ARTICLE

MACHINE LEARNING, AUTOMATED SUSPICION ALGORITHMS, AND THE FOURTH AMENDMENT

MICHAEL L. RICH†

At the conceptual intersection of machine learning and government data collection lie Automated Suspicion Algorithms, or ASAs, which are created by applying machine learning methods to collections of government data with the purpose of identifying individuals likely to be engaged in criminal activity. The novel promise of ASAs is that they can identify data-supported correlations between innocent conduct and criminal activity and help police prevent crime. ASAs present a novel doctrinal challenge as well, as they intrude on a step of the Fourth Amendment's individualized suspicion analysis, previously the sole province of human actors: the determination of when reasonable suspicion or probable cause can be inferred from established facts. This Article analyzes ASAs under existing Fourth Amendment doctrine for the benefit of courts that will soon be asked to deal with ASAs. In the process, this Article reveals the inadequacies of existing doctrine for handling these new technologies and proposes extrajudicial means for ensuring that ASAs are accurate and effective.

† Jennings Professor of Law, Elon University School of Law. B.A., University of Delaware; J.D., Stanford Law School. I particularly thank Amy Minardo, Bennett Capers, Simon Stern, Orin Kerr, the participants in CrimFest! 2014, and the faculties at Washington University School of Law and Stetson University College of Law for their thoughtful comments. My thanks to Michael Costolo and Brittany Puckett for their invaluable research assistance. And I am especially grateful to Dr. Daniel J. George, Dr. Michael N. Ferrandino, and their colleagues, without whom I would not be here to write this author's note.

INTRODUCTION

One day soon, a machine will identify likely criminal activity and, with the beep of an e-mail delivery, the buzz of an alarm, or the silent creation of a report, tell police where to find it. Already, a computer program analyzes massive quantities of securities trading data and notifies the Securities and Exchange Commission of investors who might be engaged in insider trading.[1] Computer systems connected to networks of video cameras alert police when bags are abandoned on subway platforms,[2] when people on a street corner

1 *See* Mary Jo White, Chair, SEC, Keynote Address at the 41st Annual Securities Regulation Institute (Jan. 27, 2014), http://www.sec.gov/News/Speech/Detail/Speech/1370540677500 [https://perma.cc/M7YV-33PR] (describing the SEC's NEAT program, which can identify and analyze issuer trading activity around times of major corporate events). Similarly, another algorithm compares medical billing data against previously identified suspicious billing patterns to uncover likely instances of fraud. *See* Colin Caffrey, *Can a Computer Read a Doctor's Mind? Whether Using Data Mining as Proof in Healthcare Fraud Cases Is Consistent with the Law of Evidence*, 30 N. ILL. U. L. REV. 509, 510-11 (2010).

2 FIRETIDE, CITY OF CHICAGO: FIRETIDE WIRELESS MESH KEY TO CITY-WIDE VIDEO SECURITY DEVELOPMENT (2007), http://www.firetide.com/files/9014/0122/6078/City_of_Chicago.pdf [https://perma.cc/62J6-QTRC].

interact multiple times in a short period,[3] or when a single individual visits multiple cars in a parking structure.[4] The federal government has field tested a device that screens individuals and predicts whether, based on physiological data, the individual intends to commit a terrorist act.[5] Researchers at Carnegie Mellon, funded by the Defense Advanced Research Projects Agency, are developing computer systems to index and analyze the text and images in online advertisements for sex services to identify likely sex traffickers and their victims.[6] While these current technologies generally follow a comprehensible logic—looking for facts that we understand to correlate with criminal conduct—technologies of the near future will analyze more data than a human being could and unearth connections that evade obvious logic.[7] In other words, soon a computer may spit out a person's name, address, and social security number along with the probability that the person is engaged in a certain criminal activity, with no further explanation.[8]

These emergent technologies arise from the intersection of two trends: the collection of massive troves of individualized data about people in the United States and the explosive growth of a field of computer science known as machine learning.[9] With respect to the former, these data come from a

3 Russell Nichols, *Smart Cameras Aim to Stop Crimes Before They Occur*, GOV'T TECH. (Oct. 26, 2010), http://www.govtech.com/featured/Smart-Cameras-Aim-to-Stop-Crimes-Before-They-Occur.html [https://perma.cc/DZ62-JNHB]; Digital Justice, *Digisensory Technologies Avista Smart Sensors*, YOUTUBE (Sept. 14, 2012), https://www.youtube.com/watch?v=JamGobiS5wg [https://perma.cc/6EQF-GMZL].

4 Diane Cardwell, *At Newark Airport, the Lights Are On, and They're Watching You*, N.Y. TIMES (Feb. 17, 2014), http://www.nytimes.com/2014/02/18/business/at-newark-airport-the-lights-are-on-and-theyre-watching-you.html [https://perma.cc/L2NQ-XPKW].

5 Sharon Weinberger, *Terrorist "Pre-Crime" Detector Field Tested in United States*, NATURE (May 27, 2011), http://www.nature.com/news/2011/110527/full/news.2011.323.html [https://perma.cc/QUR5-EJMY].

6 Byron Spice, *Carnegie Mellon Developing Online Tools to Detect and Identify Sex Traffickers*, CARNEGIE MELLON U. NEWS (Jan. 13, 2015), http://www.cmu.edu/news/stories/archives/2015/january/detecting-sex-traffickers.html [https://perma.cc/5W4D-EPF2].

7 *See* Margaret Hu, *Small Data Surveillance v. Big Data Cybersurveillance*, 42 PEPP. L. REV. 773, 803 (2015) ("As we transition from a small data world to a big data world, it appears that the government may be at the earliest stages of attempting to merge small data evidence and big data evidence for prosecutorial purposes.").

8 *See* Tal Z. Zarsky, *Transparent Predictions*, 2013 U. ILL. L. REV. 1503, 1519-20 (explaining how automated predictions can be generated in processes "which [are] not explainable in human language," such that "[i]t would be difficult for the government to provide a detailed response when asked why an individual was singled out to receive differentiated treatment by an automated recommendation system"). As a side note, this impending capability has captured the imagination of popular culture. Three current television series feature analogous technologies. *See Minority Report* (FOX); *Person of Interest* (CBS); *The Player* (NBC).

9 For a discussion of the scope of government data collection and sharing, see Danielle Keats Citron & Frank Pasquale, *Network Accountability for the Domestic Intelligence Apparatus*, 62 HASTINGS L.J. 1441, 1448-55 (2011) (discussing the role of fusion centers in an effective "information sharing environment" between government agencies); Deven R. Desai, *Constitutional Limits on Surveillance:*

nearly unlimited variety of public and private sources, including video cameras, crime scene gunshot detectors, license plate readers, automatic tollbooth payment systems, and social media websites.[10] Government bodies from the municipal to the federal level are all involved in this "data vacuuming."[11] Moreover, private companies are increasingly making personal data available to governments including to law enforcement agencies.[12] With a mixture of resignation and pessimism, this Article takes the government's past and future collection of enormous quantities of personal data as a given and instead examines the government's use of those data.[13]

Meanwhile, researchers have made colossal strides in recent years in machine learning, "the systematic study of algorithms and systems that improve their knowledge or performance with experience."[14] Machine learning is particularly useful for revealing otherwise unrecognizable patterns in complex

Associational Freedom in the Age of Data Hoarding, 90 NOTRE DAME L. REV. 579, 626-27 (2014) (detailing expansive data collection by the FBI and NSA).

10 Citron & Pasquale, *supra* note 9, at 1451 (discussing the range of information sources that feed into government fusion centers).

11 Nancy L. Rosenblum, *Governing Beyond Imagination: The "World Historical" Sources of Democratic Dysfunction*, 94 B.U. L. REV. 649, 658 (2014); *see also, e.g.*, Devlin Barrett, *U.S. Spies on Millions of Drivers*, WALL ST. J. (Jan. 26, 2015), http://www.wsj.com/articles/u-s-spies-on-millions-of-cars-1422314779 [https://perma.cc/WR2J-6KYU] (discussing DEA efforts to track license plates to combat drug trafficking efforts); Somini Sengupta, *Privacy Fears Grow as Cities Increase Surveillance*, N.Y. TIMES (Oct. 13, 2013), http://nyti.ms/18lwttl [https://perma.cc/UH92-HKZX] (discussing centralized data collection systems in Oakland and New York City); Hilton Collins, *Video Camera Networks Link Real-Time Partners in Crime-Solving*, GOV'T TECH. (Feb. 1, 2012), http://www.govtech.com/public-safety/Video-Camera-Networks-Link-Real-Time-Partners-in-Crime-Solving.html [https://perma.cc/6LG8-3TYU] (discussing similar systems in Chicago, Atlanta, and Memphis).

12 *See, e.g.*, Stephen Russo, *Creating a Safer Planet with Smarter Analytics Solutions*, IBM BIG DATA & ANALYTICS HUB (June 22, 2015), http://www.ibmbigdatahub.com/blog/creating-safer-planet-smarter-analytics-solutions [https://perma.cc/W73W-YX59] ("The introduction of COPLINK on the Cloud is a major milestone for both IBM and the law enforcement community. This gives law enforcement agencies of any size access to the world's largest network of law enforcement data, comprising more than 1 billion documents.").

13 This assumption should not suggest approval. Fortunately, many have already discussed the panoply of concerns, often revolving around individual privacy and the Fourth Amendment, raised by such data gathering. *See, e.g.*, EXEC. OFFICE OF THE PRESIDENT, BIG DATA: SEIZING OPPORTUNITIES, PRESERVING VALUES 53-55 (2014), http://www.whitehouse.gov/sites/default/files/docs/big_data_privacy_report_may_1_2014.pdf (discussing privacy concerns arising from collection of "big data" about private citizens). *See generally* Paul Ohm, *The Fourth Amendment in a World Without Privacy*, 81 MISS. L.J. 1309 (2012) (discussing how the Fourth Amendment can continue to protect citizens in light of the collection of "big data" by private companies and the third-party doctrine under the Fourth Amendment); Jeffrey M. Skopek, *Reasonable Expectations of Anonymity*, 101 VA. L. REV. 691 (arguing that, in light of widespread data collection, the Fourth Amendment should be interpreted to protect reasonable expectations of anonymity in addition to reasonable expectations of privacy).

14 PETER FLACH, MACHINE LEARNING: THE ART AND SCIENCE OF ALGORITHMS THAT MAKE SENSE OF DATA 3 (2012).

processes underlying observable phenomena.[15] Specifically, machine learning techniques help computer systems learn about an underlying process and its patterns by creating a useful mathematical approximation of how the process works.[16] This approximation can then be applied to new data to predict future occurrences of the same phenomena.[17] For instance, machine learning methods are used to examine patient records and create algorithms that can help doctors diagnose illnesses or provide prognoses.[18]

At least on a conceptual level, machine learning and crime fighting are a perfect match. The interaction of forces that cause people to commit crimes is incomprehensibly complex. Criminologists have sought for decades to use data to understand that interaction and identify the most likely criminal offenders.[19] Statistical models that aim to identify the criminally inclined based on quantifiable personal characteristics have become influential in the contexts of pretrial release, probation, and parole.[20] Similarly, police departments have recently begun to use statistical models to predict where in their jurisdictions certain crimes are likely to occur.[21] Machine learning provides a way to go one step further and use data to identify likely criminals among the general population without the need to disentangle the Gordian knot of causal forces.

This Article addresses technologies that apply machine learning techniques to the "data hoards" available to law enforcement in order to predict individual criminality.[22] Some of these technologies are already in use or are in advanced stages of development.[23] Nascent examples are even more numerous, including: using past offender and crime scene data to create more accurate profiles of unknown offenders,[24] leveraging behavioral data to identify individuals who

---

15 ETHEM ALPAYDIN, INTRODUCTION TO MACHINE LEARNING 2 (3d ed. 2014).

16 *Id.*

17 *Id.*

18 IGOR KONONENKO & MATJAŽ KUKAR, MACHINE LEARNING AND DATA MINING: INTRODUCTION TO PRINCIPLES AND ALGORITHMS 25 (2007).

19 *See, e.g.*, Kevin Miller, *Total Surveillance, Big Data, and Predictive Crime Technology: Privacy's Perfect Storm*, 19 J. TECH. L. & POL'Y 105, 114 (2014) (describing actuarial techniques in place since the 1930s in the field of criminology).

20 Andrew Guthrie Ferguson, *Predictive Policing and Reasonable Suspicion*, 62 EMORY L.J. 259, 285 (2012); *see also* RICHARD BERK, CRIMINAL JUSTICE FORECASTS OF RISK: A MACHINE LEARNING APPROACH 3-4 (2012) (describing Pennsylvania's use of a "computerized risk-assessment model" as part of pre-sentencing and early release decisions); Shima Baradaran & Frank L. McIntyre, *Predicting Violence*, 90 TEX. L. REV. 497, 529-31 (2012) (undertaking statistical analysis of defendants to identify characteristics that are reliable predictors of crime).

21 *See* Ferguson, *supra* note 20, at 265-70 (providing examples).

22 *See* Desai, *supra* note 9, at 583 (describing the FBI's data gathering process).

23 *See supra* notes 1–6 and accompanying text for examples.

24 K. Baumgartner et al., *Constructing Bayesian Networks for Criminal Profiling from Limited Data*, 21 KNOWLEDGE-BASED SYS. 563, 564-66 (2008).

are attempting to conceal their true—and potentially nefarious—intent,[25] and analyzing past corporate financial statements to create algorithms that can determine from the language used in a financial statement whether the company is likely engaged in fraud.[26]

This Article refers to programs like these—programs created through machine learning processes that seek to predict individual criminality—as Automated Suspicion Algorithms, or ASAs. ASAs share three defining characteristics as implied by the name. First, they are based on *algorithms*, which can be broadly defined as sequences of instructions to convert an input into an output.[27] In this case, ASAs convert data about an individual and her behavior into predictions of the likelihood that she is engaged in criminal conduct.[28] Second, ASAs assess individuals based on *suspicion* of criminal activity in that they engage in probabilistic predictions that rely on patterns detected in imperfect information.[29] Third, ASAs *automate* the process of identifying suspicious individuals from data: they comb through data for factors that correlate to criminal activity, assess the weight of each factor and how it relates to other factors, use the results to predict criminality from new data, and continuously improve their performance over time.[30] The automated creation of rules that predict criminality distinguishes ASAs from computer systems that might merely automate the application of a pre-existing police profile of criminality.[31]

---

25 Judee K. Burgoon et al., *Detecting Concealment of Intent in Transportation Screening: A Proof of Concept*, 10 IEEE TRANSACTIONS ON INTELLIGENT TRANSP. SYS. 103 (2009).

26 Sean L. Humpherys et al., *Identification of Fraudulent Financial Statements Using Linguistic Credibility Analysis*, 50 DECISION SUPPORT SYS. 585 (2011). This article is part of an issue dedicated to quantitative methods for detecting financial fraud. *See Quantitative Methods for Detection of Financial Fraud,* 50 DECISION SUPPORT SYS. 557 (2011).

27 ALPAYDIN, *supra* note 15, at 2.

28 *Id.* at 2-3.

29 *See infra* notes 83–87 and accompanying text (discussing how machine learning algorithms may handle imperfect data); *see also* United States v. Knights, 534 U.S. 112, 121 (2001) ("The degree of individualized suspicion required of a search is a determination of when there is a sufficiently high probability that criminal conduct is occurring to make the intrusion on the individual's privacy interest reasonable."); Harry Surden, *Machine Learning and Law*, 89 WASH. L. REV. 87, 98 (2014) (noting that because machine learning algorithms rely on patterns, "they necessarily are under- and over-inclusive relative to the phenomenon they are representing").

30 *See* ALPAYDIN, *supra* note 15, at 5 (providing the example of credit scoring, by which an algorithm examines past data to create rules that generate predictions of future risk when presented with new data). The fact that the processing of data is automated should not obscure the fact, however, that humans are involved in the process of programming and training the ASA. *See infra* Part I.

31 For instance, makers of a lighting and security system at the Newark Liberty International Airport claim that it can alert security if an individual stops at numerous cars in a parking lot, presumably because it is common wisdom in law enforcement that such behavior suggests criminal intent. Cardwell, *supra* note 4. This is an example of the application of a preexisting profile, where

Of course, from fingerprints to field testing kits to DNA matching, law enforcement has always tried to find ways to use the newest technologies.[32] As a result, attorneys, judges, and commentators are quite familiar with the role that technologies play in helping police ascertain the basic facts about a crime: the who, what, when, where, and why. A field test for cocaine, for instance, tells police whether a certain substance is contraband. A DNA match confirms that a suspect was at a crime scene. But determining these historical facts is only the first step in deciding whether individualized suspicion exists sufficient to justify a search or seizure under the Fourth Amendment.[33]

Until now, the second step in determining the existence of individualized suspicion—deciding whether the historical facts give rise to probable cause or reasonable suspicion[34]—has remained the sole province of human actors. The Supreme Court has held that determinations about the existence of probable cause and reasonable suspicion ultimately depend on reason,[35] "common sense,"[36] and police experience.[37] The Court has also made clear that individualized suspicion is ultimately about "probabilities," though in the next breath we learn that probabilities "are not technical."[38] The promise of ASAs is that they can answer the individualized suspicion question by providing data-derived probabilities of whether crime is afoot; the novel problem they present is how those statistical probabilities fit in the "practical, nontechnical conception" of individualized suspicion articulated by the Supreme Court.[39]

human experience and logic, not a machine learning algorithm learning iteratively from data, generate the rule predicting the targeted criminal activity.

32 *See* Erin Murphy, *Databases, Doctrine & Constitutional Criminal Procedure*, 37 FORDHAM URB. L.J. 803, 805-07 (2010) (reciting, briefly, the history of law enforcement databases).

33 *See* Ornelas v. United States, 517 U.S. 690, 696 (1996) ("The first part of the [Fourth Amendment] analysis involves only a determination of historical facts, but the second is a mixed question of law and fact . . . .").

34 *Id.*

35 *See* Terry v. Ohio, 392 U.S. 1, 21 (1968) (holding that Fourth Amendment searches and seizures must be justified by facts and "rational inferences from those facts").

36 *See* Illinois v. Gates, 462 U.S. 213, 244 (1983) ("[W]e think it suffices for the practical, common-sense judgment called for in making a probable-cause determination.").

37 *See* United States v. Cortez, 449 U.S. 411, 418 (1981) ("[A] trained officer draws inferences and makes deductions . . . that might well elude an untrained person.").

38 Brinegar v. United States, 338 U.S. 160, 175 (1949).

39 *Id.* at 176.

ASAs are coming,[40] and courts will soon be asked to consider how their output should factor into the individualized suspicion analysis.[41] The initial goal of this Article is to provide courts with a framework for that analysis.[42] Yet setting out this framework teaches broader lessons about how emergent technologies interact with the Fourth Amendment. First, we learn that ASAs push the limits of the Court's current approach to the Fourth Amendment in areas that have already raised red flags among scholars. One is the ongoing metamorphosis of the collective knowledge doctrine into what some call the "constructive knowledge" doctrine.[43] The former allows knowledge to be imputed

40 *See* Reed E. Hundt, *Making No Secrets About It*, 10 ISJLP 581, 588 (2014) (asserting that "government now routinely asks computers to suggest who has committed crimes"); Zarsky, *supra* note 8, at 1506 (noting that governments "are increasingly curious to figure out what we will do next and take action, rather than wait and investigate what has already happened and suffer the possible consequences"). Though the author has found no reported case that addresses this issue, the Virginia Supreme Court came tantalizingly close in *Commonwealth v. Smith*, 709 S.E.2d 139 (Va. 2011). There, officers conducted a *Terry* frisk of a suspect solely on the basis of an alert they received from a police database known as PISTOL (Police Information System Totally On Line) that stated that the suspect was "probably armed and a narcotics seller/user." *Id.* at 143. The alert issued because an officer had made a record in the system of the suspect's prior arrest for weapons and drug crimes. *Id.* The lower court record did not reveal whether the alert issued because of an automated decision of the PISTOL software or the choice of the officer who created the earlier arrest record. Smith v. Commonwealth, 683 S.E.2d 316, 318 n.1 (Va. Ct. App. 2009), *rev'd* 709 S.E.2d 139 (Va. 2011). Regardless, the Virginia Supreme Court relied upon the constructive knowledge doctrine, discussed *infra* Section III.A, in refusing to suppress evidence obtained as a result of the frisk. *See Smith*, 709 S.E.2d at 144 (imputing knowledge of the defendant's criminal record to the officers "based on the language appearing in the PISTOL alert," which was "critical in determining whether the officers had reasonable suspicion").

41 The Fourth Amendment implications of automated predictions about criminality have largely escaped in-depth analysis by scholars, though some have recognized the importance of the question. For instance, Daniel J. Steinbock raises, but does not answer, the overarching question discussed herein. Daniel J. Steinbock, *Data Matching, Data Mining, and Due Process*, 40 GA. L. REV. 1, 30 (2005). Additionally, Erin Murphy frames the issue as "the challenges that large-scale databasing pose[] to conventional constitutional analysis." Murphy, *supra* note 32, at 804. Her discussion of issues raised by databases is extremely useful in the context of ASAs, but she does not attempt to examine the individualized suspicion analysis in any depth. *Id.* at 826 ("[M]y aim is more to think about the meaning of databases than the meaning of constitutional doctrine."). Finally, Andrew Guthrie Ferguson asks, but again does not resolve, many of the questions addressed herein. *See* Andrew Guthrie Ferguson, *Big Data and Predictive Reasonable Suspicion*, 163 U. PA. L. REV. 327, 383-84 (2015) (discussing how "big data invites provocative questions about whether such predictive tips should factor into the reasonable suspicion calculus").

42 As the discussion herein ultimately concludes courts are likely not the best place for rulemaking with respect to ASAs. *See generally* Orin S. Kerr, *The Fourth Amendment and New Technologies: Constitutional Myths and the Case for Caution*, 102 MICH. L. REV. 801 (2004) (articulating at length arguments against judicial rulemaking with respect to new technologies). Nonetheless, even if rules are drafted in the near future to address police use of ASAs, defendants will still make Fourth Amendment arguments, and courts will need to know how to address them.

43 *See infra* Section III.A; *see also* Simon Stern, *Constructive Knowledge, Probable Cause, and Administrative Decisionmaking*, 82 NOTRE DAME L. REV. 1085, 1089-94 (2007) (contrasting the "collective-knowledge rule" with the "constructive-knowledge rule"); Daniel Poniatowski, Comment,

between officers, so one officer may instruct another to conduct a search or seizure without having to explain why.[44] The latter permits a search based on the aggregated knowledge of law enforcement personnel generally, even if no one officer possessed enough knowledge to make an individualized suspicion assessment.[45]

Another area of concern is the integration of statistical data in the individualized suspicion analysis,[46] which the Supreme Court recently tackled in the context of drug dogs.[47] A third area implicated by ASAs is the Supreme Court's holding that errors in police databases require the exclusion of evidence only in cases of gross negligence or systemic misconduct.[48] Taken together, these issues establish a second, overarching point: ASA accuracy cannot be regulated through the courts alone; rather, extrajudicial action is needed to ensure that ASAs are created, maintained, used, and updated accurately and effectively.

Part I of this Article provides a brief background of machine learning and how it could be applied to create ASAs. Part II sketches out the Fourth Amendment's individualized suspicion analysis, with a particular focus on the two steps articulated by the Supreme Court. Part III tackles the question of whether an ASA's prediction should be sufficient to establish individualized suspicion and concludes that it should not. Part IV discusses how ASAs should be integrated into the totality-of-the-circumstances analysis. Part V addresses how courts should handle ASA errors, and specifically when such errors should lead to exclusion of evidence. The Article concludes by pulling together lessons from the prior discussion and proposing extrajudicial means of ensuring ASA accuracy.

A Constructive Problem: Redemption of Unlawful Arrests Via Fusion Centers, 2014 WIS. L. REV. 831, 834-35 (discussing the "constructive-knowledge" doctrine and its roots in the Supreme Court seeking "to promote law enforcement efficiency").

44 Stern, *supra* note 43, at 1089.

45 *Id.* at 1092-93.

46 *See, e.g.*, Erica Goldberg, *Getting Beyond Intuition in the Probable Cause Inquiry*, 17 LEWIS & CLARK L. REV. 789, 808-10 (2013) (discussing the errors that often arise when courts attempt to incorporate statistical data in the individualized suspicion determination); Richard E. Myers II, *Detector Dogs and Probable Cause*, 14 GEO. MASON L. REV. 1, 4, 13 (2006) (explaining how misunderstandings about the accuracy of a drug dog's alert lead courts to ascribe them more evidentiary value than they deserve).

47 Florida v. Harris, 133 S. Ct. 1050, 1056-57 (2013); *see also* Kit Kinports, *The Dog Days of Fourth Amendment Jurisprudence*, 108 NW. U. L. REV. 64, 64-69 (2013) (critiquing *Harris*).

48 Herring v. United States, 555 U.S. 135, 146 (2009); Arizona v. Evans, 514 U.S. 1, 15-16 (1995).

I. MACHINE LEARNING AND ASAS

"Machine learning" is part of a nest of concepts in the artificial intelligence arena, including "data mining,"[49] "knowledge discovery in databases,"[50] and "big data,"[51] that are often used interchangeably and confusingly in academia, government, and popular media.[52] For the sake of clarity, in this Article "machine learning" refers to the study of algorithms that analyze data in order to help computer systems become more accurate over time when completing a task.[53] This continuous improvement on a given task is the "learning" referenced in "machine learning," and it differs from the more holistic concept referred to when people speak of human learning.[54] In particular, machine learning does not require a computer to engage in higher-order cognitive skills like reasoning or understanding of abstract concepts.[55] Rather, machine learning applies inductive techniques to often-large sets of data to "learn" rules that are appropriate to a task.[56] In other words, the "intelligence" of a machine learning algorithm is oriented to outcomes, not process: a "smart" algorithm reaches consistently accurate results on the chosen task even if the algorithm does not "think" like a person.[57]

---

49 For instance, Daniel Solove calls the process undertaken by algorithms like ASAs "data mining." *See* Daniel J. Solove, *Data Mining and the Security-Liberty Debate*, 75 U. CHI. L. REV. 343, 343 (2008) ("Data mining involves creating profiles by collecting and combining personal data, and analyzing it for particular patterns of behavior deemed to be suspicious."). That term is not used herein, however, because it is not consistently defined in the literature, *see, e.g.*, Jeffrey W. Seifert, *Data Mining and the Search for Security: Challenges for Connecting the Dots and Databases*, 21 GOV'T INFO. Q. 461, 462 (2004) ("[W]hile data mining is widely mentioned in a growing number of bills, laws, reports, and other policy documents, an agreed upon definition or conceptualization of data mining appears to be generally lacking within the policy community."), and it fails to capture the pattern detection that is crucial in the criminal law context.

50 *See* KONONENKO & KUKAR¸ *supra* note 18, at 2-3 (discussing the conceptual interaction between knowledge discovery in databases and data mining).

51 The term "big data" is ubiquitous in legal academic literature. *See, e.g.*, Ferguson, *supra* note 41; Hu, *supra* note 7. Unfortunately, its ubiquity has led to imprecision and confusion about what the term precisely means. *Id.* at 794.

52 This Article does not undertake the Herculean task of resolving the terminological confusion, though others have done so. *See* Liane Colonna, *A Taxonomy and Classification of Data Mining*, 16 SMU SCI. & TECH. L. REV. 309, 313-29 (2013) (defining a wealth of terms relating to machine learning and data mining).

53 *See id.* at 320 ("[M]achine learning is concerned with the development of algorithms and techniques for building computer systems that can automatically improve with experience . . . ."); *see also* FLACH, *supra* note 14, at 3 ("Machine learning is the systematic study of algorithms and systems that improve their knowledge or performance with experience."); Surden, *supra* note 29, at 89 ("'Machine learning' refers to a subfield of computer science concerned with computer programs that are able to learn from experience and thus improve their performance over time.").

54 Surden, *supra* note 29, at 89.

55 *Id.* at 95-96.

56 *Id.* at 91 n.21.

57 *Id.* at 95-96.

Machine learning methods are particularly good at helping computers look at a complex set of data and model the underlying processes that generated those data.[58] The models generated through machine learning can then be applied to new data in order to predict future outcomes.[59] One of the most common tasks to which machine learning algorithms are applied is the "classification" of "objects," a catchall concept that can include anything, including people, about which one might collect data.[60] Classification is an example of what is called "supervised" machine learning, by which an algorithm learns from data that has already been "labeled" with the target "feature."[61] Features, in turn, are the "language" that machine learning algorithms uses to describe the objects within its domain.[62] The only technological limit on the kind of characteristic that can be a feature is that it must be measurable.[63] The machine learning process then creates a model based on the labeled dataset that can be used to predict the proper classification of future objects.[64]

---

58 KONONENKO & KUKAR, *supra* note 18, at 1.

59 ALPAYDIN, *supra* note 15, at 2.

60 *See* FLACH, *supra* note 14, at 13-14.

61 *Id.* at 14-15. Alternatively, machine learning can occur in an "unsupervised" or "semi-supervised" environment, where all or much of the data used to train an algorithm is unlabeled. *Id.*; Steven M. Bellovin et al., *When Enough Is Enough: Location Tracking, Mosaic Theory, and Machine Learning*, 8 N.Y.U. J.L. & LIBERTY 556, 590-96 (2014) (providing an overview of unsupervised, supervised, and semi-supervised machine learning approaches). In the law enforcement context, this would mean that machine learning methods are used to examine data about individuals where it is unknown whether the individual is engaged in criminal conduct. In such a case, a machine learning algorithm may engage in "clustering," by which similar instances are grouped together and the attributes of that group are defined. ALPAYDIN, *supra* note 15, at 155. Then, a law enforcement expert may use the groupings to gain a better understanding of a given population and develop appropriate strategies for each group. *See id.* (discussing a similar strategy in the context of a commercial enterprise analyzing its customer base). Police already interact with groups differently based on assumptions about their likelihood to engage in criminal conduct. *See, e.g.*, Wendy Ruderman, *To Stem Juvenile Robberies, Police Trail Youths Before the Crime*, N.Y. TIMES (Mar. 3, 2013), http://nyti.ms/WDjRH2 [https://perma.cc/US3X-H9Y5] (reporting on a New York City Police Department project to conduct early intervention with juveniles who are believed to be likely to engage in violent crime). The Fourth Amendment typically requires individualized suspicion based on something more than group membership, however. David A. Harris, *Using Race or Ethnicity as a Factor in Assessing the Reasonableness of Fourth Amendment Activity: Description, Yes; Prediction, No*, 73 MISS. L.J. 423, 442 (2003). Thus, this Article will focus on supervised machine learning ASAs, which are more likely to be able to provide the required individualized suspicion. Nonetheless, many of the insights unearthed herein would also apply to unsupervised machine learning.

62 FLACH, *supra* note 14, at 13.

63 *Cf. id.* at 38-39 (noting that the "main ingredient" of machine learning algorithms are "features," which "can be thought of as a kind of measurement").

64 *See id.* at 13 ("We should not normally have to go back to the domain objects themselves once we have a suitable feature representation, which is why features play such an important role in machine learning.").

More specifically,[65] in supervised machine learning the initial set of labeled data is typically subdivided into three parts: a "training set"; a "verification set" or "validation set"; and a "test set."[66] During the development of a model, the algorithm first learns an initial group of classification rules by analyzing the training set.[67] These rules are then applied to a validation or verification set, and the results are used to optimize the rules' parameters.[68] Finally, the optimized rules are applied to the test set, and the results establish both a "confidence" level and a "support" level for each rule.[69] The support level of a rule describes the percentage of objects in the test set to which the rule applies.[70] Rules with a low support level are less likely to be statistically significant.[71] Thus, to restrict which rules the algorithm will use to ensure predictions are made only on the basis of statistically significant correlations, programmers often require rules to meet a minimum support level.[72] The confidence level of a rule describes how often objects in the test set follow the rule.[73] It is, in essence, a measure of the strength of the algorithm's prediction.[74]

Machine learning methods are currently used in a wide variety of classification tasks, including identification of "spam" e-mails, optimization of productions processes, diagnosis of diseases, risk evaluation, image classification, and game playing.[75] Law enforcement's task of ferreting out crime is also one of classification: distinguishing the guilty from the innocent.[76] Or, more precisely in the Fourth Amendment context, the job of a police officer on the beat is to

---

65 This discussion is meant only as a high-level overview of a deep and complex field of mathematics. It is not intended to comprehensively discuss all of the issues that might arise in the construction of an ASA; rather, it aims to highlight some substantial concerns and the sort of information that might be helpful in identifying others.

66 *See* ALPAYDIN, *supra* note 15, at 40 (describing validation sets and test sets); FLACH, *supra* note 14, at 50 (describing training sets and test sets); KONONENKO & KUKAR, *supra* note 18, at 85 (describing validation sets).

67 FLACH, *supra* note 14, at 50.

68 KONONENKO & KUKAR, *supra* note 18, at 85.

69 *See* FLACH, *supra* note 14, at 182-84 (illustrating the creation of association rules—rules in "if *X* then *Y*" form—and explaining how to measure confidence in such rules).

70 *Id.* at 182.

71 *See* Zarsky, *supra* note 8, at 1525 (noting that a rule with limited support is "probably statistically insignificant"). Still, it is possible for a rule with a low support level to be statistically significant. *Id.*

72 *Id.*

73 *See id.* ("[The confidence level] relates to the level of 'false positives' in the process . . . .").

74 *See id.* ("[The confidence level] refers to the degree of accuracy of the rule produced.").

75 *See* ALPAYDIN, *supra* note 15, at 4-14 (providing examples); FLACH, *supra* note 14, at 1-12 (discussing "spam" identification in detail); KONONENKO & KUKAR, *supra* note 18, at 24-29 (providing examples of "typical" and "successful" applications).

76 *See* JEROME H. SKOLNICK, JUSTICE WITHOUT TRIAL: LAW ENFORCEMENT IN DEMOCRATIC SOCIETY 196 (1966) ("[T]he policeman tends to emphasize his own expertness and specialized abilities . . . to estimate accurately the guilt or innocence of suspects.").

separate those who are likely criminals from those who are likely innocent.[77] Thus, the machine learning task of classification would seem to complement the police officer's objectives.

Indeed, the outline of how an ASA could be created is straightforward. One would begin with historical data about people containing a variety of features that might be relevant to predicting a certain kind of criminal activity, perhaps including their immutable personal characteristics (e.g., age, gender, race, religion), demographic information (e.g., address, salary, occupation), and specific activities (e.g., presence on a certain street corner at a certain time, patterns of flights, or specifics of tax returns).[78] These data would also be labeled to indicate whether each included person was known to be engaged in the targeted criminal conduct or not. Machine learning methods would then be applied to these data to create a model that an ASA could apply to new data to predict which individuals are likely to be engaged in the targeted criminal activity. The confidence level of the model would determine the confidence level of the ASA's prediction that a given individual is engaged in criminal conduct.

Machine learning algorithms are not perfect, however, and mistaken predictions stem from four general sources. First, when machine learning methods are used to model complex causal systems, they necessarily rely upon approximations.[79] The causes of criminal conduct are sufficiently complex to motivate entire fields of study, but the ASA does not become a criminologist, psychologist, police officer, or sociologist. Instead, machine learning methods use patterns and correlations within the data to make a (perhaps highly educated) guess about what differentiates criminals from non-criminals.[80] Because these patterns and correlations are mere estimates of the more complex underlying

---

77 *See* Hill v. California, 401 U.S. 797, 804 (1971) ("[S]ufficient probability, not certainty, is the touchstone of reasonableness under the Fourth Amendment . . . .").

78 While there would be legal limits on the use of some personal characteristics as features in an ASA, such as the Equal Protection Clause's ban on intentional discrimination, the Fourth Amendment does not appear to impose any such restriction. *See* Whren v. United States, 517 U.S. 806, 813 (1996) ("We of course agree with petitioners that the Constitution prohibits selective enforcement of the law based on considerations such as race. But the constitutional basis for objecting to intentionally discriminatory application of laws is the Equal Protection Clause, not the Fourth Amendment."). Thus, this Article does not discuss the exceptionally difficult question of what types of data should be used to program an ASA.

79 *See* Surden, *supra* note 29, at 97 (explaining that approximation occurs "through algorithms that employ heuristics and proxies"); *see also* ALPAYDIN, *supra* note 15, at 1-2 ("We may not be able to identify the process completely, but we believe we can construct a good and useful approximation. That approximation may not explain everything, but may still be able to account for some part of the data." (emphasis omitted)).

80 *See* Surden, *supra* note 29, at 97 (discussing how machine learning uses a "strategy that has proven to be successful in automating a number of complex tasks: detecting proxies, patterns, or heuristics that reliably produce useful outcomes in complex tasks that, in humans, normally require intelligence").

phenomenon, they are inevitably inaccurate in some instances.[81] Such inaccuracies can be reduced, however, if the set of training data is large and representative.[82]

Second, inaccuracies in supervised machine learning models may come from "noise" in the training data.[83] In other words, the training data may contain information about the people described therein that is wrong.[84] For instance, a database containing the training data for an ASA targeting auto theft may list as a feature each individual's age. The database may list a particular individual as thirty years old when she was really forty years old, or perhaps list the individual as having been engaged in auto theft when she really was not. Though these kinds of noise differ in terms of their source,[85] they both can cause the machine learning process to create inaccurate models.[86] Inaccuracies resulting from noise can be mitigated by avoiding "overfitting," where machine learning methods try to match a model perfectly to the training data, and by the use of distinct test sets that were not used to train the algorithm.[87]

Third, inaccuracies can arise if an algorithm's training data is not representative of all instances of the relevant event or object in the world.[88] For instance, if our ASA—meant to identify likely auto theft—is trained on data only from a single city, the ASA may be less accurate when applied nationally if auto thieves have different criminal methods in different locales. Similarly, machine learning methods typically assume that the near future will be substantially similar to the time when the sample data were collected.[89] Thus, if the methods of auto thieves change over time, perhaps in response to police action or new technologies, our auto-theft-detecting ASA must continue to learn from new data in order to remain accurate.

---

81 *Id.* at 98.

82 *See id.* at 105-06 ("[M]achine learning algorithms often require a relatively large sample of past examples before robust generalizations can be inferred. To the extent that the number of examples (e.g., past case data) are too few, such an algorithm may not be able to detect patterns that are reliable predictors.").

83 *See* ALPAYDIN, *supra* note 15, at 30-32 (discussing how to reduce "noise"); FLACH, *supra* note 14, at 50 (same).

84 *See* ALPAYDIN, *supra* note 15, at 30 (explaining the existence of noise as additional "attributes [that] may be hidden or latent" because "they may be unobservable" (emphasis omitted)).

85 The former is called instance noise, and the latter is called label noise. FLACH, *supra* note 14, at 50.

86 *See* Surden, *supra* note 29, at 106 (explaining that inaccurate data would "produce inaccurate results because the training data was nonrepresentative" of the data generally).

87 *Id.*; *see also* FLACH, *supra* note 14, at 50.

88 *See* FLACH, *supra* note 14, at 55 ("We typically only have access to the true classes of a small fraction of the instance space and so an estimate is all we can hope to get. It is therefore important that the test set is as representative as possible."); Surden, *supra* note 29, at 106 ("[I]t is undesirable for a machine learning algorithm to detect patterns in the training data that are so finely tuned to the idiosyncrasies or biases in the training set such that they are not predictive of future, novel scenarios.").

89 ALPAYDIN, *supra* note 15, at 2.

Fourth, the choices made by humans throughout the machine learning process can cause inaccuracies in the final predictions of a machine learning algorithm.[90] At the outset, decisions must be made about what features of the objects in question should be used to construct the model.[91] In other words, before an ASA can be developed, a person must decide what facts might matter in determining whether certain behavior or characteristics are indicative of criminal conduct and how such facts can be described. For example, if an ASA is meant to detect suspicious bank transactions, should we look at the timing of each transaction? If so, is it the time of day that matters, the temporal distance of the transaction from other similar transactions, or some other time-related characteristic? The selection of the features to be analyzed is "absolutely crucial" to the success of the machine learning process.[92]

Next, data analysts must construct the training dataset.[93] This requires decisions about which databases to use, how to normalize data from different databases so that all the objects are described in terms of the same set of features, and whether to reject data that a given analyst believes is wrong or insignificant.[94] These decisions and others allow human assumptions about what correlations *should* exist in the data to color the outcome.[95] The algorithm must then be trained, a process that requires a decision about how different kinds of potential errors should be weighted.[96] For instance, an ASA programmer would need to decide whether it is worse for an innocent person to be treated as a likely criminal than for the police to ignore a person engaged in the targeted activity, and, if so, how much worse.[97] Implementing this decision will adjust the frequency with which the model predicts criminality. Finally, once an algorithm is generated, a person must answer numerous questions about its application in the field.[98] For example, how certain must a prediction be before it is reported to the police? What information will the ASA convey to the police

90 Zarsky, *supra* note 8, at 1518-19; *see also* Colonna, *supra* note 52, at 335-37 (discussing the role of various human actors in a data mining process).

91 FLACH, *supra* note 14, at 41.

92 *Id*.

93 *See* Zarsky, *supra* note 8, at 1518 (outlining the process of "data mining," an analytical technique used to "identify patterns that describe events and the links among them").

94 *Id*.

95 *See id*. at 1552 (noting that "human decisions carry particular risks of their own—such as hidden and internal biases that might be premised upon bigotry").

96 *See* KONONENKO & KUKAR, *supra* note 18, at 71 (creating a "cost matrix" to account for the occurrence in certain predictive models when "the costs of misclassifying examples from some class may be much greater than the costs of misclassifying examples from other classes").

97 *See, e.g.*, Richard Berk, *Algorithmic Criminology*, 2 SECURITY INFORMATICS, no. 5, 2013, at 1, 8 (assuming, in the probationary context, a cost ratio between false negatives and false positives of 20 to 1).

98 *See* Zarsky, *supra* note 8, at 1519 (highlighting the "opportunities for exercising human discretion" in the generation of a predictive model).

about the prediction? Such decisions will impact a model's accuracy and operation when it is put into practice.

A person also must decide whether and to what extent the machine learning algorithm will be comprehensible to humans.[99] Absent an intentional decision to the contrary, machine learning tends to create models that are so complex that they become "black boxes," where even the original programmers of the algorithm have little idea exactly how or why the generated model creates accurate predictions.[100] On the other hand, when an algorithm is interpretable, an outside observer can understand what factors the algorithm relies on to make its predictions and how much weight it gives to each factor.[101] Interpretability comes at a cost, however, as an interpretable model is necessarily simpler—and thus often less accurate—than a black box model.[102] It is certainly plausible that in the context of ASAs, society may ultimately decide to bear this cost. Yet when it comes to crime detection, the political cost of interpretability, measured in crimes unprevented and criminals uncaught, may well be quite high, thus making a black box ASA a far more attractive option.

II. INDIVIDUALIZED SUSPICION, OLD ALGORITHMS, AND ASAS

This Part lays out the existing doctrine that governs the finding of individualized suspicion to justify either a search or seizure under the Fourth Amendment. First, it articulates the two sequential steps that a police officer, magistrate, or court must undertake when determining whether probable cause or reasonable suspicion exists in a given case. Second, it establishes that ASAs play a different role in the individualized suspicion analysis than traditional algorithmic data.

A. *The Two-Step Individualized Suspicion Analysis*

In most circumstances,[103] the police must have individualized suspicion that a person is engaged in criminal conduct before they can search or seize that person.[104] The two prototypical levels of individualized suspicion are reasonable suspicion, which is required to conduct a limited search or brief seizure as

99 *See id.* at 1566 (describing the importance of "interpretability" to the success of a predictive model).

100 Edward K. Cheng, *Being Pragmatic About Forensic Linguistics*, 21 J.L. & POL'Y 541, 548 (2013).

101 ALPAYDIN, *supra* note 15, at 225-26.

102 *See* Zarsky, *supra* note 8, at 1520 ("Mandating interpretability might render the process less complex and therefore less accurate.").

103 Though the Supreme Court has authorized suspicionless searches and seizures to serve "special needs, beyond the normal need for law enforcement," City of Indianapolis v. Edmond, 531 U.S. 32, 37 (2000), the focus in this Article is on policing that serves typical crime-prevention goals.

104 Skinner v. Ry. Labor Execs.' Ass'n, 489 U.S. 602, 624 (1989).

articulated in *Terry v. Ohio*,[105] and probable cause, which is required for a "full-blown" arrest or more intrusive search.[106] To determine whether individualized suspicion exists, courts and police must look at "the totality of the circumstances—the whole picture."[107] The Court adopted the totality-of-the-circumstances approach in *Illinois v. Gates* to overturn a line of precedent that had been interpreted to limit when anonymous tips could be used to establish probable cause.[108] The Court instructed that rather than applying "[r]igid legal rules" in the individualized suspicion analysis, police and magistrates must engage in a "balanced assessment of the relative weights" of all the relevant evidence.[109] In a similar vein, the totality-of-the-circumstances approach requires police and magistrates to consider exculpatory evidence, along with any incriminating facts, in determining whether individualized suspicion exists.[110]

The totality-of-the-circumstances analysis involves two distinct, sequential steps:

> The principal components of a determination of reasonable suspicion or probable cause will be [(1)] the events which occurred leading up to the stop or search, and then [(2)] the decision whether these historical facts, viewed from the standpoint of an objectively reasonable police officer, amount to reasonable suspicion or to probable cause.[111]

The "events which occurred leading up to the stop or search" answer basic who, what, where, and when questions about the crime and the suspect: Who

---

105 392 U.S. 1 (1968).

106 *See id.* at 19, 21 ("[I]n justifying the particular intrusion the police officer must be able to point to specific and articulable facts which, taken together with rational inferences from those facts, reasonably warrant that intrusion.").

107 Navarette v. California, 134 S. Ct. 1683, 1687 (2014) (citing United States v. Cortez, 449 U.S. 411, 417 (1981)). The police's suspicion also must be individualized, in the sense that "belief of guilt must be particularized with respect to the person to be searched or seized." Maryland v. Pringle, 540 U.S. 366, 371 (2003); *see also* United States v. Cortez, 449 U.S. 411, 418 (1981) ("[A]n assessment of the whole picture . . . must raise a suspicion that the particular individual being stopped is engaged in wrongdoing."). However, the justifications and precise contours of the individualized suspicion requirement are unclear, both in the courts, and among academics. *See* Jane Bambauer, *Hassle*, 113 MICH. L. REV. 461, 468-80 (2015) (laying out various explanatory theories of individualized suspicion).

108 462 U.S. 213, 233 (1983).

109 *Id.* at 232, 234.

110 *See* Wilder v. Turner, 490 F.3d 810, 814 (10th Cir. 2007) ("We determine probable cause from the totality of the circumstances taking into account both inculpatory as well as exculpatory evidence."); Broam v. Bogan, 320 F.3d 1023, 1032 (9th Cir. 2003) ("An officer is not entitled to a qualified immunity defense, however, where exculpatory evidence is ignored that would negate a finding of probable cause."); Gardenhire v. Schubert, 205 F.3d 303, 318 (6th Cir. 2000) ("A police officer has probable cause only when he discovers reasonably reliable information that the suspect has committed a crime. And, in obtaining such reliable information, an officer cannot look only at the evidence of guilt while ignoring all exculpatory evidence." (citation omitted)).

111 Ornelas v. United States, 517 U.S. 690, 696 (1996).

is she? What did she do? Where is she? When did she engage in the relevant conduct?[112] The sources of this information are as diverse as human experience would suggest: direct observation by law enforcement personnel, tips from informants, and documentary evidence are but a few. The question for an officer, magistrate, or court at this stage is relatively straightforward: Was law enforcement's information sufficiently reliable for a reasonable officer to rely upon on in determining the historical facts?[113] The methods used to evaluate the reliability of a given piece of evidence differ depending on the nature of the evidence, and the evaluation can be quite difficult. Nonetheless, courts have extensive experience with such questions.[114]

The second step is more complicated because it presents a mixed question of law and fact.[115] An officer, magistrate, or court must decide, given the historical facts upon which a reasonable officer would rely, "whether the facts satisfy the relevant . . . constitutional standard."[116] Determinations about what behavior is adequately indicative of criminal conduct must be "practical" and "commonsense"[117] and based upon "inferences about human behavior."[118] In addition to historical facts, these inferences may be informed by "background facts" about the community at issue that are unlikely to be the subject of proof.[119] Courts are also instructed to defer to police experience and training when deciding whether individualized suspicion exists. For instance, the Court in *United States v. Brignoni-Ponce* recognized that "the officer is entitled to assess the facts in light of his experience" in detecting the criminal conduct at issue.[120] In *United States v. Arvizu*, the Court reiterated that the individualized suspicion analysis "allows officers to draw on their own experience and specialized training

112 *Id.*

113 *See* Illinois v. Rodriguez, 497 U.S. 177, 185-86 (1990) ("[W]hat is generally demanded of the many factual determinations that must regularly be made by agents of the government . . . is not that they always be correct, but that they always be reasonable.").

114 As just one example, the question of how to assess the reliability of an informant has long occupied both courts and commentators. For instance, the Supreme Court's decision in *Illinois v. Gates*, 462 U.S. 213 (1983), has been cited by courts more than 3200 times for the proposition that corroboration of an informant's tip is an important factor in establishing the tip's reliability. Search Results, WESTLAW NEXT, http://next.westlaw.com (search "462 U.S. 213" and locate headnote fourteen, "Searches and Seizures: Reliability or Credibility; Corroboration") (last visited Jan. 23, 2016).

115 *Ornelas*, 517 U.S. at 696-97.

116 *Id.* (quoting Pullman-Standard v. Swint, 456 U.S. 273, 289 n.19 (1982)).

117 *Gates*, 462 U.S. at 230.

118 Illinois v. Wardlow, 528 U.S. 119, 125 (2000).

119 For instance, in *Ornelas*, police stopped the occupants of a car with California license plates at a Milwaukee hotel in December. 517 U.S. at 699-700. The Court found that background facts like the geographical location of Milwaukee and its winter weather conditions permitted the inference that the defendants were not on vacation, but rather were in the city either to conduct business or visit family or friends. *Id.*

120 422 U.S. 873, 885 (1975).

to make inferences from and deductions about the cumulative information available to them that 'might well elude an untrained person.'"[121] Taken together, these rulings teach that the level of suspicion arising from a given set of facts "may vary depending on what a police officer knew based on her training, experience, and familiarity with the neighborhood."[122]

The Court's guidance on the inference of suspicion from historical facts leaves numerous ambiguities unresolved. First, the Court has intentionally declined to state with numerical precision how likely criminal conduct must be to satisfy the reasonable suspicion and probable cause standards. The individualized suspicion analysis "does not deal with hard certainties, but with probabilities,"[123] yet the Court has rejected any attempts to quantify the relevant probabilities.[124] Second, courts and police have little guidance on how to weigh various kinds of data in deciding whether individualized suspicion exists. Because the hard questions of suspicion involve predicting criminal conduct from noncriminal behavior, "the relevant inquiry is not whether particular conduct is 'innocent' or 'guilty,' but the degree of suspicion that attaches to particular types of noncriminal acts."[125] Yet courts rarely possess empirical data that might prove or disprove a correlation between certain conduct and criminal activity.[126] And even when they do, courts are typically

121 534 U.S. 266, 273 (2002) (quoting United States v. Cortez, 449 U.S. 411, 418 (1981)).

122 Kit Kinports, *Veteran Police Officers and Three-Dollar Steaks: The Subjective/Objective Dimensions of Probable Cause and Reasonable Suspicion*, 12 U. PA. J. CONST. L. 751, 755-56 (2010).

123 *Cortez*, 449 U.S. at 418 (1981).

124 *See* Maryland v. Pringle, 540 U.S. 366, 371 (2003) ("The probable-cause standard is incapable of precise definition or quantification into percentages because it deals with probabilities and depends on the totality of the circumstances."); *see also* Goldberg, *supra* note 46, at 794 (calling for the establishment of a minimum numerical threshold for probable cause).

125 Illinois v. Gates, 462 U.S. 213, 244 n.13 (1983).

126 *See* Illinois v. Wardlow, 528 U.S. 119, 124-25 (2000) ("In reviewing the propriety of an officer's conduct, courts do not have available empirical studies dealing with inferences drawn from suspicious behavior, and we cannot reasonably demand scientific certainty from judges or law enforcement officials where none exists."); *cf.* Andrew E. Taslitz, *Cybersurveillance Without Restraint? The Meaning and Social Value of the Probable Cause and Reasonable Suspicion Standards in Governmental Access to Third-Party Electronic Records*, 103 J. CRIM. L. & CRIMINOLOGY 839, 862-63 (2013) ("Absent [generalized, objective probability] data, it is hard to see how a specific number can exist to serve as an anchor."). *But cf.* Tracey L. Meares & Bernard E. Harcourt, *Foreword: Transparent Adjudication and Social Science Research in Constitutional Criminal Procedure*, 90 J. CRIM. L. & CRIMINOLOGY 733, 750-52 (2000) (arguing that courts should utilize social science and empirical research because doing so would improve constitutional decisionmaking); David Rudovsky & Lawrence Rosenthal, *Debate: The Constitutionality of Stop-and-Frisk in New York City*, 162 U. PA. L. REV. ONLINE 117, 119 (2013), http://www.pennlawreview.com/debates/12-2013/Stop-and-Frisk.pdf [https://perma.cc/9GGP-2Q32] ("The Court has not required police or prosecutors to demonstrate by empirical data that the characteristics relied upon—for example, that the suspect was acting suspiciously, had fled from police, had bulges in his pockets, or was engaged in 'furtive movements'—are actually predictive of criminal conduct.").

untrained in how to assess that data.[127] Finally, the Supreme Court has not explained how courts should decide whether to defer to police experience in a given case and how much deference to give.[128]

B. *Algorithms in the Individualized Suspicion Analysis: The Old and the New*

Law enforcement officials have used the output of automated algorithms for decades.[129] Breathalyzers run on algorithms that state the amount of alcohol in an individual's blood based on the amount of alcohol in a sample of that individual's breath.[130] Radar guns send radio waves at a certain frequency in the direction of a moving automobile, measure the frequency of reflected waves that return, and calculate the speed of the automobile based on the change in frequency.[131] A DNA sample from a crime scene can be matched

127 *See* Andrew Jurs, *Judicial Analysis of Complex & Cutting-Edge Science in the* Daubert *Era: Epidemiologic Risk Assessment as a Test Case for Reform Strategies*, 42 CONN. L. REV. 49, 73-75 (2009) (describing studies that show "wide variance" in judges' capacities to handle statistical evidence and that "support[] the conclusion that judges fare poorly with statistical analysis"). For instance, in *Navarette v. California*, the Supreme Court asked whether a reliable tip that a truck had nearly run another vehicle off the road created reasonable suspicion that the driver of the truck was intoxicated. 134 S. Ct. 1683, 1690 (2014). In concluding that it did, the Court relied on a pamphlet on the visual detection of impaired motorists issued by the National Highway Traffic Safety Administration. *Id.* at 1691. This pamphlet lists probabilities that a driver exhibiting certain behaviors is intoxicated, including "[a]lmost striking a vehicle or other object." NAT'L HIGHWAY TRAFFIC SAFETY ADMIN., U.S. DEP'T OF TRANPS., THE VISUAL DETECTION OF DWI MOTORISTS 5 (2010), http://nhtsa.gov/staticfiles/nti/pdf/808677.pdf [https://perma.cc/LQ2D-CUSW]. The pamphlet is light on an explanation of the foundation for its claimed probabilities, stating only that the pamphlet is based on a prior NHTSA study and "3 field studies involving hundreds of officers and more than 12,000 enforcement stops." *Id.* at 4. Yet the Court seemed unconcerned with exploring the reliability of this statistical information, despite the weight it placed on the pamphlet. *See Navarette*, 134 S. Ct. at 1691 (concluding that running another car off the highway "bears too great a resemblance to paradigmatic manifestations of drunk driving to be dismissed as an isolated example of recklessness"). A comparison of the Court's questioning of the available data with the analysis of the same data, *see* Joshua C. Teitelbaum, *Probabilistic Reasoning in* Navarette v. California, 62 UCLA L. REV. DISCOURSE 158 (2014), reveals the shallowness of the Court's examination.

128 *Cf.* L. Song Richardson, *Police Efficiency and the Fourth Amendment*, 87 IND. L.J. 1143, 1155 (2012) (noting that "courts consistently fail to determine whether the inferences drawn by the officer conducting the stop are actually *entitled* to any weight").

129 *See* Maryland v. King, 133 S. Ct. 1958, 1966 (2013) (noting that the first use of forensic DNA analysis by a court occurred in 1986).

130 *See* Okorie Okorocha & Matthew Strandmark, *Alcohol Breath Testing: Is There Reasonable Doubt?*, 27 SYRACUSE J. SCI. & TECH. L. 124, 130 (2012) (discussing the commonly accepted "Partition Ratio" between alcohol in breath and in blood).

131 *See* Ryan V. Cox & Carl Fors, *Admitting Light Detection and Ranging (LIDAR) Evidence in Texas: A Call for Statewide Judicial Notice*, 42 ST. MARY'S L.J. 837, 842-43 (2011). The other main speed detection device, known as a LIDAR gun, uses a simple algorithm to calculate the speed of an automobile based on the time it takes repeated pulses of light to reflect back to the device. *Id.* at 849.

against stored DNA profiles using search algorithms.[132] Emerging algorithmic biometric technologies aim to enhance the ability of police to identify suspects and track their movements.[133]

These traditional technologies can be exceptionally helpful to police in establishing the "historical facts" of what happened, when it happened, and who was involved.[134] In DNA matching, algorithmic searches of databases reveal either who was at the scene of a given crime or whether a given person was at the scene of other unsolved crimes.[135] Radar guns show how fast a vehicle is moving at a given moment.[136] Newer biometric technologies can provide substantially more information about a suspect's location and movements.[137] All of these technologies help police establish facts that can be ascertained to a definable level of certainty: for example, the quantity of alcohol in a driver's bloodstream can be certain within some calibration level,[138] or the identity of DNA found at a crime scene can be determined to some statistical level of confidence.[139] And because these technologies answer questions of fact, a court can focus its analysis on the familiar question of the accuracy of the technology used to determine the fact at issue.[140]

Unlike the output of traditional technologies, the output of an ASA is directed at the mixed question of law and fact of whether the historical facts

---

132 *Frequently Asked Questions (FAQs) on the CODIS Program and the National DNA Index System*, FBI, http://www.fbi.gov/about-us/lab/biometric-analysis/codis/codis-and-ndis-fact-sheet [https://perma.cc/T3BN-4KE5] (last visited Jan. 23, 2016). Technology that automates the process of extracting DNA from buccal swabs is currently being tested as well. *Rapid DNA or Rapid DNA Analysis*, FBI, http://www.fbi.gov/about-us/lab/biometric-analysis/codis/rapid-dna-analysis [https://perma.cc/M6UU-E9TN] (last visited Jan. 23, 2016).

133 *See* Margaret Hu, *Biometric ID Cybersurveillance*, 88 IND. L.J. 1475, 1500-08 (2013) (reviewing examples of government use of biometric identification devices).

134 *See, e.g.*, *King*, 133 S. Ct. at 1966 (noting that police used a DNA match to obtain search warrant); United States v. Flores, No. 4:08CR3059, 2008 WL 4104136, at *2 (D. Neb. Aug. 28, 2008) (describing how a state trooper used a radar gun to detect speeding).

135 Jason Kreag, *Letting Innocence Suffer: The Need for Defense Access to the Law Enforcement DNA Database*, 36 CARDOZO L. REV. 805, 815-16 (2015).

136 *See supra* note 131 and accompanying text.

137 *See* Hu, *supra* note 133, at 1490-92 (discussing "identity verification" and "identity determination" systems).

138 *See* COMM. ON IDENTIFYING THE NEEDS OF THE FORENSIC SCIENCES COMMUNITY, NAT'L RESEARCH COUNCIL, STRENGTHENING FORENSIC SCIENCE IN THE UNITED STATES: A PATH FORWARD 117 (2009) (discussing how methods that measure the level of blood alcohol do so within a confidence interval).

139 *See, e.g.*, People v. Nelson, 185 P.3d 49, 53 (Cal. 2008) ("[T]he prosecution presented evidence that the DNA profile on the vaginal swab would occur at random among unrelated individuals in about one in 950 sextillion African-Americans, one in 130 septillion Caucasians, and one in 930 sextillion Hispanics.").

140 *See, e.g.*, Hall v. State, 297 S.W.3d 294, 298 (Tex. Crim. App. 2009) (upholding suppression of evidence on the ground that the state provided no evidence that a speed gun "supplies reasonably trustworthy information").

are sufficient to establish reasonable suspicion or probable cause.[141] ASAs look at data from other sources and predict the probability that an observed person with a certain set of "features"[142] is engaged in criminal conduct.[143] In providing a prediction of criminality, the ASA's examination of data overlaps with the second step in the individualized suspicion analysis.[144] As such, ASAs provide a kind of data to the Fourth Amendment analysis that serves an analytically different role than the output of traditional algorithms.

To illustrate this distinction, consider the case of *People v. Nelson* from the California Supreme Court.[145] In *Nelson*, a nineteen-year-old college student disappeared after telephoning her mother to report that her car would not start.[146] The victim's body was found two days later.[147] After more than twenty-five years, police were able to obtain a sample of a suspect's DNA and match it to DNA collected near where the victim's body was found.[148] Almost conclusively, the DNA match established the historical fact that the defendant had been at the location where the body was found close enough in time that the DNA that he left behind had not degraded or otherwise disappeared.[149]

Yet this historical fact, standing alone, does not tell police how likely it was that the defendant was guilty. Rather, to connect Nelson to the murder, more facts are needed. In *Nelson*, that "more" included: that the victim had been raped before she was killed, that the DNA sample was collected from semen on her body and clothing, and that the victim was seen in a car matching one owned by the defendant shortly before her death.[150] Traditionally, a human being must consider the DNA match together with those additional facts to decide that a sufficient probability of guilt existed to justify arresting Nelson for the murder. The novelty of an ASA is its potential to step into the shoes of that human being by analyzing groups of disparate facts together and drawing conclusions about the probability of an individual's guilt.

---

141 An ASA could be incorporated into a larger computer system that collects data, which is then fed into the ASA. *Cf.* Cardwell, *supra* note 4 (discussing a computer system that would both record video images and analyze them to decide whether crime is likely occurring). Nonetheless, for the sake of analytical precision, these two functions should be considered separately.

142 This term is used in the technical sense described above, not in reference to the physical features of a suspect. *See supra* notes 62–63 and accompanying text.

143 *See* Brinegar v. United States, 338 U.S. 160, 175 (1949) ("In dealing with probable cause, however, as the very name implies, we deal with probabilities.").

144 The extent of the overlap is discussed *infra* Part III.

145 185 P.3d 49 (Cal. 2008).

146 *Id.* at 53.

147 *Id.*

148 *Id.*

149 *Id.* at 60.

150 *Id.* at 53.

III. THE INSUFFICIENCY OF AN ASA'S PREDICTION

Say that an ASA predicts a 60% likelihood that a specific person is selling drugs on a street corner, and a police officer, upon receiving the prediction, stops the suspect, frisks him, and finds drugs. If the defendant challenges the stop and frisk, can the prosecution rely solely on the ASA's prediction, or does the Fourth Amendment require something more? The "collective knowledge" doctrine, which allows one police officer to engage in a search or seizure based on the instruction of another officer who knows facts that establish individualized suspicion,[151] provides a framework for answering this question. If the ASA's prediction is the equivalent of an officer's instruction, then under the constructive knowledge doctrine an officer would be justified in acting on that prediction, standing alone. The first Section of this Part lays out the scope and operation of the collective knowledge doctrine, including how some courts have extended the doctrine to apply to constructive knowledge, and scholars' criticisms of the expanded doctrine. The second Section explores the application of the doctrine to an ASA's output. This Part makes two arguments: first, that the expanded "constructive knowledge" doctrine, as applied to ASAs, would eviscerate the individualized suspicion requirement; and second, that an ASA's prediction is not sufficient to create individualized suspicion.

A. *The Collective and Constructive Knowledge Doctrines*

In *Whiteley v. Warden*, the Court held that when one officer asks another officer to help her with the execution of a warrant, the second officer is entitled to presume that the first officer provided a magistrate with sufficient information to justify a finding of probable cause.[152] The Court expanded this rule in *United States v. Hensley* beyond situations involving a warrant to allow an officer to rely on a flyer or bulletin if: (1) the officer acted in "objective reliance" on the flyer or bulletin;[153] and (2) the flyer or bulletin was based on articulable facts sufficient to establish the necessary individualized suspicion.[154] Lower courts have since applied the collective knowledge rule to justify searches and seizures in a wide variety of situations in which an officer is instructed to undertake the search but is not provided information sufficient to independently find the proper level of individualized suspicion.[155]

---

151 United States v. Lyons, 687 F.3d 754, 766 (6th Cir. 2012).

152 401 U.S. 560, 568 (1971).

153 United States v. Williams, 627 F.3d 247, 252 (7th Cir. 2010).

154 United States v. Hensley, 469 U.S. 221, 232 (1985). The search or seizure also must not exceed the scope justified by the underlying individualized suspicion. *Williams*, 627 F.3d at 252-53.

155 For a thorough discussion of the various permutations of the collective knowledge rule, see Stern, *supra* note 43, at 1094-1105.

The rationale behind the collective knowledge rule is largely pragmatic: "[E]ffective law enforcement cannot be conducted unless police officers can act on directions and information transmitted by one officer to another and . . . officers . . . cannot be expected to cross-examine their fellow officers about the foundation for transmitted information."[156] Requiring that the officer who engages in a search or seizure must herself have the necessary individualized suspicion would be a "crippling restriction[] on our law enforcement."[157] Instead, it is sufficient that at some point, an individual trained in making individualized suspicion determinations, whether a magistrate or a law enforcement officer, had sufficient knowledge to conclude that the individual be seized or searched.[158] Mandating that a person trained in individualized suspicion determinations find probable cause or reasonable suspicion seems to ensure that reliance on the instruction to stop is objectively reasonable.[159] In addition, an individual searched or seized pursuant to the collective knowledge doctrine has "minimal" interests at stake.[160] Because the suspect could have been seized by one officer, she loses little in the way of security or privacy when she is stopped by another officer at the instruction of the first.[161]

While the constructive knowledge doctrine applies the general idea underlying the collective knowledge doctrine of police reliance on other officers' knowledge, it does so without the same strict requirements.[162] The broadest view of the constructive knowledge doctrine, and the one most relevant here, is one where no one officer possesses facts sufficient to establish the needed individualized suspicion, but the aggregation of several officers' knowledge would meet the standard.[163]

Specifically, this version of the doctrine omits both the requirement that a single individual trained in individualized suspicion assessments evaluate the facts and the need for the knowledgeable officers to have communicated with

156 *Hensley*, 469 U.S. at 231 (quoting United States v. Robinson, 536 F.2d 1298, 1299 (9th Cir. 1976)).

157 United States v. Lyons, 687 F.3d 754, 766 (6th Cir. 2012).

158 *See* United States v. Colon, 250 F.3d 130, 135-36 (2d Cir. 2001) ("A primary focus in the imputed knowledge cases is whether the law enforcement officers initiating the search or arrest, on whose instructions or information the actual searching or arresting officers relied, had information that would provide reasonable suspicion or probable cause to search or arrest the suspect.").

159 *See id.* at 138 (finding that police officer reliance on information known to a civilian 911 operator was not objectively reasonable because the operator was not trained in making individualized suspicion determinations).

160 *Hensley*, 469 U.S. at 232.

161 *See* Stern, *supra* note 43, at 1090 (explaining that the collective knowledge doctrine increases the efficiency of a police department without diminishing the protections afforded by the probable-cause doctrine).

162 *Id.* at 1105.

163 *See id.* at 1106-09 & nn.80-101 (discussing cases in which courts applied the constructive knowledge doctrine).

each other.[164] Nevertheless, courts generally limit the scope of the constructive knowledge doctrine to officers who are working closely together.[165]

Academics and dissenting judges have criticized the constructive knowledge doctrine for not meaningfully enhancing law enforcement expediency, reasoning that police communication is inexpensive and increases accuracy.[166] Moreover, the constructive knowledge doctrine removes the concept of "belief" and the perspective of a "reasonable officer" from the definitions of probable cause and reasonable suspicion.[167] After all, a court cannot inquire into whether "facts and circumstances within the officer's knowledge . . . are sufficient to warrant a prudent person, or one of reasonable caution, in believing . . . that the suspect" is engaged in criminal conduct if no single officer knew the information and could believe something about it.[168] Finally, as massive quantities of information become readily available to law enforcement agencies through fusion centers and communication technologies,[169] a broad reading of the constructive knowledge doctrine would render the individualized suspicion requirement meaningless in most situations.[170] This threat has led one scholar to suggest that the constructive knowledge doctrine would turn the police into "something like *Star Trek*'s Borg Collective," in that officers would be able to rely upon what is known by any other officer anywhere, at least for the purposes of providing a post hoc justification for a search or seizure.[171]

B. *Applying the Doctrines to ASAs*

The power of ASAs to analyze large quantities of data in making their predictions underscores the threat that the constructive knowledge doctrine poses to the Fourth Amendment's individualized suspicion requirement.[172] Where

164 *Id.* at 1116-17.

165 *Id.* at 1108-09.

166 *Id.* at 1111-12 & 1111 n.107.

167 *Id.* at 1112-13.

168 Michigan v. DeFillippo, 443 U.S. 31, 37 (1979).

169 *See* Citron & Pasquale, *supra* note 9, at 1448-55 (discussing the role of fusion centers in collecting and distributing information to law enforcement agencies).

170 *See* Poniatowski, *supra* note 43, at 842 ("The communication/teamwork requirement of the constructive-knowledge doctrine has been applied leniently by some courts, threatening to eviscerate the [probable cause] requirement altogether." (footnote omitted)).

171 Stern, *supra* note 43, at 1114.

172 Of course, there are those who argue that the individualized suspicion requirement fails to advance the interests underlying the Fourth Amendment in some situations and should be replaced. *See, e.g.*, Bernard E. Harcourt & Tracey L. Meares, *Randomization and the Fourth Amendment*, 78 U. CHI. L. REV. 809, 816 (2011) (contending that "individualized suspicion" should be abandoned in favor of randomized searches). These proposals may have merit, but the constructive knowledge doctrine threatens to create an even more dismal state of affairs than what currently exists, as the individualized suspicion doctrine would survive, but without the few teeth it now has.

police have access to the massive troves of information contained in fusion centers, the doctrine already opens the door to "arrest first, justify later" policing.[173] But permitting police to claim constructive awareness of an ASA's predictions of criminality without any requirement that the predictions be communicated to the officer conducting a search or seizure would further encourage police to ignore individualized suspicion requirements.[174] Particularly as criminal laws have proliferated to the point that "everyone is a criminal if prosecutors look hard enough,"[175] applying the constructive knowledge doctrine to ASAs could permit the police to stop anyone and later find a prediction of crime to justify the intrusion.

Applying the collective knowledge doctrine to ASAs, however, presents a less immediately discomfiting dystopia. Upon receipt of an ASA's prediction, police could search or seize a person identified by an ASA without engaging in any independent assessment of the facts to determine whether individualized suspicion exists. Reliance by the police officer would be permitted if the ASA's prediction were analogous to an instruction to arrest by an individual trained in making individualized suspicion determinations.[176] In some sense, an ASA is very well trained in making individualized suspicion determinations, as it can provide a quantifiable prediction of criminality based on the available data (e.g., there is a 60% chance that the person on a certain street corner is dealing drugs).[177] So long as we have reason to believe that the ASA is accurate,[178] the ASA's prediction is arguably analogous to an assertion of the existence of probable cause or reasonable suspicion by a person trained in making such assessments. The Court, after all, has repeatedly explained that individualized suspicion deals with probabilities,[179] and an ASA can quantify those probabilities like no technologies before it.

This analogy between an ASA and a trained person fails for two related reasons, however. First, it depends on a fundamental misunderstanding of the question that the individualized suspicion standard asks. Second, the analogy

---

173 Poniatowski, *supra* note 43, at 851.

174 *Cf. id.* at 849 ("At the heart of this risk is an arresting officer's reliance on the hope that information exists that would condemn a suspect.").

175 Glenn Harlan Reynolds, *Ham Sandwich Nation: Due Process When Everything Is a Crime*, 113 COLUM. L. REV. SIDEBAR 102, 104 (2013).

176 *See supra* note 158 and accompanying text.

177 When a machine learning algorithm generates an association rule from data—that is, it predicts that when certain antecedent conditions are satisfied, some consequent condition will also exist—a confidence measurement describes the accuracy of the rule when the antecedent conditions are satisfied. KONONENKO & KUKAR, *supra* note 18, at 233-34.

178 The factors that go into determining the accuracy of an ASA are discussed *infra* Section IV.C.

179 New Jersey v. T.L.O., 469 U.S. 325, 346 (1985) (citing Hill v. California, 401 U.S. 797, 804 (1971); Illinois v. Gates, 462 U.S. 213, 232 (1983); United States v. Cortez, 449 U.S. 411, 418 (1981); Brinegar v. United States, 338 U.S. 160, 176 (1949)).

fails to appreciate differences in how humans and machines examine factual situations. To see these flaws, we must start by recalling that the probable cause and reasonable suspicion determinations require a consideration of the totality of the circumstances.[180] As its name suggests, the totality-of-the-circumstances approach demands a consideration of *all* evidence relevant to the question of how likely it is that the targeted individual is engaged in criminal activity, including exculpatory evidence.[181]

For an ASA's prediction to be sufficient to justify a search or seizure, it too must engage in a totality-of-the-circumstances analysis. But, at least under current technological constraints, ASAs are fundamentally incapable of doing so. As with any machine learning process, an ASA is only as good as the data its programmers choose to provide it, either in training or in real-world application.[182] This is because the data provided to an ASA constitutes the sum total of what the algorithm "knows" about the world; the ASA cannot identify new types of relevant data that are not currently contained in its dataset and then seek out those data.[183] Thus, an ASA trained on a small dataset "knows" very little, while an ASA trained on an enormously robust dataset "knows" quite a lot.[184] But even the latter ASA is limited in making its predictions to analysis of the data within its dataset, and it cannot consider other facts that might be relevant but that were not included. In contrast, human beings are always at least potentially capable of including a new piece of relevant information in an analysis.[185]

This distinction matters enormously for the capacity of an ASA to engage in a totality-of-the-circumstances analysis. The kinds of information that might be relevant to an individualized suspicion determination are infinite.[186] While

---

180 *See* Alabama v. White, 496 U.S. 325, 330 (1990) ("[T]otality of the circumstances . . . must be taken into account when evaluating whether there is reasonable suspicion." (citation omitted)); *see also* Illinois v. Gates, 462 U.S. 213, 230-31 (1983).

181 *See supra* note 110 and accompanying text.

182 *See supra* notes 83–89 and accompanying text; *see also* Surden, *supra* note 29, at 106 ("In general, machine learning algorithms are only as good as the data that they are given to analyze.").

183 While "active learning," where a machine learning algorithm chooses the data with which to be trained, is a goal of a sub-field of machine learning research, even an active learning algorithm is limited to considering the data provided to it. *See* Burr Settles, *Active Learning Literature Survey* 27 (Univ. of Wis.-Madison, Computer Sciences Technical Report 1648, 2010), http://burrsettles.com/pub/settles.activelearning.pdf [https://perma.cc/A38G-E3MT].

184 *See supra* notes 69–74 and accompanying text.

185 This is not to say that humans do not suffer from cognitive biases that may substantially undermine their ability to make accurate individualized suspicion determinations. *See* Christopher Slobogin, *Why Liberals Should Chuck the Exclusionary Rule*, 1999 U. ILL. L. REV. 363, 403-4 (discussing heuristics that interfere with the ability of judges to make accurate probable cause determinations).

186 This observation can be confirmed by any criminal procedure professor who has taught a student fond of imagining his or her own hypotheticals.

an ASA may be trained with a database that contains all the facts that are most relevant in a large majority of cases, that database cannot contain all the facts that are relevant in every case.[187] As a result, an ASA cannot consider the "whole picture" regarding a person's potential criminality as required by the Fourth Amendment.[188]

To illustrate this point, imagine an ASA targeting the selling of narcotics on street corners. The ASA has access to information from a variety of inputs, such as closed-circuit cameras, license-plate readers, and facial recognition technology. Based on both historic and real-time data from these sources, it predicts when specific individuals are engaging in hand-to-hand drug transactions. One day it issues an alert predicting that an individual is more likely than not selling narcotics on a street corner. A patrol officer in uniform is dispatched to investigate and witnesses the suspect and passers-by briefly exchanging items by hand. As she approaches the suspect, the officer makes two observations. First, she notes that the suspect sees her and does not change his behavior. Second, she sees a passer-by drop an item recently received from the suspect on the ground, picks the item up, and notes that it is a flyer for a church event.

Both observed facts tend to diminish the likelihood that the suspect is engaged in criminal activity, but neither are captured in the ASA's dataset. A totality-of-the-circumstances analysis of individualized suspicion must account for these facts, however, and our ASA has failed to do so.[189] But now that we know the identified facts matter, the ASA can be programmed to incorporate them in future predictions. Yet this does not resolve the underlying problem that the ASA must consider *every* fact that *might* impact the existence of individualized suspicion. To do so the ASA must either be able to process all known information or have been programmed in advance to "know" all potentially relevant information. Neither is feasible: the former requires more processing power than is currently available and the latter requires impossible foresight. Thus, a person trained in making individualized suspicion determinations must be the final assessor of the totality-of-the-circumstances, including both the ASA's prediction and any other relevant available data, in order to decide whether the probable cause or reasonable suspicion standards are met.[190]

187 An ASA's capacity to consider all potentially relevant facts is also limited because it can only consider "features" that are quantifiable. *See supra* note 63 and accompanying text.

188 Navarette v. California, 134 S. Ct. 1683, 1687 (2014) (quoting United States v. Cortez, 449 U.S. 411, 417 (1981)).

189 The ASA also fails to take in new information after issuing its prediction and adjust that prediction accordingly. With sufficient computing power, however, this flaw could be corrected.

190 The human capacity to be open to consideration of new relevant data is related to what Orin Kerr has called "instinct" or "intuition" in arguing against the quantification of individualized suspicion standards. *See* Orin Kerr, *Why Courts Should Not Quantify Probable Cause*, *in* THE POLITICAL

Requiring a human to assess the totality-of-the-circumstances, however, may reduce the overall accuracy of police searches and seizures. While some of the additional evidence that a human being will consider may clearly confirm or rebut the ASA's prediction, the human officer may not be able to accurately assess the impact of other evidence on the analysis.

For instance, imagine that an ASA predicts that a specific individual, who has been going from car to car in a parking lot and then spends five minutes trying to get into one vehicle before walking away, has a 52% chance of being engaged in auto theft. If the ASA is accurate, the odds establish probable cause to arrest the suspect.[191] An officer is told of the ASA's prediction and approaches the individual near the parking lot. The officer asks him for an explanation, and the individual provides a story that innocently explains his actions. If the officer finds the story credible, that explanation would destroy the officer's probable cause. The officer could not validly arrest the suspect considering the totality of all the circumstances known to him.

Yet, there are serious reasons to doubt the officer's ability to evaluate accurately the totality of the circumstances in many cases. First, studies have shown that police, like laypeople, are not good lie detectors.[192] Other cognitive roadblocks also may hinder an officer's capacity to make accurate individualized suspicion determinations.[193] For example, an officer's initial perception of a suspect's criminality may be overly influenced by the suspect's facial expression when approached by the officer or the suspect's nervous reaction if the officer

---

HEART OF CRIMINAL PROCEDURE: ESSAYS ON THEMES OF WILLIAM J. STUNTZ 131, 138 (Michael Klarman et al. eds., 2012). Kerr's "instinct" involves the recognition that sometimes there is important information missing from the facts currently available in an individualized suspicion analysis. *See id.* at 138-39 (noting that the desired evidence may tie the suspect more closely to the crime or its absence may suggest something important about police motives). The missing data may be inculpatory, exculpatory, or relevant from a policy standpoint. *Id.* Yet courts have typically resisted imposing any duty on police officers to investigate further once known facts are sufficient to establish individualized suspicion. *See, e.g.*, Ahlers v. Schebil, 188 F.3d 365, 371 (6th Cir. 1999) ("Once probable cause is established, an officer is under no duty to investigate further or to look for additional evidence which may exculpate the accused."). Thus, the fact that machine learning algorithms lack Kerr's "instinct" to look for more information does not provide an independently sufficient reason why ASAs cannot undertake a totality-of-the-circumstances analysis.

191 *See* Craig S. Lerner, *The Reasonableness of Probable Cause*, 81 TEX. L. REV. 951, 996 (2003) (explaining that probable cause, if quantified, falls somewhere between 0.01% and 90% certainty).

192 *See* Andrew E. Taslitz, *Police Are People Too: Cognitive Obstacles to, and Opportunities for, Police Getting the Individualized Suspicion Judgment Right*, 8 OHIO ST. J. CRIM. L. 7, 27-29 (2010) (citing research that found that police are no better than laypeople at detecting deception). *See generally* Eugenio Garrido et al., *Police Officers' Credibility Judgments: Accuracy and Estimated Ability*, 39 INT'L J. PSYCHOL. 254 (2004) (summarizing past experiments that showed that both police and undergraduates are little better than chance at detecting lies, and confirming this result based on new research).

193 *See generally* Taslitz, *supra* note 192.

appears unfriendly.[194] Once the officer forms a negative impression of the suspect, human nature makes the officer resistant to changing it.[195]

In addition to troubles with credibility determinations, facts that society may not want to be part of the analysis—like a suspect's race, religion, or national origin—influence the officer's assessment of criminality. Racial minorities, and particularly African-American males, have long been stereotyped as "violent, criminal, and dangerous."[196] These stereotypes can unconsciously impact how police assess criminality. Whites react negatively to faces displaying features typically associated with African-Americans.[197] African-Americans draw attention more quickly than Whites.[198] Observers viewing ambiguous behavior interpret the behavior differently depending on the race of the observed person.[199] Negative implicit biases are also prevalent with respect to non-White races other than African-Americans, as well as traits other than race, including religion and national origin.[200] L. Song Richardson has argued convincingly that these unconscious biases infect police assessments of individualized suspicion.[201]

Finally, incorporating the output of an ASA into the totality-of-the-circumstances analysis in an accurate and meaningful way is likely to be quite challenging.[202] For all these reasons, requiring police to be open to additional data and to include such data in their totality-of-the-circumstances analysis for each suspect will likely lead to more police errors: namely, searches and seizures of the innocent and instances of the guilty going free.

This result is certainly not ideal, but the Fourth Amendment and its individualized suspicion standards are not in place to maximize police accuracy; rather, they aim to ensure individualized justice.[203] In other words, the Fourth Amendment would not be satisfied if a police agency conducted ten searches, five on suspects who were almost certainly engaged in criminal activity and five on suspects who almost certainly were not, on the ground

194 *See id.* at 23.

195 *Id.* at 29.

196 Richardson, *supra* note 128, at 1147.

197 Taslitz, *supra* note 192, at 19.

198 Richardson, *supra* note 128, at 1150.

199 *See id.* at 1148-49 (summarizing research finding that "the threshold for labeling ambiguous behavior as aggressive or violent is lower for Blacks than for Whites").

200 *See* Brian A. Nosek et al., *Pervasiveness and Correlates of Implicit Attitudes and Stereotypes*, 18 EUR. REV. SOC. PSYCHOL. 36, 43-57 (2007) (summarizing results of Implicit Association Tests with respect to a broad range of individual traits).

201 *See generally* Richardson, *supra* note 128.

202 *See infra* Part IV.

203 *See* ANDREW E. TASLITZ, RECONSTRUCTING THE FOURTH AMENDMENT: A HISTORY OF SEARCH AND SEIZURE, 1789-1868, at 45-54 (2006) (describing the evolution of the Fourth Amendment).

that *on average* probable cause existed.[204] Rather, probable cause must exist for each suspect.[205] Put another way, in most circumstances the Fourth Amendment entitles *each* suspect to an assessment of whether individualized suspicion exists based on *all* available facts relating to her potential guilt.[206] Recent approaches to probable cause and reasonable suspicion may have undermined the individualized nature of these standards,[207] but the requirement of a totality-of-the-circumstances analysis remains, even if that requirement means that police will make more mistakes.

IV. INCLUDING ASAS IN THE TOTALITY-OF-THE-CIRCUMSTANCES ANALYSIS

If an ASA's output alone cannot satisfy the individualized suspicion requirements of the Fourth Amendment, it must be considered as a part of the totality-of-the-circumstances analysis. As explained above, ASAs are unique data sources in that they aim to assist in the second step of the individualized suspicion analysis by providing information about when one should infer criminality from certain historical facts.[208] Though they can be the source of bad law, when it comes to new technologies and the Fourth Amendment, analogies to existing data sources are the currency of the realm.[209] A good analogy should help courts and police identify the factors that will help them separate reliable ASAs from unreliable ones. This Part will explore three potential analogies.[210] First, ASAs are similar to police profiles, such as

---

204 *See id.* at 49 ("What matters most . . . is that probable cause required specific, trustworthy information to make real the implicit aspiration toward *individualized* justice." (emphasis added)).

205 The Fourth Amendment has sometimes been interpreted by the Court in a manner consistent with it being a collective, rather than an individual, right. *See, e.g.*, Thomas K. Clancy, *The Fourth Amendment as a Collective Right*, 43 TEX. TECH L. REV. 255, 263-94 (2010). Nonetheless, this interpretation is generally limited to situations outside of police interdiction of ordinary criminal activity. *See id.* at 273-90 (discussing examples of situations where the Fourth Amendment is viewed as a collective right).

206 As Taslitz correctly notes, the fact that there are exceptions to the individualized suspicion requirement does not mean that the Fourth Amendment's commitment to individualized justice is any less important. Andrew E. Taslitz, *Search and Seizure History as Conversation: A Reply to Bruce P. Smith*, 6 OHIO ST. J. CRIM. L. 765, 775 (2009).

207 *See generally* David A. Harris, *Particularized Suspicion, Categorical Judgments: Supreme Court Rhetoric Versus Lower Court Reality Under* Terry v. Ohio, 72 ST. JOHN'S L. REV. 975 (1998) (critiquing how lower courts have permitted categorical judgments to satisfy the Fourth Amendment's individualized suspicion standards).

208 *See supra* Section II.B.

209 *See* Kerr, *supra* note 42, at 875-76 ("Judges struggle to understand even the basic facts of [new] technologies, and often must rely on the crutch of questionable metaphors to aid their comprehension. Judges generally will not know whether those metaphors are accurate . . . ." (footnote omitted)).

210 A skeptic may also argue that an ASA's prediction is an "inchoate and unparticularized suspicion or 'hunch,'" which is insufficient to generate reasonable suspicion or probable cause. Terry

those frequently used to identify drug couriers, human traffickers, child abusers, or terrorists,[211] as they utilize historical information to identify traits that are commonly held by criminals with the goal of predicting future criminality.[212] Second, algorithms are akin to informants in that people outside of law enforcement are providing information to police, albeit indirectly through the ASA.[213] Third, algorithms are similar to drug-sniffing dogs in that both resemble "black boxes" that create outputs from known inputs and potentially uncertain processes.[214] This Part addresses each analogy in turn. The first two are ultimately unhelpful for substantive and procedural reasons. The third is more useful, though it is also imperfect. This Part concludes by identifying lingering challenges around incorporating ASAs into the totality-of-the-circumstances analysis.

A. *Algorithms as Police Profiles*

Traditional police profiles are "abstract[s] of characteristics thought typical of persons" engaged in certain criminal activity.[215] These characteristics often include traits or behavior that are legal and innocent when considered individually, but that become suspicious in a given context or when viewed together.[216] For instance, in *United States v. Sokolow*, the profile of a drug courier on an airplane included innocent facts such as: (1) paying for plane tickets in cash; (2) traveling under a name that does not match the name listed with

---

v. Ohio, 392 U.S. 1, 27 (1968). The analogy is unhelpful as an analytical tool, however, as the concept of a "hunch" as used by courts, tends to be little more than a talismanic signifier for situations where sufficient individualized suspicion does not exist. In this vein, it is instructive to note that Craig Lerner, the foremost legal scholar on "hunches," provides perhaps the most complete definition without referencing a single case. Craig S. Lerner, *An Introduction to Police Hunches*, 4 J.L. ECON. & POL'Y 1, 3-5 (2007).

211 *See* Tung Yin, *The Probative Values and Pitfalls of Drug Courier Profiles as Probabilistic Evidence*, 5 TEX. F. ON C.L. & C.R. 141, 144 (2000) (citing cases in which such profiles have been used); *see also* United States v. Ortiz, 714 F. Supp. 1569, 1570 (C.D. Cal. 1989) (applying a terrorist profile to identify a suspicious individual in an airport).

212 *See* Ferguson, *supra* note 20, at 308-10 (comparing "predictive policing" technologies, which predict where crime is likely to occur, with a police profile).

213 *Cf. id.* at 305-08 (comparing predictive policing to an informant's tip); Elizabeth E. Joh, *Policing by Numbers: Big Data and the Fourth Amendment*, 89 WASH. L. REV. 35, 56-57 (2014) (expanding somewhat upon on Ferguson's analysis and arguing that predictive software may be more objective than inferences drawn from informants' tips).

214 *See* Goldberg, *supra* note 46, at 791 n.9 (noting the similarity between drug dogs and "other machines used by the police or forensic scientists").

215 United States v. Mendenhall, 446 U.S. 544, 547 n.1 (1980).

216 *See* United States v. Sokolow, 490 U.S. 1, 9 (1989) (holding that DEA agents had a reasonable basis to stop a defendant who fit a drug courier profile because, even though each individual factor in the profile was "consistent with innocent travel," "taken together they amount to reasonable suspicion").

one's phone number; (3) traveling from a "source city" for drugs; (4) staying in the "source city" for a brief period, particularly when compared to the length of the flight to get there; (5) appearing nervous; and (6) not checking luggage.[217] Profiles like this one formalize the traditional policing process of examining an individual's noncriminal characteristics and actions to determine whether, when taken together, they create a suspicion of criminal conduct.[218] Taken in their best light, profiles consolidate and perpetuate the experience of numerous officers, thus allowing even junior officers to be smart at detecting crime, much in the way that police training instills the experience of veterans in new recruits.[219]

Though somewhat confusing, the Supreme Court's guidance on profiles suggests that they matter, but only indirectly, to the individualized suspicion analysis. First, the Court has been clear in saying that a profile *qua* profile does not justify an inference of individualized suspicion. In other words, a set of facts should receive neither any greater nor any lesser weight because those facts are contained in something that a certain law enforcement agency has called a profile of criminal activity.[220] Rather, a court must review de novo a police officer's determination that the facts contained in a police profile support the necessary individualized inference of suspicion.[221] However, to the extent that a profile is a distillation of police experience that the conjunction of certain facts is indicative of criminal conduct, that experience is entitled to some deference.[222]

---

217 *Id.* at 3.

218 *See id.* at 9-10 ("Indeed, *Terry* itself involved 'a series of acts, each of them perhaps innocent' if viewed separately, 'but which taken together warranted further investigation.'"(quoting Terry v. Ohio, 392 U.S. 1, 22 (1968))); *see also* Milton Heumann & Lance Cassak, *Profiles in Justice? Police Discretion, Symbolic Assailants, and Stereotyping*, 53 RUTGERS L. REV. 911, 918 (2001) ("[P]rofiling often comes to focus on behavior that is perfectly legal and in other contexts (perhaps even in the context at hand) purely innocent.").

219 Florida v. Royer, 460 U.S. 491, 525 n.6 (1983) (Rehnquist, J., dissenting). Criticism of police profiles is plentiful, however. *See* Mark J. Kadish, *The Drug Courier Profile: In Planes, Trains, and Automobiles; And Now in the Jury Box*, 46 AM. U. L. REV. 747, 751 n.7 (1997) (citing numerous scholars who criticize police use of drug courier profiles). *See generally* DAVID A. HARRIS, PROFILES IN INJUSTICE: WHY RACIAL PROFILING CANNOT WORK (2002).

220 *See Sokolow*, 490 U.S. at 10 ("We do not agree with respondent that our analysis is somehow changed by the agents' belief that his behavior was consistent with one of the DEA's 'drug courier profiles.'").

221 *See* Ornelas v. United States, 517 U.S. 690, 699 (1996) ("We therefore hold that as a general matter determinations of reasonable suspicion and probable cause should be reviewed *de novo* on appeal."); *see also, e.g.*, *Sokolow*, 490 U.S. at 9-10 (finding reasonable suspicion through independent analysis of the facts contained in drug courier profiles); United States v. Montoya de Hernandez, 473 U.S. 531, 541-42 (1985) (finding reasonable suspicion based on review of all facts when a customs agent stopped defendant based on behavior consistent with a drug courier profile); Reid v. Georgia, 448 U.S. 438, 440-41 (1980) (per curiam) (holding that facts that "appeared to the agent to fit the so-called 'drug courier profile'" were insufficient to create reasonable suspicion).

222 *See Ornelas*, 517 U.S. at 699 (recognizing that when "a police officer views the facts through the lens of his police experience and expertise," the inferences he draws "deserve deference"); United

Requiring courts to engage in de novo reviews of traditional profiles makes sense.[223] We expect police to examine the facts in an individual case and use typical tools of reason and logic—induction, deduction, and the like—to decide whether they suggest possible criminal activity. Yet police are "engaged in the often competitive enterprise of ferreting out crime."[224] Thus, police have some incentive to push boundaries, which in this case might mean constructing a "chameleon-like" profile that can fit any situation.[225] A judge, on the other hand, should have no skin in the game.[226] Thus, when presented with the same facts that were available to the police at the time, including background on the relevant officer's training and experience, the judge can double-check the officer's logic dispassionately. Ideally this process ensures that the police reasoned logically and reached a defensible conclusion that the facts supported a sufficient inference of criminal suspicion.[227] Similarly, the court can ensure

States v. Mendenhall, 446 U.S. 544, 564 (1980) (Powell, J., concurring) (noting, where a suspect allegedly fit a DEA profile, that "[i]n all situations the officer is entitled to assess the facts in light of his experience" (quoting United States v. Brignoni-Ponce, 422 U.S. 873, 885 (1975))).

223 It bears mention that this approach is effective only when lower courts engage in the analysis with the appropriate balance of deference and skepticism toward police decisions. However, experience shows that courts may have difficulty striking that balance. *See* David Cole, *Discretion and Discrimination Reconsidered: A Response to the New Criminal Justice Scholarship*, 87 GEO. L.J. 1059, 1077-79 (1999) (providing a catalogue of facts cited by courts as factors in DEA drug courier profiles at airports and concluding that "[s]uch a profile does not meaningfully narrow the field of potential suspects"); Sharon L. Davies, *Profiling Terror*, 1 OHIO ST. J. CRIM. L. 45, 60-61 (2003) ("Despite scholarly criticism, courts tended to uphold reliance on drug courier profiles prior to September 11 provided law enforcement agents who relied on those profiles did not consider an individual's race or ethnicity in isolation in calculating reasonable suspicion." (footnote omitted)); *see also* Charles L. Becton, *The Drug Courier Profile: "All Seems Infected That th' Infected Spy, As All Looks Yellow to the Jaundic'd Eye*," 65 N.C. L. REV. 417, 469 (1987) ("[L]ower courts have sanctioned profile stops with increasing regularity."). Similarly, courts must engage in a meaningful inquiry into an officer's training and experience to be able to properly assess its impact on the individualized suspicion analysis. Unfortunately, there is ample evidence that many courts fail to engage in such an inquiry. *See* Richardson, *supra* note 128, at 1158 ("[Courts] rarely engage in any serious attempt to think through what types of experience and training are significant in the reasonable suspicion context."). These concerns suggest there are serious issues with how the Court's approach has been implemented, but do not undercut the logic of its approach.

224 Johnson v. United States, 333 U.S. 10, 14 (1948).

225 *Sokolow*, 490 U.S. at 13 (Marshall, J., dissenting); *see also* United States v. Hooper, 935 F.2d 484, 499 (2d Cir. 1991) (Pratt, J., dissenting) (describing the drug courier profile as "laughable, because it is so fluid that it can be used to justify designating anyone a potential drug courier if the DEA agents so choose").

226 *See Johnson*, 333 U.S. at 14 (noting that the Fourth Amendment requires assessment of the facts by a "neutral and detached magistrate").

227 *Cf.* Becton, *supra* note 223, at 444 (noting that the inconsistencies between drug courier profiles "invite careful analysis of the[ir] purported logic").

that enough facts exist regarding the specific target of the search or seizure to support a finding of suspicion that is sufficiently individualized.[228]

On the surface, then, an ASA is like a police profile, in that it identifies likely criminals through the coexistence of multiple innocent facts gleaned from past experience. But substantial differences lie beneath this superficial analogy. First, ASAs derive their conclusions from hard data. In order to "learn" a correlation between certain conduct or characteristics and criminal activity, an ASA must process training data derived from real-life situations.[229] On the other hand, traditional profiles are frequently criticized for the absence of data demonstrating a person meeting the profile is likely engaged in criminal conduct.[230] Similarly, ASAs, like other machine learning algorithms, can examine exponentially more data points about a person or situation than could reasonably be listed in a traditional profile.[231]

Second, ASAs can identify more complex relationships between observable data and criminal activity than the simple checklist of a traditional profile, which is often applied without clear standards.[232] An ASA that applies even basic machine learning algorithms can not only check for the existence of particular facts, but also assign a weight to each fact depending on the strength of its correlation to criminal activity.[233] Likewise, an ASA can assess the interdependency of variables. For example, an ASA can determine the extent to which the occurrence of criminal activity depends not just on the existence of a single variable, but on the concurrent existence or non-existence of multiple variables.[234] Thinking back to the profile in *Sokolow*, an ASA might reveal that paying for a ticket in cash and not checking a bag do not, each standing alone, predict

228 *See* Morgan Cloud, *Search and Seizure by the Numbers: The Drug Courier Profile and Judicial Review of Investigative Formulas*, 65 B.U. L. REV. 843, 920 (1985) (concluding that mechanical application of drug courier profiles violates the Fourth Amendment's requirement of individualized suspicion).

229 *See supra* notes 66–69 and accompanying text. As will be discussed shortly, this reliance on hard data makes ensuring the quality of the underlying data even more important. *See infra* Part V.

230 *See* Heumann & Cassak, *supra* note 218, at 918 ("Less precise than arrest-based historical data, but also frequently mentioned as a basis for identification of profiling traits, are inferences or interpretations of facts drawn by the police officer's experience."); *see also* Cloud, *supra* note 228, at 920 (arguing that courts relying on profile characteristics "do[] so without first requiring the government to demonstrate that the profile characteristics actually identify criminals").

231 For example, an email spam filter may have a vocabulary of 10,000 terms that it uses to predict whether a given email is spam. FLACH, *supra* note 14, at 9.

232 Traditional profiles are often informal, unwritten, and do not state how many of a set list of factors must be met before the profile is satisfied. *See* Heumann & Cassak, *supra* note 218, at 919-21 (noting criticism of police profiles as overly malleable and loosely formulated).

233 *See* FLACH, *supra* note 14, at 25-32 (discussing the use of probability functions to determine the likelihood of an event occurring depending on the existence of each variable).

234 *See id.* at 44 ("One fascinating and multi-faceted aspect of features is that they may interact in various ways. Sometimes such interaction can be exploited, sometimes it can be ignored, and sometimes it poses a challenge.").

drug trafficking with any particular strength, but that the concurrence of the two factors is highly predictive. Consequently, an ASA's capacity for examining a multitude of variables and identifying complex relationships between variables means that the rules generated by an ASA may not be interpretable, even to the ASA's programmers.[235]

The differences between traditional profiles and ASAs make the Supreme Court's approach to traditional profiles unhelpful and counterproductive when applied to ASAs. To begin with, often no one will be able to explain to a reviewing court how or why the algorithm made its prediction. Thus, the court simply will be unable to double-check the ASA's work. Even when an ASA is programmed to be interpretable, the "logic" of an ASA is not of the sort that a judge can easily double-check.[236] One benefit of an ASA is its capacity to identify correlations within data that are not obvious but are statistically valid.[237] In other words, an ASA could identify a set of behaviors that correlate strongly to criminal conduct, even though the logical connection between the behavior and criminality—that is, why a criminal would engage in that behavior—is unclear to a human observer. The absence of a clear logical connection does not mean that the behavior is a bad predictor of criminality; rather, the logic explaining the correlation may be surprising, or the available dataset may fail to contain the information needed to understand it.[238] Yet a court treating the ASA like a traditional police-created profile, and therefore requiring a logical explanation for why certain facts predict criminality, might incorrectly reject the ASA's "illogical" prediction, notwithstanding the level of confidence the ASA has in the prediction.

In sum, courts treating ASAs like police profiles may demand that the ASAs be interpretable, thus undermining their effectiveness,[239] and may reject accurate predictions as "illogical." At the same time, the profile analysis would ignore the real sources of ASA inaccuracy, which typically occur in the training

235 *See supra* notes 99–102 and accompanying text (discussing the tradeoff between interpretable and black box ASAs).

236 The issue of whether the interpretability of ASAs should be mandated to facilitate understanding involves a number of policy questions that are outside the scope of this Article. *See, e.g.*, Zarsky, *supra* note 8, at 1526-30 (discussing transparency-related policy concerns such as the feasibility of regulating a process as dynamic as data mining).

237 *See, e.g.*, Colonna, *supra* note 52, at 313 (relating a "canonical anecdote" about a marketing manager for a supermarket who used data mining to discover a correlation between purchases of diapers and beer as "an example of unpredictable knowledge found in a huge dataset" (emphasis omitted)); Surden, *supra* note 29, at 107 ("Machine learning techniques are also useful for discovering hidden relationships in existing data that may otherwise be difficult to detect.").

238 *See, e.g.*, Surden, *supra* note 29, at 108-09 (suggesting that machine learning can be applied to assess legal opinions and unearth unarticulated bases for judicial decisions).

239 *See* Zarsky, *supra* note 8, at 1520 (discussing the functional costs of mandating ASA interpretability).

and programming of the algorithm.[240] Consequently, courts should look elsewhere to find useful analogies to ASAs.

B. *Algorithms as Informants*

Under the most general definition, an informant is a non-police-officer who provides information to the police.[241] Traditional informants include "jailhouse snitches, criminal accomplices, concerned citizens, and innocent eyewitnesses."[242] As a class, little ties the various categories of informants together beyond the fact that they are civilians, not trained law enforcement agents. Consequently, their information does not fall within the scope of the Fourth Amendment's collective knowledge doctrine.[243] For analytical purposes, informants can be subdivided into three categories: "(1) criminal . . . informants, (2) anonymous tipsters, and (3) citizen-informants."[244] Criminal informants generally provide police with information about their own criminal contacts in exchange for money or leniency.[245] Anonymous tipsters provide information to the police without disclosing any identifying information.[246] Citizen-informants are known civilians who provide police with information that they obtained by virtue of being the victim of or witness to a crime.[247] ASAs share characteristics with all three categories.

When assessing the weight to assign to a tip in the individualized suspicion analysis, courts look to the informant's veracity and reliability.[248] To assess the informant's veracity, courts will look to relevant data like her purported motivations for helping the police,[249] her previous history of providing accurate information,[250] and her reputation in the community.[251] The reliability of a tip traditionally

240 *See supra* notes 79–98 and accompanying text.

241 *See* Michael L. Rich, *Coerced Informants and Thirteenth Amendment Limitations on the Police-Informant Relationship*, 50 SANTA CLARA L. REV. 681, 689 (2010).

242 *Id.* at 689-90.

243 *See supra* Section III.A; *see also* United States v. Ventresca, 380 U.S. 102, 111 (1965) ("Observations of *fellow officers* of the Government engaged in a common investigation are plainly a reliable basis for a warrant applied for by one of their number." (emphasis added)).

244 Ariel C. Werner, *What's in a Name? Challenging the Citizen-Informant Doctrine*, 89 N.Y.U. L. REV. 2336, 2343 (2014).

245 *Id.* at 2343-44.

246 *Id.* at 2357.

247 *Id.* at 2341-44.

248 *See, e.g.*, Alabama v. White, 496 U.S. 325, 328-29 (1990); Illinois v. Gates, 462 U.S. 213, 230 (1983).

249 *See, e.g.*, *Gates*, 462 U.S. at 233-34; United States v. Angulo-Lopez, 791 F.2d 1394, 1397 (9th Cir. 1986) (discussing different motivations for providing a tip and how they affect the inference of trustworthiness).

250 *See, e.g.*, United States v. Bush, 647 F.2d 357, 362 (3d Cir. 1981) (noting that a track record of accurate tips is "the typical basis for a finding of veracity").

251 *See* Florida v. J. L., 529 U.S. 266, 270 (2000) (reasoning that an anonymous tip requires deeper analysis of veracity than a tip from a known informant where the informant's reputation can be assessed).

depends on whether the tip is based on personal knowledge or is otherwise so detailed that it is likely to have come from someone with first-hand information.[252] In addition, police corroboration of a tip's details substantially enhances the reliability of the tip,[253] with corroboration of predictions about future conduct seen as particularly valuable.[254]

Courts assess the veracity and reliability of criminal informants, anonymous tipsters, and citizen-informants very differently. They view criminal informants with skepticism because their criminal activities call their credibility into question.[255] Moreover, courts doubt the motivations of criminal informants, who often provide information in anticipation of receiving some benefit.[256] Additionally, since the source of an anonymous tip is, by definition, unknown to the police, an anonymous tip provides "virtually nothing" to suggest the tipster's honesty and gives "absolutely no indication" of the basis for the tipster's information.[257] Plus, anonymous tipsters are trusted less because they cannot be held responsible for fabricated allegations.[258] Given these concerns, an anonymous tip generally must be corroborated before it can be assigned much weight in the individualized suspicion analysis.[259] Corroboration of an anonymous tip's predictions about future conduct provides the most weight in establishing the tipster's credibility and basis of knowledge.[260]

Finally, tips from citizen-informants are generally accorded substantial weight. Many courts adhere to the "citizen-informant doctrine," by which information provided by "ordinary" citizens can be relied upon in the individualized suspicion analysis even without corroboration.[261] In particular,

---

252 *See, e.g.*, United States v. Nieman, 520 F.3d 834, 840 (8th Cir. 2008) (relying in part on the first-hand nature of informants' information to sustain a finding of probable cause).

253 *See Gates*, 462 U.S. at 241-44.

254 *Id.* at 245-46.

255 *See* Werner, *supra* note 244, at 2366 (describing the readiness of courts to vilify criminal informants); *see also* Max Minzner, *Putting Probability Back into Probable Cause*, 87 TEX. L. REV. 913, 942 (2009) ("[C]itizen informants are often assumed to be more reliable than those who have been involved in criminal activity.").

256 *See, e.g.*, State v. Smith, 867 S.W.2d 343, 347 (Tenn. Crim. App. 1993) ("Information supplied to officers by the traditional police informer is not given in the spirit of a concerned citizen, but often is given in exchange for some concession, payment, or simply out of revenge against the subject." (quoting State v. Paszek, 184 N.W.2d 836, 842 (Wis. 1971))).

257 *Gates*, 462 U.S. at 227.

258 *See* Florida v. J. L., 529 U.S. 266, 270 (2000).

259 *See Gates*, 462 U.S. at 227 ("The Illinois Supreme Court concluded—and we are inclined to agree—that, standing alone, the anonymous letter . . . would not provide the basis for a magistrate's determination that there was probable cause to believe contraband would be found . . . .").

260 *See J. L.*, 529 U.S. at 269 ("Anonymous tips . . . are generally less reliable than tips from known informants and can form the basis for reasonable suspicion only if accompanied by specific indicia of reliability, for example, the correct forecast of a subject's 'not easily predicted' movements.").

261 *See* Werner, *supra* note 244, at 2348-50, 2350 n.57.

courts are willing to trust non-criminal informants because they are presumed to have no reason to lie and can be punished if they falsely report a crime.[262]

ASAs fit the broad definition of informants, in that they are outside of law enforcement and provide information to police about criminal activity, but they are not neatly placed in any of the three traditional categories.[263] With respect to veracity, it is immediately obvious that ASAs are not people with personalities, codes of morality, motivations, or the capacity for honesty and dishonesty.[264] As such, a discussion of an ASA's veracity is nonsensical. But ASAs are programmed by people, and an ASA will inevitably reflect the conscious and unconscious biases and motivations of its programmers.[265] Such biases and motivations are generally irrelevant to the individualized suspicion analysis, unless they provide a reason for the police to doubt the underlying accuracy of the information.[266] In one sense, an ASA's programmers are like citizen-informants in that they are likely non-criminals, and thus should be trusted. Yet like anonymous informants, the programmers almost certainly will not be subject to criminal prosecution if an ASA's prediction is wrong, thus reducing the weight that should be given to their "tips" under the traditional informant analysis. Moreover, an ASA's creators, like many criminal informants, are motivated, at least in part, by money, thus perhaps calling their credibility into question.[267]

---

262 *Id.* at 2360 (describing how courts "tend to justify the citizen-informant doctrine" on the premise that "citizen-informants have no reason to lie" and that "if citizen-informants *did* have any reason to lie, they would be adequately deterred by the potential sanctions for falsely reporting a crime").

263 *See* Ferguson, *supra* note 20, at 305–10 (attempting, with limited success, to analogize "predictive policing" algorithms that identify where crime is likely to occur to an informant's tip).

264 *See id.* at 307 ("The computer has no biases, no past bad acts, and no agendas."). This is not to say that something meant to replicate human personality, morality, or honesty could not be programmed into an automated process. *See, e.g.*, George R. Lucas, Jr., *Automated Warfare*, 25 STAN. L. & POL'Y REV. 317, 322 (2014) (arguing that we could achieve "robot morality" by programming unmanned vehicles to follow the moral and legal demands of war as well or better than human beings). An ASA is unlikely to involve such programming, however, given its limited purpose.

265 *See* Danielle Keats Citron & Frank Pasquale, *The Scored Society: Due Process for Automated Predictions*, 89 WASH. L. REV. 1, 4 (2014) ("Because human beings program predictive algorithms, their biases and values are embedded into the software's instructions . . . ."). On this point, I differ with Andrew Guthrie Ferguson, who has said, in an analogous context, that "the computer algorithm presents none of the truth-related concerns that arise with a human informant. The computer computes what it computes, neither being true nor false." Ferguson, *supra* note 20, at 307 n.280. Ferguson's contention glosses over the important role that humans play in programming any algorithm and the potential that a programmer's biases or motives may shade the "truth" of what the algorithm computes.

266 *See* United States v. Perez, 651 F.2d 268, 271 (5th Cir. 1981) ("The motives upon which informants act in reporting crimes are generally irrelevant . . . ."); *cf.* Whren v. United States, 517 U.S. 806, 813 (1996) (holding that the subjective intentions of the police typically play no role in the probable cause analysis).

267 On the other hand, it stands to reason that a police department would be displeased with an ASA if courts found reliance on the ASA's predictions to be objectively unreasonable and

Looking to the motivations of an ASA's programmers to assess how much weight to put on an ASA's prediction is odd, however, given that data should be available about an ASA's proven or anticipated reliability. As noted, an ASA's prediction should come with a confidence level that can provide a tangible measure of the ASA's potential predictive power.[268] Moreover, data can be collected about the performance of the ASA that can then be used to establish its reliability.[269] This sort of data does have a place in the Court's informant analysis, in that tips from informants who have proven to be reliable in the past are given greater weight.[270] But the analysis that courts typically undertake when looking at the quality of an informant's past tips is not robust; instead, they often rely on a general assertion of an informant's reliability by a police affiant.[271] Moreover, the inferences that support giving weight to an informant who has proven reliable—that he has access to information about criminality and reports it truthfully—are quite different from what can be inferred from an algorithm that has proven to be statistically reliable. The former are based on human experience and thus are well within the expertise of police and judges.[272] The latter derive from complex statistical analyses that may not be interpretable by anyone,[273] much less magistrates or officers untrained in statistics.[274] Thus, even though there may be reasons based on a traditional analysis of human motivations to believe that an ASA is "credible" and "reliable," the traditional informant analysis is a poor fit for the statistical evidence that better substantiates a finding of the proper weight to give to an ASA's prediction.

---

excluded evidence obtained as the result of such reliance. *Cf.* United States v. Leon, 468 U.S. 897, 918-19 (1984) (discussing the exclusionary rule's deterrence of police misconduct). An ASA's programmers would therefore have a financial incentive to provide an ASA that is reliable enough for the police to rely on in finding individualized suspicion.

268 *See supra* note 121 and accompanying text.

269 *See* Joh, *supra* note 213, at 57 ("Software with a demonstrated history of successfully predicting high crime areas based on verifiable crime data is likely to be a highly persuasive factor in the reasonable suspicion formulation.").

270 *See* Adams v. Williams, 407 U.S. 143, 146 (1972) (noting that the tip of an informant who had provided an officer "with information in the past" made a stronger case for individualized suspicion than an anonymous tip); Draper v. United States, 358 U.S. 307, 313 (1959) (finding from the facts that the police properly relied on information from informant "whose information had always been found accurate and reliable").

271 *See* Goldberg, *supra* note 46, at 808 (describing how courts have not quantified what counts as a reliable track record, because often courts require only a general assertion by an affiant that the informant has supplied information leading to arrests).

272 Of course, the accuracy of these inferences could be tested empirically and analyzed statistically. The important point, though, is that judges rely on these inferences not because of empirical support, but rather because they comport with the judge's intuitions about how people behave.

273 *See supra* notes 99–102 and accompanying text.

274 *See generally* Joëlle Anne Moreno, *Beyond the Polemic Against Junk Science: Navigating the Oceans that Divide Science and Law with Justice Breyer at the Helm*, 81 B.U. L. REV. 1033 (2001) (discussing the challenges of expecting judges to engage in statistical analysis in the context of expert testimony).

The type of data that goes into an ASA's prediction also does not analogize well to the "basis of knowledge" analysis traditionally used for informant tips. Courts credit human informants whose tips are based on reliable information about a suspect's criminality.[275] ASAs claim no such inside information; rather, their "tips" are based on an enormous amount of past data about a large number of people and some quantity of data specifically about the suspect. The quality and quantity of this data is central to the weight that should be given to the ASA's prediction.[276] But the "basis of knowledge" analysis in the informant context provides no insight into how an ASA's prediction should be incorporated into the individualized suspicion analysis.

Finally, corroboration does not provide the same logical basis for believing in an ASA's accuracy that it does for an informant's tip. Corroboration of innocent facts in a human informant's tip, and specifically corroboration of the informant's predictions about the suspect's future behavior, is relevant because it suggests that the informant has some inside knowledge of the suspect's conduct and thus is more likely to be right about the suspect's illegal activities.[277] To the extent an officer is aware of the data that resulted in the ASA's tip, corroborating the accuracy of the data says very little about the accuracy of the ASA's prediction. Moreover, an ASA will not make any predictions about future behavior outside of the general prediction that the suspect is engaged in criminal conduct. Thus, to the extent there is anything to corroborate, it will not be particularly useful in ensuring the ASA's accuracy.

To conclude, analogizing an ASA's prediction to an informant's tip does not provide a useful analytical structure for courts. Human informants and computer algorithms are fundamentally different in ways that impact how we should assess the reliability of each source of information.

C. *Algorithms as Drug-Sniffing Dogs*

Like other machine learning algorithms, ASAs are likely to be structured as "black boxes" that take in data and spit out predictions, but whose inner workings are unknown and perhaps incomprehensible to humans.[278] Moreover, the tendency in criminal justice arenas toward secrecy in police investigative strategies

275 *See* Illinois v. Gates, 462 U.S. 213, 245 (1983) (holding that the corroboration of non-criminal details in an anonymous tip suggested that the information came from someone who "also had access to reliable information of the [defendants'] alleged illegal activities").

276 *See supra* notes 83–87 and accompanying text; *see also* Ferguson, *supra* note 20, at 317-18 (discussing concerns about the underlying data quality in the context of predictive policing).

277 *Gates*, 462 U.S. at 244.

278 *See supra* notes 99–102 and accompanying text.

suggests that ASAs are unlikely to be transparent or interpretable.[279] This lack of transparency differentiates suspicion algorithms from traditional algorithms that determine "historical facts," like DNA matching and blood-alcohol-level testing. The algorithms that underlie these determinations are relatively straightforward and explicable, and therefore they can (at least theoretically) be fully explored through expert testimony and cross-examination.[280] But even if the data used to train an ASA and the rules that the ASA creates were available to a defendant in a suppression hearing, the enormity of the data, the complexity of the rules, and the resource constraints would present formidable obstacles to a full consideration of how the ASA generated the prediction at issue.[281]

Drug-sniffing dogs are the prototypical black boxes in the individualized suspicion analysis. It is easy to understand that a dog has a heightened sense of smell and that drug-sniffing dogs are trained to recognize certain chemical compounds that are affiliated with illegal drugs.[282] But explaining how the input of the residue of an illegal drug is translated in a dog's brain into the output of an "alert" is beyond the scope of available expert testimony, in large part because "[t]he science of 'alerting' is not yet fully developed."[283] Thus, the drug dog's brain is like an ASA: we know the inputs, and we receive the outputs, but we cannot fully understand how the internal mechanism works.

The following discussion tests the validity of analogizing ASAs to drug dogs and is divided into two parts. The first outlines the approach taken by courts to drug dogs and some of the criticisms that have been leveled against it. The second discusses how the lessons from drug dogs should be applied to ASAs, in light of the similarities and differences between the two entities.

---

279 *See* Zarsky, *supra* note 8, at 1526-27 (describing government reluctance to provide transparency in law enforcement proceedings).

280 *See, e.g.*, Jay A. Zollinger, *Defense Access to State-Funded DNA Experts: Considerations of Due Process*, 85 CALIF. L. REV. 1803, 1810-13 (1997) (discussing different states' conceptions of the role of an expert in addressing DNA analysis technology).

281 *See* Kerr, *supra* note 42, at 875-76 (discussing how courts struggle to fully understand new technologies, and how this can cause misunderstandings and errors in judicial rulemaking).

282 *See* Robyn Burrows, *Judicial Confusion and the Digital Drug Dog Sniff: Pragmatic Solutions Permitting Warrantless Hashing of Known Illegal Files*, 19 GEO. MASON L. REV. 255, 280 (2011) ("[A] layman can understand that dogs have a heightened sense of smell and that the dog is trained to detect specific substances."); Myers, *supra* note 46, at 3 ("Researchers at Auburn University studying dogs' capacity to identify certain smells have found that some dogs can detect odors when the particles in the air are at a concentration of 500 ppt—that's parts per *trillion*.").

283 Myers, *supra* note 46, at 4.

1. The Law of Drug Dogs

The Supreme Court first touched on the role of drug dogs in the individualized suspicion analysis in *United States v. Place*.[284] There, the Court dubbed a canine sniff "*sui generis*" because, at least in an ideal world, it discloses only the presence of contraband, and concluded therefore that it is not a Fourth Amendment search.[285] The Court did not discuss the weight to be given to a drug dog alert in establishing probable cause or reasonable suspicion, though it did repeatedly refer to "trained" and "well-trained" canines.[286] In *Illinois v. Caballes*, a drug dog alerted on the trunk of a vehicle and "[b]ased on that alert," police searched the trunk.[287] The only question before the Court was whether the Fourth Amendment requires any individualized suspicion before a canine sniff for drugs is permitted.[288] The Court said nothing about what weight to give a dog's alert in the individualized suspicion analysis, other than to note that the trial court found the drug dog's alert to be sufficiently reliable to create probable cause.[289] Once again, the Court referred to "well-trained" drug dogs in formulating its holding, without explaining what might actually make a dog "well-trained."[290]

Dissenting in *Caballes*, Justice Souter pointed to substantial data suggesting that "[t]he infallible dog . . . is a creature of legal fiction."[291] Souter's main argument was that a dog sniff should be treated as a Fourth Amendment search, but he conceded that even a fallible dog alert can create reasonable suspicion or probable cause because "the Fourth Amendment does not demand certainty of success to justify a search for evidence or contraband."[292] Nonetheless, he recognized that "sniffing averages" differ from dog to dog.[293]

In light of *Place* and *Caballes*, federal circuit courts uniformly permitted drug dog alerts to establish individualized suspicion and thus to justify searches and seizures.[294] But before finding that a dog's alert creates probable cause, lower courts often require some evidence of the dog's reliability, typically in the form

---

284 462 U.S. 696 (1983).

285 *Id.* at 707.

286 *See, e.g.*, *id.* at 705-06 ("Moreover, the police may confine their investigation to an on-the-spot inquiry—for example, immediate exposure of the luggage to a trained narcotics detection dog . . . .").

287 543 U.S. 405, 406 (2005).

288 *Id.* at 407.

289 *Id.* at 409.

290 *Id.* at 409 ("Accordingly, the use of a well-trained narcotics-detection dog . . . during a traffic stop generally does not implicate legitimate privacy interests.").

291 *Id.* at 411 (Souter, J., dissenting).

292 *Id.* at 413.

293 *Id.* at 411-12 (collecting cases dealing with potential false positives from drug dog alerts).

294 *See* Myers, *supra* note 46, at 18 n.88 (providing a list of cases in each of the federal circuits that have concluded that an alert by a trained detector dog, alone, constitutes sufficient probable cause for a search).

of training records or certifications.[295] The quantum of evidence of a drug dog's reliability that courts will require before finding the alert sufficient to establish probable cause is unclear, and critics fear that in many cases it is too low.[296] Moreover, there are no uniform standards for the certification of drug dog reliability.[297]

The Supreme Court recently tackled the question of how much weight should be attributed to a drug dog's alert in the individualized suspicion analysis in *Florida v. Harris*.[298] There, police searched the defendant's truck based on a drug dog's alert, and the State put on evidence of both the dog's and its handler's certification and training.[299] The Florida Supreme Court found the State's evidence insufficient to establish the reliability of the alert, and instead required the State to present a wider array of evidence, including training and certification records, field performance records, evidence of the handler's training and experience, and any other objective evidence of the dog's reliability known to the officer.[300]

The Supreme Court unanimously reversed the state court's decision, holding that the "strict evidentiary checklist" created by the Florida Supreme Court ran contrary to the totality-of-the-circumstances approach required by the Fourth Amendment.[301] The Court further opined that courts should place much greater weight on a dog's training and certification records than on its field performance.[302] And the Court noted two problems with reliance on field data.[303] First, field data cannot accurately establish how often a dog fails to detect drugs, because police typically will not search a car if a dog fails to alert on the vehicle.[304] Second, the Court was concerned that a dog may alert on quantities of drugs that are otherwise undetectable, either because they are too well hidden or too small for police to find.[305]

---

295 *See, e.g.*, United States v. Sundby, 186 F.3d 873, 876 (8th Cir. 1999) ("To establish the dog's reliability, the affidavit need only state the dog has been trained and certified to detect drugs."); United States v. Diaz, 25 F.3d 392, 394 (6th Cir. 1994) ("For a positive dog reaction to support a determination of probable cause, the training and reliability of the dog must be established."). *But see* United States v. Williams, 69 F.3d 27, 28 (5th Cir. 1995) ("Because a showing of the dog's reliability is unnecessary with regard to obtaining a search warrant, *a fortiori*, a showing of the dog's reliability is not required if probable cause is developed on site as a result of a dog sniff of a vehicle.").

296 *See* Myers, *supra* note 46, at 19-24 (exploring court methods used to establish a dog's reliability).

297 *See id.* at 27; *see also, e.g.*, Robert C. Bird, *An Examination of the Training and Reliability of the Narcotics Detection Dog*, 85 KY. L.J. 405, 410–15 (1996-97) (highlighting different training and certification techniques used by the Rhode Island State Police and the United States Customs Service).

298 133 S. Ct. 1050 (2013).

299 *Id.* at 1054.

300 *Id.* at 1055 (quoting Harris v. State, 71 So. 3d 756, 775 (Fla. 2011)).

301 *Id.* at 1056.

302 *Id.* at 1056-57.

303 *Id.* at 1056.

304 *Id.*

305 *Id.*

According to the Court, such instances are not errors by the dog, though labelled as such, and thus field data may overstate the number of false positives.[306] On the other hand, a controlled training environment ensures that false negatives and false positives are accurately recorded because the dog's trainers know if drugs are present.[307] The Court thus concluded "[i]f a bona fide organization has certified a dog after testing his reliability in a controlled setting, a court can presume (subject to any conflicting evidence offered) that the dog's alert provides probable cause to search."[308] A training program that evaluates a dog's proficiency in finding drugs can also establish a dog's reliability.[309]

Harris has been justifiably criticized for a number of reasons. First, and fundamentally, *Harris* perpetuates a problem that has plagued the Court's consideration of drug dogs from the beginning: the overvaluing of one piece of data—here, the alert of a trained drug dog—in the totality-of-the-circumstances analysis.[310] *Harris* dictates that the reliability of a dog's alert should be the central (and perhaps sole) focus of a court's analysis:

> If the State has produced proof from controlled settings that a dog performs reliably in detecting drugs, and the defendant has not contested that showing, then the court should find probable cause. If, in contrast, the defendant has challenged the State's case (by disputing the reliability of the dog overall or of a particular alert), then the court should weigh the competing *evidence* . . . . The question . . . is whether all the facts surrounding a dog's alert, viewed through the lens of common sense, would make a reasonably prudent person think that a search would reveal contraband or evidence of a crime.[311]

The problem with this approach is that by focusing on drug dog reliability, courts will likely undervalue other information. Specifically, courts should consider the "prior odds" that the suspect possessed drugs, that is, the reasonable likelihood that the suspect possessed drugs before the dog alerted.[312] From a statistical standpoint, one can use an equation called Bayes' Theorem to predict, based on the prior odds and a drug dog's rates of false positives and true

306 *Id.* at 1056-57.

307 *Id.* at 1057.

308 *Id.*

309 *Id.*

310 For comprehensive statistical explanations of this concern see Goldberg, *supra* note 46, at 817-18, and Myers, *supra* note 46, at 12-18.

311 *Harris*, 133 S. Ct. at 1058.

312 In this context, these prior odds can be generated by looking at the frequency of the targeted criminal conduct in the general population or considering specific facts about the defendant. *See* Goldberg, *supra* note 46, at 819.

positives,[313] the likelihood that the dog's alert was correct in a given case.[314] According to Bayes' Theorem, even an alert from a very reliable dog does not necessarily indicate a strong likelihood that drugs will be found. For instance, if a drug dog has a false positive rate of 5% and a true positive rate of 90%, the prior odds of finding drugs must have been at least 5% for the dog's alert to make it more likely than not that drugs would be found.[315] While the Court's call in *Harris* to look at "all the facts surrounding a dog's alert"[316] could be read to include the prior odds, the fact that the *Harris* court did not itself look at any facts other than the dog's reliability make such an examination unlikely.[317]

Of course, the *Harris* Court's failure to recognize the importance of prior odds does not preclude courts from considering them, but additional obstacles stand in the way of courts wishing to do so. First, it is somewhat counterintuitive that an alert from a dog that is 95% accurate does not create a 95% likelihood that drugs will be found.[318] Thus, courts are unlikely, without persuasion, to look beyond a dog's accuracy.[319] Second, judges and police are not trained statisticians, and therefore they may be incapable, at least without additional training, of accurately incorporating prior odds into their individualized suspicion analysis.[320] Third, prior odds are often unavailable.[321] Fourth, unless they apply statistical formulae like Bayes' Theorem, police, magistrates, or courts who receive numerical probability data, like a dog's accuracy rate, are likely to give it undue

313 False positives are instances in which the dog alerts when there are no drugs present, and true positives are instances in which the dog alerts when drugs are present. In other words, the false-positive rate describes how frequently the dog alerts when it should not, and the true-positive rate describes how often the dog alerts when it should.

314 Indeed, the totality-of-the-circumstances test for individualized suspicion essentially asks courts to engage in a Bayesian analysis of all new evidence. *See* Minzner, *supra* note 255, at 920 n.32 (describing how courts could use Bayes' Theorem to describe the probability of a location containing contraband, given certain evidence).

315 *See id.* at 950 n.181 (explaining that just because a dog with a 95% accuracy rate "has alerted on a particular location does not mean that location will contain contraband 95% of the time").

316 133 S. Ct. at 1058.

317 For instance, in *Harris*, when the defendant was pulled over for an expired license plate, he was visibly nervous, unable to sit still, shaking, breathing rapidly, and had an open can of beer in his cup holder. *Id.* at 1053. It is possible that these background facts established prior odds of drug possession sufficiently high for the dog's alert to establish probable cause. But, that conclusion is not obvious, and the Court did not ask the question. *See* Minzner, *supra* note 255, at 950 n.181 (noting the differences in the background facts, and likely prior odds, in *Place* and *Caballes*).

318 *See* Myers, *supra* note 46, at 13 (describing the belief that the dog's success rate is equal to the chance that the particular vehicle contained a controlled substance as "a widely held and intuitive misconception").

319 This likely explains why, even before *Harris*, lower courts accepted a drug dog's alert as sufficient to establish probable cause. *See supra* text accompanying note 294.

320 *See* Minzner, *supra* note 255, at 952-55 (discussing "capacity objections" to the use of statistical evidence, which are based on the inability of decisionmakers to accurately consider such evidence).

321 *See* Taslitz, *supra* note 126, at 862 (noting that generalized, objective probability data regarding crime rates is rarely available).

weight because of the "anchoring effect."[322] For instance, a court may use a drug dog's 95% hit rate as a starting point in the probable cause analysis and adjust up or down from there, without understanding that because of low prior odds, the actual likelihood of finding contraband was far lower.

In addition to this fundamental logical flaw, the *Harris* Court also relies on two questionable factual premises: first, that training and certification programs are strong evidence of a dog's reliability, and second, that field performance is weak evidence of reliability.[323] The first claim is faulty for two reasons. First, there are no accepted standards for dog training.[324] Dogs are trained according to standards articulated by law enforcement agencies,[325] or they may be certified by private organizations.[326] The level of reliability required for a dog to achieve certification varies substantially depending on the regimen. The most stringent require 100% accuracy, while some certify police dogs that are reliable in controlled testing environments only 70% of the time.[327] Of course, a certification is only as strong as the underlying testing standards, and the absence of such standards should give courts pause as they assess even a certified dog's reliability.

The second problem with giving great weight to training certification is that a dog's reliability in a controlled testing environment fails to account for circumstances in the real world that cause a drug dog to alert falsely. Conscious or unconscious cues from a dog's handler may cause the dog to alert falsely.[328] Furthermore, because of poor training or temperament, a drug dog may get distracted by chaotic, real-world circumstances.[329] A dog can be trained to ignore distractions,[330] and a handler can be trained not to cue his dog, but a training certification does not provide any information about how well the dog and handler apply that training in practice.

Meanwhile, the Court's second premise—that field performance is a poor indicator of reliability—is based on a misunderstanding of the probable cause

---

322 The anchoring effect is the human tendency to grab hold of any available number as the starting point of an estimation in the face of uncertainty. *See* Cass R. Sunstein, *Hazardous Heuristics*, 70 U. CHI. L. REV. 751, 752 (2003) ("[I]n the face of uncertainty, estimates are often made from an initial value, or "anchor," which is then adjusted to produce a final answer.").

323 133 S. Ct. at 1056-57.

324 *See* Taylor Phipps, *Probable Cause on a Leash*, 23 B.U. PUB. INT. L.J. 57, 77-79 (2014) (providing examples of the "drastic[]" variation among training and certification programs).

325 Bird, *supra* note 297, at 420-21.

326 Phipps, *supra* note 324, at 78.

327 *Id.* at 78-79.

328 Myers, *supra* note 46, at 22-24.

329 *Id.* at 4.

330 *See* Bird, *supra* note 297, at 413-14 (discussing dog training exercises that train dogs to work under adverse conditions).

requirement. Specifically, the Court argues that a dog's alert on the residual odor of drugs when no contraband is actually present is not a false positive that undermines the dog's reliability.[331] Yet, "probable cause to search requires an assessment of the odds that evidence is *currently* located in the place to be searched, not that it was there at some indeterminate time in the past."[332] Moreover, while the Court is right that field records cannot accurately record when a dog fails to find drugs that are present, such false negatives are not what matter when it comes to individualized suspicion.[333] The Fourth Amendment protects individuals against unreasonable searches and seizures, not from searches and seizures that did not happen. Thus, field records are highly relevant evidence of what matters in terms of a dog's reliability: how often a drug dog's alert accurately predicts the presence of contraband.[334]

Finally, critics point out that even though *Harris* requires courts to consider evidence that might undermine a drug dog's reliability,[335] the defense is unlikely to have access to such information.[336] Police agencies often do not keep detailed records of drug dog performance in the field.[337] Even when they do, such records and other relevant evidence, like training details, are generally in the hands of the government, and defendants are likely to have a difficult time obtaining them in discovery.[338] Thus, *Harris* may, for all practical purposes, create a bright-line rule that a drug dog's alert creates probable cause.[339]

2. ASAs as Drug Dogs

As just explained, courts trying to decide whether a drug dog's alert created individualized suspicion are instructed to look at a drug dog's training reliability, any certifications, its field performance, and any facts about the

---

331 Florida v. Harris, 133 S. Ct. 1050, 1056-57 (2013).

332 Kinports, *supra* note 47, at 68; *see also* Myers, *supra* note 46, at 22 ("Perversely, the better the dog is at detecting trace amounts of the desired substance, the higher the likelihood that the dog will alert on trace amounts that are inadvertently present in materials owned by the innocent."); Phipps, *supra* note 324, at 77 ("Although detecting residual odors from weeks prior may seem like a valuable trait in a dog, it actually demonstrates that the dog is less reliable at discerning whether drugs are actually present.").

333 *See* Myers, *supra* note 46, at 15 ("For our purposes, the important number is the false positives. What we want to know is the probability the car contains drugs conditional on (or in light of) the dog alert.").

334 *See* Phipps, *supra* note 324, at 73 (collecting studies documenting drug dog accuracy in the field).

335 133 S. Ct. at 1057.

336 *See* Kinports, *supra* note 47, at 65 ("Details about training programs the dog and its handler completed are in the hands of the government, and a defendant who was not on the scene during the dog sniff cannot know whether the dog was cued by its handler or working under 'unfamiliar conditions.'").

337 *Id.*

338 *Id.* at 65-66.

339 *Id.* at 65.

specific alert.[340] Analogizing an ASA to a drug dog, a court determining whether an ASA's prediction created individualized suspicion would look at the strength of the ASA's prediction, the confidence level of that prediction as established in initial programming of the ASA, the ASA's field performance, and any specific facts about the prediction to determine if sufficient individualized suspicion existed.

This analogy brings good news and bad news. The good news is that treating an ASA like a drug dog requires courts and police to ignore what they are ill-equipped to evaluate. Courts are not expected to directly examine for soundness the biological processes in a drug dog's brain by which drug residues inhaled by the dog result in an alert.[341] Similarly, courts would not be expected to directly examine how the ASA processes the information it receives in order to create its predictions of criminality. This is good news because, as explained earlier, the most effective ASAs are likely to operate in a way that is not comprehensible even to the people who programmed the algorithm.[342] Instead, courts considering a drug dog's alert assign weight to the alert based on the quality of the inputs and the outputs.[343] Courts treating ASAs like drug dogs would do the same.

The bad news is that the flaws in the drug dog analysis may work even greater mischief if that analysis is applied to ASAs. First, the problem of prior odds remains, but it is more complicated with respect to ASAs. The prior odds are important to the individualized suspicion analysis for a drug dog alert because the facts that go into calculating those odds are not incorporated into determining the reliability of the dog's alert.[344] In other words, the facts underlying the prior odds are independent of the reliability of the dog's alert. For instance, the fact that the police stopped Harris for driving with an expired license plate is independent from the dog's reliability,[345] because the reliability of a dog's alert generally will not be calculated based on the reasons behind a given stop. Thus, the reason for the stop, since it does not go into the reliability analysis, should be included in the prior odds calculation. For instance, a court should consider how much, if at all, the fact that a vehicle was stopped for an

340 *Harris*, 133 S. Ct. at 1057-58.

341 *Id.*

342 Cheng, *supra* note 100, at 548.

343 *Harris*, 133 S. Ct. at 1057-58.

344 *See* Minzner, *supra* note 255, at 921 (recognizing that for new evidence to impact the accuracy of a probable cause determination, the new evidence must not have already been considered in the pre-existing probable cause analysis).

345 *See Harris*, 133 S. Ct. at 1053, 1059 (acknowledging that a police officer stopped Harris because of his expired license plate, but holding that a drug dog's detection of drugs was reliable based on training records).

expired license plate increases or decreases the likelihood that drugs would be found in the vehicle. On the other hand, if we know the reliability of a dog's alert when that dog alerts on a truck rather than a car, then the fact that Harris was driving a truck is not independent of the dog's reliability and should not be incorporated into the prior odds calculation.[346]

It is easy to identify the evidence that is independent of a dog's reliability, because that reliability is likely going to be described in basic terms like, "Bobo correctly alerted 71% of the time."[347] Thus, in most cases, any facts specific to an individual alert should contribute to the calculation of prior odds. On the other hand, the strength of an ASA's prediction may be based on a substantial and unknown network of facts.[348] Imagine, for instance, that an officer receives a prediction from an ASA that, given the facts analyzed by the algorithm, there is a 62% likelihood that a suspect on a street corner is engaged in drug dealing. When the officer approaches the suspect, she will necessarily learn facts about the suspect. Perhaps she will know that the suspect is a local pastor, or she may observe the color of the clothing he is wearing. She must then decide whether the totality of the circumstances creates individualized suspicion that permits her to seize the suspect. But, the officer should consider the facts that she observed only if these facts were not included in the ASA's calculation of the strength of its prediction. Otherwise, those facts will be double counted.[349]

However, it will be difficult, or even impossible, for the officer to know whether the facts she learned were already considered by the ASA. This is because ASAs can process massive amounts of data,[350] and police may not even know what kind of data is input into an ASA. For instance, is clothing color a feature considered by the ASA? If not, the officer should incorporate her observations of clothing color in her analysis; if so, she should not because it has already been considered. This problem is at least theoretically solvable by ensuring that police agencies and individual officers are informed about the types of information used by an ASA.

The officer also may not be able to determine which inputs the ASA relied on because there may be no way for an officer to know what information the ASA actually obtained about this specific suspect in making its prediction.

346 *See* Minzner, *supra* note 255, at 950 n.181 (explaining how the prior odds affect the calculation of the success rate for drug dogs).

347 United States v. Kennedy, 131 F.3d 1371, 1378 (10th Cir. 1997).

348 Note that the reliability of the drug dog is analogous to the numerical certainty expressed by the ASA in its prediction of criminality.

349 *See* Goldberg, *supra* note 46, at 833 (discussing the problem of double counting in the context of a facial recognition device).

350 *See* Hu, *supra* note 7, at 803-04 (discussing the rise of mass data collection that "may facilitate the digital construction of . . . data patterns and data analyses").

For example, even if an individual's occupation is a feature in the ASA's model, did the ASA make a positive identification of the suspect and "know" that he was a pastor before making its prediction? Again, if not, the officer should incorporate that information into her analysis; if so, she should ignore it because it has already been considered. Whether the officer can access this information will depend on the extent to which the ASA is interpretable. Without interpretability, it will be impossible for the officer to make an accurate assessment of individualized suspicion. Moreover, even if this information is available, it must be constantly communicated and updated to officers acting on the prediction.

Additionally, relying only on training performance and certification to assess the reliability of ASAs is problematic for several reasons. The first is that the programming of ASAs is far more complex than the training of drug dogs. Teams of programmers will inevitably make hundreds or thousands of decisions throughout the programming process, and each of these decisions may create errors in the ASA.[351] This complexity ratchets up the importance of robust, meaningful standards for the creation and training of ASAs that minimize error and maximize effectiveness. Only with such standards can certification of an ASA provide substantial guarantees of accuracy in the individualized suspicion analysis. Unfortunately, the lack of concern that the Court showed in *Harris* about certification standards provides little reason to believe that courts would require more from ASAs.[352]

Even with robust certification procedures, field performance data are crucial in the evaluation of an ASA's reliability. A dog sniff is a straightforward process: a dog sniffs air to determine whether it contains trace amounts of narcotics. Though the circumstances in which a drug dog and handler seek to achieve this goal can vary substantially, the input and task remain essentially the same.[353] The input and task of an ASA, on the other hand, vary substantially depending on the circumstances and over time. The same crime may be committed in different ways in different places. Thus, the data used to train an ASA may

351 *See supra* notes 90–98 and accompanying text (discussing human sources of error in machine learning algorithms).

352 *See* Florida v. Harris, 133 S. Ct. 1050, 1054-1056 (2013) (finding a dog's training sufficient based on its quantity and reports that the dog performed "really good" and "satisfactorily" in training).

353 In fact, criminals try to make the drug dog's job more difficult by, for example, masking the odor of illegal drugs with some other substance, such as talcum powder or perfume. *See* David S. Rudstein, *"Touchy" "Feely"—Is There a Constitutional Difference? The Constitutionality of "Prepping" a Passenger's Luggage for a Human or Canine Sniff After* Bond v. United States, 70 U. CIN. L. REV. 191, 200-05 (2001) (describing methods used by law enforcement officers to detect passengers who attempt to mask the scent of drugs in their luggage). If criminal avoidance methods do change substantially over time, that presents yet another argument for the use of field performance records in the drug dog context.

not lead to reliable results in all places.[354] More importantly, how crimes are committed changes over time, particularly in response to law enforcement activities.[355] Consequently, even with robust, accurate, and representative training data, changes in crime patterns over time will inevitably lead to less reliable predictions. This diminishing reliability can only be captured through field performance data. Therefore, even crediting the *Harris* Court's critique of field performance data in the dog sniff context, courts and police must consider field performance data when assessing the reliability of an ASA's prediction.[356]

Finally, just as the state controls the data needed to undermine a drug dog's reliability, the government also likely will possess the information a defendant would need to challenge an ASA's prediction in court. In particular, a defendant would want information about (1) an ASA's reliability in general, including any certification it received, its training performance, and its field performance; and (2) the ASA's application to the defendant's case, including the facts that the ASA incorporated into its analysis, and how those facts were used. Access to all of this information would be necessary to ensure the most robust individualized suspicion counterargument.

Yet prosecutors are likely to resist the disclosure of information about how an ASA works on the grounds that such information may be used by criminals to "game the system" or avoid engaging in the forms of conduct that the ASA uses as a proxy for criminality.[357] This argument has substantial currency under

---

354 This is an example of an error that would arise from training data being insufficiently representative of real-world situations. *See supra* notes 88–89 and accompanying text. But even if training data is representative, an ASA may be particularly unreliable in geographic locations with peculiar crime patterns. *See* Adam Benforado, *The Geography of Criminal Law*, 31 CARDOZO L. REV. 823, 837-45 (2010) (discussing how characteristics of the physical environment influence an individual's decision of whether and how to commit a crime).

355 *See* William J. Stuntz, *Race, Class, and Drugs*, 98 COLUM. L. REV. 1795, 1804 (1998) (noting, with respect to illegal drug sales, that "as soon as law enforcement agencies adapt to a particular distribution pattern, the sellers have an incentive to change the pattern. The result is a cat-and-mouse game, with different forms of a given drug sold by different actors in different ways at different times and places").

356 Moreover, assessing the reliability of an ASA over time is the only way to ensure that the ASA continues to learn from new data and improve.

357 *See* Zarsky, *supra* note 8, at 1554 ("[K]nowledge of the inner workings of the automated prediction models in the hands of adversaries will allow them to 'game the system.'"). Prosecutors also may argue that information about ASAs is a trade secret that must be protected from disclosure. *See* Danielle Keats Citron, *Technological Due Process*, 85 WASH. U. L. REV. 1249, 1291-93 (2008) (discussing the use of the trade secret argument to prevent disclosure by government attorneys in other contexts); *see also* David S. Levine, *Secrecy and Unaccountability: Trade Secrets in Our Public Infrastructure*, 59 FLA. L. REV. 135, 138-40 (2007) (discussing trade secrecy relating to governmental function, such as voting machines); *Id.* at 171-72 ("While many people may not give [government secrecy] much thought, it is difficult to ignore such concerns because we interact with this infrastructure—roads, the Internet, governmental actions like law enforcement—on a daily basis."). The arguments for and against this position are beyond the scope of this Article.

the law,[358] and is forceful in some contexts. For example, information about the specific facts relevant to an ASA prediction of criminality and data about how the ASA weighs those facts would be most useful to future criminals. However, this information would also be the most detrimental to law enforcement because criminals could use it to change their behavior in targeted ways to avoid detection. Thus, information regarding specific relevant facts and their weight is the most deserving of protection on this ground.

However, the fear of enabling future criminals to "game the system" does not apply with the same force to all kinds of information that a defendant may want. Generalized information about an ASA's certification and reliability, for instance, reveals nothing about how the ASA works and provides no guidance for future criminals seeking to avoid detection. Therefore, because these specific pieces of information do not facilitate "gaming" of the system, this information should be provided to defendants. Similarly, knowing the types of information that an ASA incorporates into its analysis may be of limited utility to future criminals, but without knowing more about which facts matter and how they matter, criminals trying to alter their behavior to escape suspicion will be stumbling in the dark. Thus, rather than rejecting all discovery requests by defendants seeking to challenge the validity of a search or seizure based on an ASA's prediction, courts should carefully consider what information can be shared without adversely impacting law enforcement interests.[359]

3. Conclusion

A number of lessons emerge from the application of the analogy of drug dogs to ASAs under current Fourth Amendment jurisprudence. First, courts must recognize that an ASA's prediction, like any prediction of criminality, is only a part of the totality-of-the-circumstances analysis, and litigants must be prepared to educate courts about the importance of facts other than an ASA's numerical prediction in determining the existence of individualized suspicion. Second, interpretability of ASAs remains a central issue. While a "black box" ASA is more likely to provide accurate predictions, information about how an ASA works is necessary for the most accurate and complete Fourth Amendment

358 *See* Zarsky, *supra* note 8, at 1553-54 ("This powerful rationale [that transparency would enable avoidance] is reflected in current law. Every disclosure law has a law-enforcement exemption clause.").

359 *See* Zarsky, *supra* note 8, at 1555 ("Yet perhaps the most salient context for this pro-opacity argument is elsewhere in the 'usage' stage—at the point at which the government uses a mix of criteria, factors, behaviors, and attributes as proxies to identify wrongdoings. Here, the opacity argument is perhaps most intuitive—if government discloses the lists of proxies used, adversaries will simply avoid these proxies. They will, however, still engage in unlawful conduct. Therefore, providing information regarding these steps of the process should be prohibited." (footnote omitted)).

analysis. Experts in relevant fields, like machine learning, law, and law enforcement, should come together to consider how to balance these concerns most effectively. Third, as ASAs become more widespread, these same subject-matter experts must work together to propagate standards for the development of accurate and effective ASAs.

In addition, courts must require ASAs to be certified in accordance with these standards before an ASA's prediction can be used to establish individualized suspicion,[360] or at least weigh the absence of a certification heavily in the individualized suspicion analysis. Fourth, police agencies must maintain data about the performance of ASAs in the field,[361] and ASA standards should mandate ASAs be programmed to make the collection of such data straightforward. Finally, on a case-by-case level, defendants must be prepared to argue for full disclosure of training and field performance data for ASAs, as well as discovery into the kinds of data that the ASA uses. Courts, in turn, must be willing to order disclosure of this information despite prosecution arguments to the contrary.[362]

V. ASA ERRORS

The previous Parts establish that ASAs do not replace the role of people in making the determination of whether probable cause or reasonable suspicion exists, but that they can be considered, with some caveats, to be similar to drug dogs in the totality-of-the-circumstances analysis. A question remains, however: How should ASA errors be handled under the Fourth Amendment?

Before answering this question it is important, as a preliminary matter, to define what is and is not an ASA error. Specifically, an ASA is not "wrong" every time an officer searches or seizes a suspect in reliance on an ASA's prediction and finds no evidence of criminal conduct. Like any predictive machine learning algorithm, an ASA can make only probabilistic predictions.[363] Similarly, probable cause and reasonable suspicion are themselves probabilistic

360 Given the Court's aversion to creating bright-line rules, requiring certification may be implausible without legislation. *See* Florida v. Harris, 133 S. Ct. 1050, 1055-56 (2013) ("We have rejected rigid rules, bright-line tests, and mechanistic inquiries in favor of a more flexible, all-things-considered approach.").

361 This same recommendation has been made with respect to drug dogs. *See* Myers, *supra* note 46, at 33 (urging courts to mandate data collection on the use of search dogs and their accuracy rates in the field given the government's use of these dogs to override Fourth Amendment rights).

362 *But cf.* Erin Murphy, *The Politics of Privacy in the Criminal Justice System: Information Disclosure, the Fourth Amendment, and Statutory Law Enforcement Exemptions*, 111 MICH. L. REV. 485, 503 (2013) (recognizing the power of law enforcement interests and the wide array of statutory exemptions that accommodate these interests).

363 Remember that machine learning algorithms inevitably learn approximations of complex underlying phenomena (like whatever causes individuals to commit crimes). *See supra* notes 80–82 and accompanying text. Thus, errors, in the sense of false positives in particular instances, are inevitable.

predictions,[364] and a search or seizure based on the existence of probable cause that does not ultimately lead to the discovery of criminal conduct is not necessarily a Fourth Amendment violation.[365] Rather, an ASA is wrong when it provides a prediction of criminality that it should not have provided.[366] Such error can arise in two general ways: first, the ASA could be working with inaccurate data;[367] or second, the ASA's error could arise from human error during the programming process.[368]

Moreover, for current purposes, we care only about false positives, where a person is predicted to be a criminal on information insufficient to establish the necessary individualized suspicion, because only in those instances is the Fourth Amendment's ban on unreasonable searches and seizures violated.

The practical application of the Fourth Amendment to ASAs is further complicated by doctrinal limitations on available remedies. The exclusionary rule is viewed by many as "the only remedy effective to redress a Fourth Amendment violation."[369] Others contend that civil liability under § 1983 or administrative remedies are effective.[370] Without attempting to resolve this heated debate, the focus here will be on the exclusionary rule, which is the predominant remedy for Fourth Amendment violations in criminal cases.

The "good faith" or "reasonable reliance" doctrine imposes substantial limits on the exclusionary rule. The doctrine was born in the case of *United States v. Leon*, where the Court found that "evidence obtained in objectively reasonable reliance on a subsequently invalidated search warrant" would not be subject to the exclusionary rule.[371] In *Arizona v. Evans*, the Court extended the doctrine into the digital realm, holding the exclusionary rule is inapplicable to evidence found during an unconstitutional search conducted in reasonable reliance on an incorrect database entry that a judicial court clerk had failed to correct.[372]

Then, in *Herring v. United States*, the Court applied the doctrine to refuse to suppress evidence obtained during an unconstitutional search undertaken in reliance on an expired warrant.[373] The warrant had been recalled five months

364 Illinois v. Gates, 462 U.S. 213, 231-32 (1983); United States v. Cortez, 449 U.S. 411, 418 (1981).

365 *See* United States v. Arvizu, 534 U.S. 266, 277 (2002) ("A determination that reasonable suspicion exists . . . need not rule out the possibility of innocent conduct.").

366 For example, if an ASA predicted a 60% chance that an individual was engaged in certain criminal conduct, but the ASA should have predicted a 34% chance of criminality, then the ASA's prediction was wrong.

367 *See supra* notes 83–89 and accompanying text.

368 *See supra* notes 90–98 and accompanying text.

369 Herring v. United States, 555 U.S. 135, 153 (2009) (Ginsberg, J., dissenting).

370 Hudson v. Michigan, 547 U.S. 586, 597-99 (2006).

371 468 U.S. 897, 922 (1984).

372 514 U.S. 1, 15-16 (1995).

373 555 U.S. 135, 138-39 (2009).

earlier, and a law enforcement official had failed to update the database to reflect the recall.[374] While the database error itself was negligent, the searching officer's reliance on the database was objectively reasonable.[375] The Court explained that when an error is "attenuated" from the search or seizure, such as an error in entering data or maintaining a database, the exclusionary rule will apply only if the error is "deliberate, reckless, or grossly negligent," or it involves "recurring or systemic negligence."[376] Thus, because the database error was attenuated from the arrest, the negligence of the law enforcement official in maintaining the database did not require suppression of the evidence.[377] Moreover, the defendant's failure to show that errors in the warrant database were "routine or widespread" meant that the evidence would not be suppressed on that ground either.[378] Finally, the arresting officer's objective reasonableness in relying on the information he received about the warrant also did not require suppression given the Court's holding in *Leon*.[379]

While the precise impact of *Herring* on Fourth Amendment doctrine is unclear,[380] its application to ASA errors is straightforward. Certainly, if the police employee's error in *Herring* was considered "attenuated" from the arrest, the provision of bad data to an ASA or mistakes in programming would also be considered attenuated from any search or seizure.[381] Thus, any ASA errors would require suppression only if they were the result of deliberate, reckless, or grossly negligent misconduct, or of routine or systemic negligence.

This easy application of the doctrine glosses over looming concerns about the practical impact of *Herring* on the regulation of ASAs. In her dissent, Justice Ginsburg presciently recognizes the impact of the Court's holding on more far-reaching and complex computer systems than the manual-entry warrant system before the Court. She notes first that "[e]lectronic databases form the

374 *Id.* at 138.

375 *Id.* at 140.

376 *Id.* at 144.

377 *Id.* at 137.

378 *Id.* at 147 ("But there is no evidence that errors in Dale County's system are routine or widespread."). Though the Court presents the absence of this evidence in the passive voice, the failure clearly lies with the defendant who would be expected to provide such evidence.

379 *Id.* at 146.

380 *See* Jennifer E. Laurin, *Trawling for* Herring*: Lessons in Doctrinal Borrowing and Convergence*, 111 COLUM. L. REV. 670, 671 (2011) ("The academic response to *Herring* has by and large been negative; however, to date, it has consisted as much of general puzzlement as critique.").

381 Unfortunately, the *Herring* court provides no definition for "attenuation" in its opinion. *See id.* at 687 ("Assessing the limiting work done by the Court's references to 'attenuation' is complicated by the opinion's silence as to the meaning of the term."). As used in *Herring*, however, attenuation does not seem to require the passage of time, nor intervening events that make a finding of "but for" causation too remote, "nor a disconnect between the constitutional interests protected and the harm suffered by the defendant." *Id.* at 687-88.

nervous system of contemporary criminal justice operations."[382] Yet such systems are inadequately monitored and contain numerous errors.[383] Moreover, if a defendant must show deliberate, reckless, or grossly negligent conduct, or routine or systemic errors, to get relief in cases of computer errors, then the defendant likely needs discovery or even an opportunity to audit police databases in order to prove these kinds of errors.[384] Finally, without some threat of evidence being suppressed, police may not have sufficient incentives to maintain accurate databases and computer systems.[385]

In a similar vein, Erin Murphy has articulated a number of "shared features that inhere across databases generally," some of which are applicable to ASAs.[386] First, she recognizes that databases are best regulated at a structural level, rather than on a case-by-case basis.[387] This is because databases are constructed by numerous people, spread out across time and geography, with different roles and motivations.[388] A single defendant and her counsel do not have the resources, or the motivation, to stare down this massive, interlocking structure, see the problems, and push for large-scale solutions.[389] Second, Murphy observes that databases "tend to operate anonymously" and their "content is typically shrouded in secrecy."[390] Litigation, as a presumptively public event, runs contrary to this obscurity, and thus "[i]t is too much to require courts, or to expect the Constitution, to demand full transparency in the methods of database administrators."[391] Finally, Murphy notes that "it is far easier to do harm, and far greater harm can be done, through mere benign neglect of database systems than through intentional manipulation."[392]

Justice Ginsburg's and Erin Murphy's concerns are likely to be realized if police use of ASAs is left unregulated. By encroaching on the second step of

---

382 *Herring*, 555 U.S. at 155 (Ginsburg, J., dissenting).

383 *Id.*; *see also* Christopher Slobogin, *Government Data Mining and the Fourth Amendment*, 75 U. CHI. L. REV. 317, 324 (2008) ("Most fundamentally, the information in the records accessed through data mining can be inaccurate.").

384 *See Herring*, 555 U.S. at 157 (Ginsburg, J., dissenting) (noting that the majority failed to identify how a defendant is supposed to make the required showing that "deliberate or reckless conduct is afoot," and that possible answers of an entitlement to discovery or an audit of police databases would entail "considerable administrative burden").

385 *Id.* at 156.

386 Murphy, *supra* note 32, at 826.

387 *Id.*

388 *Id.* at 827–28.

389 *See id.* at 828-29 ("Discovery, compulsory process, or cross-examination in a single case yields little opportunity to identify and uncover, much less broadly correct, errors apt to occur (and be visible) only from scrutiny on a systemic level.").

390 *Id.* at 831 (emphasis omitted).

391 *Id.* at 832.

392 *Id.* at 835.

the individualized suspicion analysis, ASAs can become the "nervous system" of the criminal justice system. Unless forced into an interpretable model, an ASA's operations are not just obscure, they are completely opaque. The result is that inquiry into these operations is not only difficult and impractical for defendants, but impossible. Moreover, ASAs can potentially analyze so much data that discovery of the underlying databases would overwhelm even the most industrious and well-funded defense counsel.[393] As a result, an ASA's individual prediction errors would be difficult to uncover in most cases. Even if an ASA's errors could be found, it would be effectively impossible for a defendant to make the showing of either willful misdeeds or routine errors that would be necessary to suppress evidence or encourage reform. Thus, bad data and benign neglect could flourish in the ecosystem of an ASA if the only oversight comes from case-by-case Fourth Amendment adjudication.[394]

The certainty of ASA errors, therefore, militates in favor of systemic oversight. Again, the creation of standards governing the programming of ASAs by experts in relevant subject-matter fields, such as machine learning, law, and law enforcement, is of paramount importance. These standards should cover the training and programming of the ASAs, as well as provide continued oversight to ensure that ASAs learn from new data so errors are minimized and effectiveness is maximized. However, standards alone are not enough. Courts must also be willing to require ASAs to meet these standards and to exclude evidence obtained by ASAs that do not. In particular, courts must not require defendants to make the almost-impossible showing that an ASA's specific failure to meet the standards led to some articulable harm to the defendant.[395] Fortunately, *Herring* continued to limit the good faith exception to the exclusionary rule to situations where police act in objectively reasonable reliance on the evidence in question.[396] Courts should give teeth to this limitation by recognizing that when ASA standards exist, reliance by police on an ASA that does not meet those standards is unreasonable.

---

393 *See* Ferguson, *supra* note 20, at 354-60 (describing the volume of data available to police).

394 *See also* Citron & Pasquale, *supra* note 9, at 1481-83 (discussing political reasons why courts may be ineffectual in monitoring "fusion centers" that accumulate and use massive amounts of data in counterterrorism).

395 *See* Murphy, *supra* note 32, at 822-23 (noting that the normal requirement of proof of a specific articulable harm to the defendant from a database may be impossible to meet because of "the diffused and decentralized nature of databases" and the incentive of those in charge of the database to protect themselves from blame).

396 *See* Herring v. United States, 555 U.S. 135, 142 (2009) ("[T]he exclusionary rule does not apply if the police acted 'in objectively reasonable reliance' on the subsequently invalidated search warrant." (quoting United States v. Leon, 468 U.S. 897, 922 (1984))).

CONCLUSION

The overarching lesson of the preceding discussion is that ASAs are not an easy fit for existing Fourth Amendment doctrine. While it is possible that courts will undertake the "major reorientation in constitutional thinking" that ASAs and similar networked technologies demand,[397] such a substantial shift seems unlikely, at least before ASAs enter the mainstream. Instead, police, magistrates, and litigants must find ways to analyze them logically within existing doctrine. This Article provides a beginning framework for that analysis. First, though ASAs intrude on the second step of the individualized suspicion analysis, they cannot replace a human being when it comes to consideration of the totality of the circumstances in each case. Instead, their predictions are merely another fact, albeit perhaps a weighty one, in that analysis. Second, an ASA's predictions are best analogized to a drug dog's alert in the totality-of-the-circumstances analysis. Unfortunately, flaws in current Supreme Court doctrine on drug dogs and the unique characteristics of ASAs suggest that more work needs to be done by legal scholars and experts in machine learning, law, and policing. In particular, uniform and robust standards are needed for the programming, training, and continued use of ASAs—including monitoring of the data used by ASAs—to maximize ASA accuracy and minimize errors resulting from bad data and programming errors. In addition, courts must be prepared to give defendants latitude in their questioning of ASA reliability and to exclude evidence obtained as a result of officer reliance on uncertified ASAs or ASAs that are not proven to be reliable.

397 Murphy, *supra* note 32, at 829.

* * * * *