assessments, continuous monitoring, detailed design logs and training data, and robust human oversight requirements.¹⁰⁵ The Act also emphasizes transparency, as operators of high risk systems must be able to explain outputs and ensure automated recommendations are not substituted for human judgement.¹⁰⁶ AI systems within the entirety of the criminal justice system are presumptively high risk, because they can substantially interfere with liberty, due process, and equality.¹⁰⁷ The Act identifies AI systems used to assess whether a person is likely to become a victim of a crime, polygraph-type tools, systems used to evaluate the reliability of evidence, and systems used to assess the risk that a person will offend or reoffend as high-risk.¹⁰⁸ The EU AI Act goes even further with the unacceptable risk category, where the use is prohibited entirely except under extreme emergency conditions.¹⁰⁹ This most notably includes real-time biometric identification programs being used in public locations.¹¹⁰ The Act also prohibits the use of AI systems used to predict whether a person will commit a criminal offense when it is based solely on profiling or the person's personality traits or characteristics.¹¹¹ This ban does not extend to systems that support human assessment based on objective and verifiable facts connected to criminal activity.¹¹² The Act further prohibits AI systems that create or expand facial-recognition databases using images from the internet or closed-circuit television footage.¹¹³ These bans demonstrate that the European Union has found that uses of AI in law enforcement present risks so severe that transparency and oversight are not

103 Regulation (EU) 2024/1689 of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), OJ 2024 L 202.

104 Ibid.

105 Ibid., at 19.

106 Ibid.

107 Ibid., at 9.

108 Ibid., at 17.

109 Ibid.

110 Ibid.

111 Ibid., at 51.

112 Ibid.

113 Ibid.

enough to protect individuals. The distinction between prohibited and high-risk uses show that the Act does not treat all tools the same, while acknowledging the risks to justice when they are used.

The United States current approach is more fragmented and focuses on reacting to harm, whereas the EU AI Act attempts to prevent harm from occurring. The United States currently does not have a comprehensive federal statute. AI is governed by constitutional doctrine, evidence rules, and state/local regulations. As shown in *State v. Loomis*, courts acknowledge algorithmic risks but permit their continued use.¹¹⁴ This reliance on courts rather than clear standards leaves defendants depending on the goodwill and discretion of judges and vendors.

While the EU AI Act may not be a perfect template for regulation in the United States, it provides a clear example that with innovation comes regulation. This is clear by the Act's presumptive high-risk approach to AI systems in criminal justice. It furthers the point that the United States needs to regulate these tools by enforcing meaningful limitations and requiring transparency from both vendors and those who use the systems.

Overall, transparency is the tool used to regulate governmental power and overreach. Intentional openness encourages better police behavior when recording data, eases administrators' analysis of the information, and reduces error systemically.¹¹⁵ It will also create trust in society that the technology is safe, especially when it is federally regulated.

II Statutory and Regulatory Pathways

Legislatures should base enforcement of a right to know on current public records statutes, because those regulations settled the balance between government transparency, administrative feasibility, and legitimate confidentiality concerns. The current law enforcement exemptions that allow withholding information on predictive techniques should be limited by legislators. Transparency law would be in line with contemporary tools if it were required that algorithmic systems used for policy or deployment decisions affecting individuals were public information.¹¹⁶

Existing municipal surveillance ordinances demonstrate that this type of transparency is possible. For example, in Massachusetts, Boston's Ordinance on Surveillance Oversight and Information Sharing expressly includes predictive-policing software within its definition of surveillance technology.¹¹⁷ Before City departments may use surveillance technology, they must obtain City Council approval and are later required to submit a surveillance technology impact report.¹¹⁸ That report must describe the technology and how it works, explain its intended purpose, identify the surveillance data that is gathered, the locations the technology will be deployed, and address its potential effects on privacy, civil rights, and civil liberties.¹¹⁹ The ordinance

¹¹⁴ *Loomis*, supra note 39, at 752-753.

¹¹⁵ Ferguson, supra note 2, at 1167.

¹¹⁶ Citron and Pasquale, supra note 17, at 26.

¹¹⁷ Boston, Mass., Code § 16-63.2(19) (2021).

¹¹⁸ *Ibid.*, at § 16-63.2

¹¹⁹ *Ibid.*, at § 16-63.3(D)(2)(d)-(l).

further requires surveillance technology impact reports and the technology-specific use policies be submitted to the City Council to be made publicly available.¹²⁰

The city of Oakland in California has adopted a similar framework through its Surveillance Technology Ordinance. That ordinance expressly includes predictive-policing software designed to forecast criminal activity or criminality within its definition of surveillance technology.¹²¹ Oakland requires a Surveillance Impact Report that must disclose how the technology was used, whether and how often the data acquired was shared with outside entities, location of the technology, potential civil-rights and civil-liberties consequences, the results of internal audits, data breaches, and crime statistics to determine the effect of the tools.¹²² Oakland separately requires a publicly released and legally enforceable Surveillance Use Policy addressing the authorized use, purpose, data collection, public access, training, auditing, and oversight.¹²³

These ordinances require government entities to disclose what technology is being used, what it does at a functional level, what information it processes, what governs its use, and what risks occur while using the technology. They also require public oversight, so individuals are aware of what technology is being used in their communities. The existence of these frameworks shows that advanced approval, public impact reporting, use policies, and oversight are mechanisms capable of helping ensure protection from invasions of privacy and civil liberties.

Beyond public records law, procurement regulations also provide a robust, straightforward mechanism to increase transparency. Government agencies already limit contracts with requirements for compliance with labor standards, data security, and auditing. Similar terms in AI contracts can require vendors to provide documentation about their design, training data, and performance metrics as a requirement for governmental use. Law enforcement agencies can prevent disputes over proprietary claims and public accountability by incorporating transparency requirements into the contracting stage. Mandating impact evaluations and audits can also improve transparency throughout the procurement process. Audits should scan the entire system, from collection to analysis, and how it is maintained. Professional data analysts should complete these audits and include a notice mechanism to ensure results cannot be hidden or skewed.¹²⁴ Predictive systems should be independently evaluated to detect bias, assess accuracy, and decide whether use is necessary.¹²⁵ For the system to become fully transparent, metrics must be publicly released. This will ensure the programs are effective at meeting their goals, like setting a target for crime rate reduction or comparing AI predictions to human predictions.¹²⁶ These requirements are not

120 Ibid., at § 16.63.3(B)(3)(e).

121 Oakland, Cal., Mun. Code § 9.64.010(11) (2018).

122 Ibid., at § 9.64.010(1)(A)-(K).

123 Ibid., at § 9.64.010(16)(A)-(K).

124 Ferguson, *supra* note 2, at 1167.

125 National Academies of Sciences, Engineering and Medicine, *Law Enforcement Use of Predictive Policing Approaches: Proceedings of a Workshop* (2025), ch. 4, at 40.

126 Ferguson, *supra* note 2, at 1168.

substitutes for complete disclosure, they provide additional layers of protection to reduce the harm caused by unchecked algorithms.

III The Role of Courts

Without a constitutional amendment to expressly guarantee a right to know, courts become essential for the practical implementation of transparency. Rulings on evidence admissibility, discovery obligations, and constitutional review serve as the backdrop for a transparency requirement. Based on this precedent, courts are able to translate the abstract principles into enforceable requirements. Courts can condition the use of predictive algorithms with the disclosure of sufficient information, allowing meaningful challenges when used as sources for probable cause, sentencing decisions, and other outcomes that impact liberty.¹²⁷

This approach aligns with the Supreme Court's insistence that reliability and contestability are prerequisites to legitimate state evidence.¹²⁸ The Supreme Court has emphasized that evidence cannot be asserted on accuracy alone, but it must also be reliable and contestable.¹²⁹ Evidence has more weight on a verdict when it can be examined, challenged, and rebutted throughout a trial. Just as courts exclude forensic evidence that cannot be examined or tested, they may exclude algorithmic outputs when their methodologies, error rates, and foundations are secret.¹³⁰ Conditioning admissibility on transparency incentivizes agencies to use systems compatible with constitutional protections and share the results with the public for continued use. Vendors would be encouraged to create tools with automatic disclosure on a regular basis, rather than relying on claims of secrecy to protect their programs from scrutiny. Eventually, this will promote structural accountability throughout the entirety of the justice system. Overall, courts can enforce a practical right to know that preserves constitutional protections while allowing the use of modern technology.

IV Addressing Counterarguments

Those against transparency frequently argue that disclosure will compromise public safety by enabling individuals to evade law enforcement or manipulate predictive systems. That objection has a legitimate foundation and should not be reduced to a preference for secrecy. For example, the FOIA Exemption 7(E) reflects the concern that disclosure of certain law enforcement techniques and procedures could enable individuals to circumvent investigative methods.¹³¹ A meaningful right to know must distinguish between information necessary for accountability and disclosure of operational information that would create a concrete risk of evasion. Error rates, validation studies, known limitations, general categories of input data, and policies governing reliance would allow courts and the public to evaluate

¹²⁷ Ferguson, *supra* note 4, at 300-305.

¹²⁸ Daubert, *supra* note 47, at 593-597.

¹²⁹ *Ibid.*

¹³⁰ Melendez-Diaz, *supra* note 48, at 310-312.

¹³¹ Freedom of Information Act, *supra* note 22.

whether a predictive system is reliable and constrained to a justifiable use. Those forms of information differ from real-time surveillance locations, specific investigative targets, or technical instructions capable of helping individuals circumvent investigative techniques. *Miles* provides a useful model for drawing this distinction. The Court required disclosure identifying the facial-recognition tools used and explaining how law enforcement officers employed them during the investigation, while protecting source code.¹³² That decision demonstrates how courts can distinguish information necessary to evaluate government action from information requiring a heightened standard of confidentiality. The same principle should govern broader transparency requirements. Confidentiality should be tied to an identifiable operational risk. A system can remain secure while being subjected to external evaluation of reliability, error rates, biases, and permissible uses. Complete secrecy prevents defendants, courts, legislators, and affected communities from discovering errors or unconstitutional practices.

Similarly, secrecy has not been a necessary condition for effective policing, primarily because previous systems have relied on widely known practices and historical data. Leaders and officers can only be accountable if people clearly understand what the predictive technology is doing. After the Shreveport Police Department began using predictive tools, one officer admitted he began stopping and questioning juveniles more and admitted to walking around apartment complexes, discussing criminal activities in the area with residents.¹³³ He also visited people he knew to be involved in crimes, like parolees, probationers, and truants, to learn about criminal activities in the neighborhood.¹³⁴ The result of individualization and unethical policing outweighs the fear of manipulating AI systems to evade law enforcement.

Some vendors claim that transparency will hinder innovation. These suppliers, however, are unable to defend complete secrecy. Private businesses routinely work in regulated circumstances where their performance evaluations, safety data, and test results must be disclosed.¹³⁵ This does not mean providing access to codes or confidential information. Rather, transparency would require access to the decision-making process, data collected, and methodology used by the programs. Transparency is a requirement of legitimacy, not an excessive burden, because vendors are financially profiting from the instruments that shape coercive state power.

Finally, concerns have been raised about the administrative burden and review process. This burden cannot outweigh constitutional guarantees. As the Supreme Court has repeatedly emphasized, efficiency alone cannot justify procedures that increase the risk of erroneous or discriminatory deprivation of liberty.¹³⁶ A way to mitigate the burden is to provide training for agencies on the use of predictive tools to ensure compliance with citizens' rights to freedom. This would also enable

132 *Miles*, supra note 74, at 232-234.

133 Ferguson, supra note 2, at 1173.

134 *Ibid.*

135 *Ruckelshaus v. Monsanto Co.*, 467 U.S. 986, 1006-1007 (1984) (Recognizing that submitting certain to regulators is a condition of market access, even when the information is commercially sensitive.).

136 *Mathews*, supra note 33, at 335.

immediate responses to the systems in practice, enabling change quickly without the need for administrative review after issues inevitably arise. The costs of transparency are outweighed by the constitutional and democratic interests at stake.

Discretionary disclosure must be replaced with required transparency to protect liberty from AI intrusion effectively. Regulations about predictive policing technologies must require transparency, limit secrecy, and promote contestability to protect citizens' rights effectively. By incorporating transparency into regulations, contracts, and judicial doctrine, algorithmic tools can promote public safety without evading accountability.

F Conclusion

Predictive policing has changed how states use their power in many ways. For example, law enforcement agencies now use computer algorithms to surveil, determine stops, and rank enforcement priorities. Law enforcement has also transformed from personal judgments into AI-generated probabilities that compromise individual liberty. Most of these changes have been implemented without public knowledge. Transparency for everyday citizens is essential to ensure they are protected under constitutional doctrine, because AI use is often unseen and unheard. Because of this secrecy, constitutional rights are reduced to formalities and are unable to truly shield citizens from discriminatory police tactics.

Given the potential for these tools to undermine liberty, disclosure of how predictive algorithms operate, what data they rely on, and how their outputs affect decision-making is a minor demand. Algorithmic policing should not be an exception to the law's historical requirement of visibility, and its use should be clear. Affirming the right to know ensures that innovation will remain inferior to accountability and the Constitution. Overall, transparency is not an obstacle that hinders effective policing. It is a protection for the communities that need it most. Without transparency, liberty is being left to an algorithmic black box, closed to both courts and citizens.