

# Adversarial Scrutiny of Evidentiary Statistical Software

Rediet Abebe

Department of Electrical Engineering  
& Computer Sciences, University of  
California, Berkeley  
Berkeley, CA, USA

Moritz Hardt

Max Planck Institute for Intelligent  
Systems, Tübingen  
Tübingen, Germany

Angela Jin

Department of Electrical Engineering  
& Computer Sciences, University of  
California, Berkeley  
Berkeley, CA, USA

John Miller

Department of Electrical Engineering  
& Computer Sciences, University of  
California, Berkeley  
Berkeley, CA, USA

Ludwig Schmidt

Paul G. Allen School of Computer  
Science & Engineering, University of  
Washington  
Seattle, WA, USA

Rebecca Wexler

School of Law, University of  
California, Berkeley  
Berkeley, CA, USA

## ABSTRACT

The U.S. criminal legal system increasingly relies on software output to convict and incarcerate people. In a large number of cases each year, the government makes these consequential decisions based on evidence from statistical software—such as probabilistic genotyping, environmental audio detection and toolmark analysis tools—that the defense counsel cannot fully cross-examine or scrutinize. This undermines the commitments of the adversarial criminal legal system, which relies on the defense’s ability to probe and test the prosecution’s case to safeguard individual rights.

Responding to this need to adversarially scrutinize output from such software, we propose *robust adversarial testing* as a framework to examine the validity of evidentiary statistical software. We define and operationalize this notion of robust adversarial testing for defense use by drawing on a large body of recent work in robust machine learning and algorithmic fairness. We demonstrate how this framework both standardizes the process for scrutinizing such tools and empowers defense lawyers to examine their validity for instances most relevant to the case at hand. We further discuss existing structural and institutional challenges within the U.S. criminal legal system which may create barriers for implementing this framework and close with a discussion on policy changes that could help address these concerns.

## CCS CONCEPTS

• **Applied computing** → **Law**; • **Computing methodologies** → **Machine learning**; • **Software and its engineering**;

## KEYWORDS

evidentiary software, statistical software, adversarial scrutiny, black-box software, robust machine learning

## ACM Reference Format:

Rediet Abebe, Moritz Hardt, Angela Jin, John Miller, Ludwig Schmidt, and Rebecca Wexler. 2022. Adversarial Scrutiny of Evidentiary Statistical Software. In *2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*, June 21–24, 2022, Seoul, Republic of Korea. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3531146.3533228>

## 1 INTRODUCTION

The U.S. criminal legal system increasingly uses output from software tools to prosecute the criminally accused. Statistical software is especially gaining prominence and has been deployed in a range of settings, from probabilistic genotyping to environmental audio detection, and forensic bullet analysis. At present, while the government uses output from such software as evidence, defense counsel are often unable to fully cross-examine or scrutinize these tools. This limitation, coupled with the accelerated use of machine learning and statistically-driven software, presents a clear and growing danger to the legal truth-seeking process.

The U.S. uses an adversarial system, where law enforcement is tasked with investigating evidence of guilt, and defense counsel is tasked with probing the prosecution’s case and investigating evidence of innocence. By pitting these two roles against one another, the system as a whole seeks legal truth and the protection of individual rights. In the context of evidentiary software, however, numerous factors prevent defense counsel from effectively performing their duty to scrutinize the prosecution’s case.

One such factor is uncertainty over what adversarial scrutiny of a software program entails [42]. The U.S. Constitution guarantees criminal defendants the right to confront and cross-examine the witnesses against them [8]. Yet, courts, counsel, and legal scholars have fractured over how and whether to apply these rights to output from software [92]. Compounding this uncertainty, there are no technical standards to adversarially scrutinize output from software used as evidence of guilt. Existing mechanisms for doing so are insufficiently powerful and incomprehensive to safeguard against erroneous output.

In this work, we tackle this issue of evaluating the validity of evidentiary statistical software by proposing a technical conceptual framework—*robust adversarial testing*. We draw on a growing body of work in robust machine learning to formally define robust adversarial testing and describe how to operationalize it within the criminal legal system. This framework connects the notion of



This work is licensed under a [Creative Commons Attribution International 4.0 License](https://creativecommons.org/licenses/by/4.0/).

FAccT '22, June 21–24, 2022, Seoul, Republic of Korea  
© 2022 Copyright held by the owner/author(s).  
ACM ISBN 978-1-4503-9352-2/22/06.  
<https://doi.org/10.1145/3531146.3533228>

adversarialism in the legal system with adversarial robustness in machine learning. The connection is of mutual benefit: Defense counsel can repurpose tools and insights from robust machine learning in service of adversarial scrutiny. The specific requirements of defense counsel, in turn, contribute well-motivated research and engineering opportunities to the machine learning community. In doing so, this framework provides a process that is concrete enough to be adopted for adversarial scrutiny of statistical software across the legal system and flexible enough to empower the defense to adversarially select input for testing on a case-by-case basis.

Specifically, our proposal of robust adversarial testing rests on two components: First, the defense can adversarially select a family of input to the statistical software that is most relevant to the defendant’s case. This could entail, for instance, testing the software’s performance for specific demographic groups. It also allows the defense to adversarially select relevant performance measures such as the rate of false inclusion of the defendant in a forensic sample. Through these two levers, the defense may adversarially test the validity of the statistical software on the input most relevant to the case at hand.

In addition to providing a framework for adversarially scrutinizing evidentiary statistical software, our proposal adds to and complements a growing body of work in the accountability community that tackles challenging tasks around auditing algorithmic and software tools used at various stages of the criminal legal system. While our motivation in this work is evidentiary software, the framework here can apply more broadly to statistical software used at any stage of a legal case. As such, we take a step towards systematizing and formalizing a key component of audits of statistical software used in the law more generally.

Our work also goes beyond existing research in algorithmic audits and accountability in its focus on evidentiary software. Existing work thus far has primarily focused on tools predominantly used at the investigation, pretrial, and sentencing stages, such as risk assessment tools and facial recognition technology [13, 19, 28, 33, 39, 41, 59, 65, 68, 70, 82, 88, 99, 112]. While output from these tools may at times be used as evidence—and indeed there is a credible concern that they will be widely adopted as evidentiary tools—this is not their current primary use. As a result, the algorithmic accountability literature has not kept up pace with evaluating evidentiary software. In presenting our adversarial scrutiny framework, we highlight a path forward for the machine learning community to consider evaluating a broader range of software tools and to contribute resources that could empower defense counsel to more thoroughly scrutinize evidence of guilt.

The rest of this paper is structured as follows: In Section 2, we set the stage for our framework by introducing the role of adversarial scrutiny in the U.S. criminal legal system. In Section 3, we provide a brief introduction to probabilistic genotyping software, which we use as a running example throughout the remainder of the paper. In Section 4, we consider existing mechanisms for scrutinizing evidentiary software and lay out their limitations in assessing statistical software in particular. We present our proposal in Section 5. We define and introduce adversarial scrutiny of evidentiary statistical software and its operationalization in the U.S. criminal legal system. While our primary contribution is to conceptualize the legal concept of adversarial scrutiny in a technical formulation,

we acknowledge that existing structural and institutional barriers may prevent defense counsel from implementing our framework in practice. We describe some of these challenges that the law and policy communities should simultaneously tackle in Section 6.

## 2 LEGAL BACKGROUND: ADVERSARIAL SCRUTINY IN THE LAW

In this section, we introduce the reader to the concept of adversarial scrutiny in the law. We begin with a general commentary about adversarialism and explain why criminal defense adversarial scrutiny is essential for forensic evidence in particular. We then describe some of the legal rules and procedures that are designed to facilitate this scrutiny at three different stages of a criminal case: plea negotiations, admissibility hearings, and trial. Importantly, these examples are meant to be illustrative and not comprehensive. Criminal defense counsel should perform adversarial scrutiny at *all* stages of a case.

The U.S. uses an adversarial criminal legal system. This means law enforcement, including prosecutors and the police, are responsible for investigating evidence of guilt. Criminal defense counsel, on the other hand, are responsible for investigating evidence of innocence and, crucially, for identifying flaws and weaknesses in the prosecution’s evidence. Criminal defense counsel are the sole actors in the entire legal system with a duty to probe, test, and challenge the prosecution’s evidence [125]. We call this duty *adversarial scrutiny*.

Defense counsel are supposed to perform adversarial scrutiny not as neutral inquisitors, but rather as “zealous advocates” for the criminally accused [113]. In this sense, defense counsel are similar to white hat hackers who are supposed to focus on finding flaws and weaknesses in a computer system. At trial, both the prosecution and defense present their evidence and arguments to the jury, which ultimately decides whether the prosecution has proven the defendant’s guilt beyond a reasonable doubt. Similarly, in pretrial and post-conviction proceedings, both sides present evidence and arguments to the judge, who ultimately decides the legal issues in the case. By pitting the prosecution and defense advocates against one another, the system as a whole seeks legal truth and the protection of individual rights [87].

Adversarial scrutiny of forensic evidence can be at once technically complex and exceptionally urgent. Invalid forensic evidence has contributed to a substantial portion of the wrongful convictions of individuals who were later exonerated [49, 97]. Indeed, prior studies have identified invalid forensics as the “second most common contributing factor” in wrongful convictions [122]. It is therefore essential for defense counsel to perform effective adversarial scrutiny when prosecutors use evidentiary statistical software as forensic evidence.

The need for criminal defense counsel to perform adversarial scrutiny of forensic evidence can apply at any stage of a case. For example, the Sixth Amendment constitutional right to effective assistance of counsel during pretrial plea negotiations may require

defense counsel to independently investigate and test the prosecution’s forensic evidence early in a case [45].<sup>1</sup> Accordingly, as soon as possible after a case begins, the defense should demand access to and relevant information about any evidentiary statistical software via statutory discovery and constitutional due process *Brady* rules, which require the prosecution to disclose certain information to the defense [51], as well as by serving pretrial subpoenas on developers or other third parties for additional relevant information that the prosecution may not know or possess [51]. These discovery and subpoena rules should enable the defense to obtain information to perform adversarial scrutiny at this and subsequent stages of the case.

As another example, the need for adversarial scrutiny may also arise during pretrial admissibility hearings, in which the prosecution and defense argue before a judge about what evidence they should be allowed to present to the jury at an upcoming trial. Based on the arguments from both sides, the judge will determine whether the evidence satisfies the requirements for admissibility, which include relevance [4] and lack of unfair prejudice [5], among other things.<sup>2</sup> For expert testimony based on evidentiary statistical software, the rules of evidence require an additional showing that the software is a method that has been reliably applied in the defendant’s case [6].

There are two leading standards across federal and state jurisdictions for this additional showing: *Daubert* and *Frye* [26]. *Daubert* gives judges flexibility in how they choose to make the reliability determination but encourages them to consider the following series of factors: “whether the theory or technique in question can be (and has been) tested, whether it has been subjected to peer review and publication, its known or potential error rate and the existence and maintenance of standards controlling its operation and whether it has attracted widespread acceptance within a relevant scientific community” [36]. The test in *Frye* jurisdictions focuses on one factor: whether a method is “sufficiently established to have gained general acceptance in the particular field in which it belongs” [44]. Critics have castigated judges for failing to perform their proper gate-keeping function to keep unreliable forensics out of court, particularly in criminal cases [48, 50, 54, 57, 64, 109]. Such failures emphasize the urgent need for effective defense adversarial scrutiny in admissibility hearings so as to better educate judges about the flaws in prosecution forensic evidence and persuade them to do a better job of excluding faulty forensics from trial.

If the defense loses an admissibility challenge and the judge permits the prosecution to present forensic evidence to the jury, then this brings us to our third example of a point in time for adversarial scrutiny: trial. At trial, defense counsel should tell the jury about any flaws and weaknesses in the prosecution’s evidence. Even if the *Daubert* and *Frye* standards were applied correctly, much imperfect evidence would still be admitted at trial because the requirements for admissibility in the rules of evidence are relatively lenient, including for expert testimony [16]. This liberal thrust of the rules in favor of admissibility reflects an overarching policy preference

for trusting juries to properly weigh evidence, including by discounting less-reliable evidence, rather than empowering judges to keep that evidence entirely secret from the jury. Further, trial courts have substantial discretion on admissibility determinations and any potential errors in their judgments are reviewed under a deferential standard by the appellate courts [105]. As a result, trials can be flooded with under-scrutinized forensic evidence that defense counsel should contest before the jury.

Several legal rules and procedures facilitate such contestations at trial. For instance, the Sixth Amendment to the U.S. Constitution guarantees criminal defendants the right to confront and cross-examine the witnesses against them [8], including prosecution witnesses who testify during trial based on the output of evidentiary software [18, 83, 126]. Legal scholars and courts are currently struggling to determine whether, to what extent, and how, this right may also entitle defendants to directly cross-examine and confront evidentiary software systems themselves [25, 103]. Some have suggested that lawmakers or the Constitution might require software developers, designers, or operators to take the witness stand [103, 114]. Others have advocated for defense access to source code, or defense access to “enhanced discovery” that would afford increased opportunities to test prosecution software [25]. To date, however, few have offered technical details about what the defense should do with this information after they obtain it.

In this paper, we present a technical framework to operationalize defense adversarial scrutiny of evidentiary statistical software. As the preceding paragraphs explain, the criminal legal system depends on an adversarial framework that tasks defense counsel with contesting the prosecution’s forensic evidence. The existing legal adversarial framework motivates our technical proposal, which we present in detail in Section 5. We argue that defense adversarial scrutiny of evidentiary statistical software, or careful examination directed to identifying flaws or weaknesses in the software output, is crucial to protect defendants’ rights and to the truth-seeking goals of the criminal legal system as a whole.

### 3 A CASE STUDY: PROBABILISTIC GENOTYPING SOFTWARE

In this section, we discuss probabilistic genotyping software (PGS) tools, a popular class of evidentiary statistical software, which we will use as a running example to ground the discussions and definitions that follow.

PGS tools are widely used in criminal cases. According to the creator STRmix, this tool has been used in over 220,000 cases worldwide [1]. Likewise, the Forensic Statistical Tool (FST), was used in over 1,300 cases from 2011 to 2017 [66]. Like many statistical decision-making tools, at their core, PGS tools implement a likelihood ratio test [30, 55]. Consequently, though we focus on PGS for concreteness, our discussion and approach apply to a much broader class of evidentiary software based on statistical techniques. Note, in discussing these tools, our goal here is not to take a stance on the validity of these tools in any specific instance but rather to provide a concrete framework to scrutinize these tools.

Roughly, forensic DNA analysis begins with a biological sample of interest obtained from a crime scene and a known reference sample taken from the defendant. The observed crime scene sample

<sup>1</sup>We thank Brandon Garrett for highlighting the significance of the Sixth Amendment right to effective assistance of counsel in establishing a constitutional basis for defense testing of prosecution forensic evidence, especially pre-plea.

<sup>2</sup>The Appendix describes the variety of rules that govern the admissibility of evidentiary statistical software output.

varies in complexity from simple settings involving DNA from a single person to complex cases involving a mixture of DNA from multiple contributors, e.g., a sample from a door handle that was potentially touched by multiple people. We commonly refer to the observed sample as a “mixture.” From these samples, an analyst extracts DNA profiles using standard laboratory techniques. PGS tools take, as input, both the observed DNA profile from the crime scene and the reference DNA profile from the defendant. They then attempt to estimate a likelihood ratio that compares the probability of observing the crime scene DNA if the reference candidate were a contributor to the mixture versus a random (unrelated) individual. PGS tools rely on, among other things, population genetic databases and laboratory-specific parameters (e.g., probability of allele drop-in) to estimate this likelihood ratio. Specific details are beyond the scope of this paper, and we refer the interested reader to [20, 30, 81] and the references therein for more background.

Despite their widespread use, the validity of PGS tools is an ongoing research direction and the source of much debate [20, 80, 81, 93, 95]. For instance, in a high-profile 2016 case, two different leading PGS tools produced different conclusions when used to analyze the same DNA sample [95, 96]. Many of the questions around validity stem from the fact that it depends on the complexity of the samples and the specific settings in which it is used [20, 81, 95]. Concretely, Butler et al. [20] detail more than 20 distinct factors which all influence the performance and validity of PGS tools. These factors include, among others, (i) the number of contributors to a mixture, (ii) the amount of DNA collected, (iii) the relative proportion of material from each contributor, (iv) the quality of the genetic material, (v) the presence of close relatives in the sample, and (vi) calibration of software parameters to laboratory equipment. In another study, Matthews et al. [81] show how variations in the parameters of the underlying statistical model, e.g., a parameter summarizing the genetic diversity in the population, can cause substantial changes in the likelihood ratios produced by different tools. Understanding the tool’s validity in any particular setting, therefore, requires understanding how these factors influence its performance.

## 4 LIMITS OF CURRENT ASSESSMENT MECHANISMS

At present, courts, counsel, and commentators have discussed a variety of mechanisms for assessing the validity of evidentiary software that defense counsel may apply to statistical tools. We focus our discussion on three frequently considered mechanisms: source code review, validation studies by relevant scientific communities, and direct testing by the defense. In this section, we consider how defense counsel might face limitations when using each mechanism to scrutinize the validity of statistical software.<sup>3</sup> We focus our discussion on PGS tools to provide specific examples, but our arguments apply to other evidentiary statistical software as well. In discussing these limitations, we motivate our proposed framework that we introduce in Section 5.

<sup>3</sup>Given our specific focus on statistical software, we do not aim to present a comprehensive list of limitations for each mechanism.

### 4.1 Source Code Review by the Defense

Several groups have called for defense counsel to assess the validity of evidentiary software using source code review [14, 27, 60, 74, 89, 108]. While examining the source code may help uncover programming mistakes, or “miscodes,” looking at the source code without also running the software executable may reveal little about the statistical tool’s performance (e.g., accuracy or error rates), consequently leaving out a crucial metric for assessing the validity of statistical software output used in a case.<sup>4</sup> Illustrating this limitation in the context of PGS tools, scrutinizing the source code can verify whether the program correctly implements a fully continuous PG model. Without also having the ability to test the statistical software, however, defense counsel cannot fully determine the tool’s performance (e.g., error rates) in any setting, including that of the defendant’s case. Given this limitation, empirical testing is crucial for assessing the tool’s performance and, consequently, for assessing the validity of the tool’s output in any given case.

### 4.2 Validation Studies by Relevant Scientific Communities

Defense counsel may also use empirical assessment mechanisms to gauge the tool’s performance. One such mechanism the criminal legal system already relies upon is validation studies. However, these studies face numerous limitations.

Many groups have rightfully pointed out that existing validation studies, even those that are peer-reviewed, do not rigorously assess evidentiary software because the authors or reviewers of such studies have financial or professional interests in ensuring that forensic laboratories use the software. For instance, amongst the eight TrueAllele validation studies that Butler et al. [20] enumerate, seven studies include Mark Perlin, the CEO of Cybergenetics Inc. which produces TrueAllele, as a co-author. The remaining study acknowledges Perlin and two Cybergenetics employees for their guidance and comments. In light of observations like these, many argue that independent groups with no stake in the outcome of validation studies should test the software [20, 21, 63, 73, 86, 95, 102, 108, 109].

Even if more independent scientific groups were to validate evidentiary software, before conducting each study, the study designer must choose a set of test cases on which to evaluate the tool. As several commentators have pointed out, these chosen test cases are not guaranteed to align with the characteristics of any actual criminal case in which the tool is applied [20, 72, 95]. Indeed, this is true regardless of whether the parties conducting the studies are independent of the developer or not. Consequently, the tool’s performance and error rates reported in a validation study are not guaranteed to reflect the tool’s validity when applied in the criminal legal system. Scholars have observed that evaluation studies of machine learning systems are often over-optimistic estimates of performance on new test cases; the performance of these systems can drop dramatically under slight variation of the test distribution [19, 32, 69, 76, 101, 127].

<sup>4</sup>E.g., Bellovin et al. [14], President’s Council of Advisors on Science and Technology [95] state that foundational validity requires that a method has been subjected to empirical testing that provides “valid estimates of the method’s accuracy”.

In sum, validation studies require the designer to choose test cases and parameters that may not match those relevant to the case at hand. This mismatch can result in significant differences in the statistical software’s performance and output, so testing the tool on input similar to the case at hand is imperative to ensure the tool, as applied in the specific case, is valid.

### 4.3 Direct Testing by the Defense

Instead of relying on other parties to empirically test the statistical software in contexts outside of any given case, defense counsel may directly test the tool by choosing input and running the executable program. Indeed, organizations like Upturn have joined defense attorneys in arguing in court for defense access to executable source code and additional resources to make this direct assessment [108], and many scholars have made similar arguments [2, 14, 74]. We agree that defense access to the executable source code is crucial for assessing software validity. We argue, however, that current discussions of testing techniques are incomplete.

In particular, current discussions do not provide defense counsel with sufficient guidance to conduct such tests. While some groups have specifically suggested that the defense test software on input relevant to the defendant’s case and have provided examples [14, 74], these discussions do not provide any overarching framework that would equip defense counsel to make their determinations when navigating steps such as choosing input that are relevant to their client’s case. To ensure that defense counsel can effectively test evidentiary statistical software, such guidance is necessary.

Acknowledging the limitations of existing proposals for testing and of the other mechanisms we described in the preceding subsections, we propose one potential framework for testing statistical software, which we present in the next section.

## 5 ROBUST ADVERSARIAL TESTING

We now present our framework for adversarial scrutiny of evidentiary statistical software. Before we introduce our mathematical definition, we first provide the underlying intuition in Section 5.1. In Section 5.2, we discuss how defense counsel can operationalize our definition.

As the previous section motivates, we would like to answer the question: *Is the tool in question accurate on input corresponding to the defendant’s case?* Accurate evidence from statistical software tools will aid the truth-seeking process of the criminal legal system and minimize the number of wrongful convictions. It is well-known however that standard statistical validation provides no simple answer to this question because statistical validation makes claims about *distributions of instances*, and not individual instances. Even if a statistical validation study establishes that a tool performs accurately across a broad distribution of instances, such a finding does *not* imply any rigorous guarantees for a previously unseen instance relevant to the defendant’s case. For example, a probabilistic genotyping tool may be highly accurate on DNA mixtures containing genetic material from two individuals, but may fail to produce correct answers when the mixture has more or an unknown number of contributors. Similarly, gunshot detection tools may be highly

accurate in otherwise quiet environments but have high error rates when the ambient environment is filled with noise from traffic, fireworks, or construction.

The difficulties surrounding the accuracy of statistical software can have direct consequences on the validity of evidence in court. Consider a hypothetical probabilistic genotyping tool that achieves perfect accuracy for 99% of DNA mixtures, but is always wrong for the remaining 1% of instances. If the defendant in question belongs to this 1% of instances, the tool will be guaranteed to fail in this case despite its seemingly impressive accuracy.<sup>5</sup>

To formalize this intuition and ground the legal concept of adversarial scrutiny in the technical reality of statistical software, we build on conceptual advances from the literature on distributional robustness in machine learning and algorithmic fairness [15, 17, 31, 40, 56, 98, 121]. At its core, robust adversarial testing requires a tool to perform well on input most relevant to the defendant’s case. For instance, the accuracy of a DNA analysis tool on mixtures with two participants may have no relevance for a case where the mixture comes from five contributors. Determining precisely what input is relevant to a defendant’s case is a key step to operationalize our definition of adversarial scrutiny. Our definition empowers the defense to adversarially select input for testing the evidentiary software in question.

Next, we provide a definition of robust adversarial testing and then discuss how to operationalize this definition in practice. Intuitively, our definition decomposes assessing the performance of evidentiary software into two parts: (i) the metric used to check performance, and (ii) the distribution on which the metric is applied. By allowing the defense to choose each component adversarially, our framework enables the defense to question evidentiary software in a way most relevant to the case at hand.

### 5.1 Formal Definition of Robust Adversarial Testing

We now state our model for adversarial testing. To ground this definition, it is helpful to keep in mind the example of probabilistic genotyping tools introduced in Section 3. A statistical tool computes a function  $A: Z \rightarrow O$ , where  $O$  is a set of output and  $Z$  is a data universe of possible input. For a PGS tool, the output space  $O$  may be the space of likelihood ratios, and the input  $Z$  are the DNA samples taken from the crime scene and the defendant.

For a given tool  $A$  and distribution  $Q$  over the data universe  $Z$ , let  $\text{check}(A, Q) \in \{\text{Pass}, \text{Fail}\}$  be a test that determines whether the tool  $A$  has *acceptable* performance on the distribution  $Q$ . For instance, the distribution  $Q$  may be a distribution over DNA mixtures with two contributors, and the check function may compute the error rate of a PGS tool  $A$  on the distribution  $Q$ , returning *Pass* only if the error rate is below some threshold. In a concrete validation study for a PGS tool, a finite dataset of test examples instantiates the distribution  $Q$ . The notion of “acceptable performance” references the idea of “general acceptance” in the law from Section 2. We envision that this check function can also be chosen adversarially by the defense to test for flaws that generally accepted evaluation

<sup>5</sup>Here and in other examples we use accuracy as a performance metric for ease of exposition. In an actual application of our framework, other performance metrics such as precision, recall, ROC curves, and others may be more meaningful depending on the tool and case in question.

metrics may hide. We will return to this point when we discuss how to operationalize the definition.

Further, let  $F$  be a family of distributions over the same data universe  $Z$ . Concretely, this could be a collection of distributions over DNA mixtures with the same number of contributors as the sample in question, but perturbed values of laboratory parameters or quantities of DNA collected. Specifically, a family of distributions  $F = \{Q_1, Q_2, Q_3\}$  could consist of three individual test distributions where:

- $Q_1$  is a distribution over DNA mixtures with two contributors and 100–200 pg of DNA.
- $Q_2$  is a distribution over DNA mixtures with two contributors where the DNA mixture was analyzed with lab equipment from a specific manufacturer.
- $Q_3$  is a distribution over DNA mixtures with two contributors and a contributor ratio between 2:1 and 4:1.

If the trial involves a DNA sample with the above characteristics (two contributors, 100–200 pg of DNA, lab equipment from the same manufacturer as in  $Q_2$ , and a likely contributor ratio between 2:1 and 4:1), each of the three distributions in the family  $F$  stress tests the PGS tool in question on input that are relevant to the case. In addition, each distribution stress tests the PGS tool differently. Hence the PGS tool’s accuracy on these distributions is likely more representative of its accuracy on the defendant’s input than the result of a validation study on a broad distribution over DNA mixtures that may only be superficially related to the DNA sample in question. As before, a dataset would instantiate each of the individual distributions  $Q_i$  in a concrete adversarial evaluation during an individual criminal case.

With our formal setup in place, we now state our main definition:

*Definition 5.1 (robust adversarial testing).* We say that a tool  $A$  passes ( $F$ , check)-robust adversarial testing if for every distribution  $Q$  in the family  $F$  we have  $\text{check}(A, Q) = \text{Pass}$ .

Importantly, since the performance check must pass for all distributions  $Q$  in the family  $F$ , it must also pass for the worst-case distribution in the family. Therefore, the defense choosing the family  $F$  corresponds to implicitly picking the worst-case distribution in the family on which to evaluate the tool. Our definition gives the defense broad latitude to choose the family  $F$ . However, nuance is required to select a family  $F$  that is relevant to the defendant’s instance and avoids trivial examples, e.g., a point-mass on a known failure case.<sup>6</sup> Selecting an appropriately rich and relevant family  $F$  is one of the key challenges involved with operationalizing this definition.

Our definition makes use of a binary check function that ultimately outputs that a tool either “passes” or “fails.” One could also consider a continuous version of our definition where we simply report the worst-case metric over the distribution family. We elected to use a binary check function to abstract away questions of statistical validity (e.g., confidence intervals can be implicitly included in

<sup>6</sup>Note that, if the defense were to choose *irrelevant* input, the prosecution could likely have the results excluded from the case because “irrelevant evidence is not admissible.” See, e.g., Fed. R. Evid. 402. Hence, the legal concept of relevance binds the defense selections in our framework. “Evidence is relevant if: (a) it has any tendency to make a fact more or less probable than it would be without the evidence; and (b) the fact is of consequence in determining the action.” Fed. R. Evid. 401.

the check function) because judges and juries must ultimately make a validity determination for a particular tool in a particular case. For a binary check function, the choice of the threshold which constitutes “passing” is often sensitive and potentially contentious [37]. In our definition, we empower the defense to select the threshold, and we envision that relevant research communities can empower the defense with guidelines for selecting this threshold. We discuss this in more detail below.

## 5.2 Operationalizing Robust Adversarial Testing

We now discuss how to implement and operationalize our formal definition. We operationalize our definition in three steps. Because our definition is adversarial, we envision each step will be carried out by the defense and specific to the case at hand. The first step involves determining the data universe. The second step is the adversarial choice of a relevant distribution of input. Finally, the third step is the choice of a reasonable performance check to measure how well the tool performs on the chosen distribution of input.

*5.2.1 Determining the Data Universe.* The first step of our procedure is to determine what, precisely, is that data universe. Conceptually, one can think of the data universe as a large database enumerating all possible inputs to the system under study, background factors, system output, and ground truth. In the case of PGS tools, this ideal database may correspond to all possible DNA samples (with differing numbers of contributors, amount of genetic material, etc.), parameters of the analysis tool, and features of the defendant (ethnicity, age, etc.), and other such considerations.

Such idealized databases are unavailable in practice. As a practical compromise, we envision the data universe as being the union of existing data from validation studies. For example, several hundred two, three, and four-person mixture samples are publicly available through the ProvedIt database from Alfonse et al. [11], and Butler et al. [20] details more than 60 validation studies of varying size testing PGS tools under different input conditions. The extent to which existing databases are complete and accessible is an ongoing challenge. However, the lack of representative data in existing databases may itself indicate that the tool has not been sufficiently validated.

*5.2.2 Adversarial Selection of the Input Distribution.* After specifying the data universe, the next step of our adversarial procedure requires the defense to choose the family of distributions on which we evaluate the tool. The choice of distributions is not just specific to the problem the software maker purports to solve. It is also specific to the defendant. The defense will choose the distributions to reflect salient information available about the defendant’s case.

For example, if the DNA mixture in the defendant’s case likely involves more than three contributors and trace amounts of DNA, then the defense might evaluate the performance of the tool on other mixtures with more than three contributors and similar quantities of DNA. As a second example, consider a hypothetical software tool for analyzing videos, e.g., to determine whether a specific individual appears in a given video segment. If the tool analyzed video footage of a supposed crime that occurred at night, where the accuracy of computer vision algorithms may be worse, it is reasonable for the

defense to interrogate the performance of the tool specifically on nighttime video. The method of determining the input distributions necessarily depends on context.

We propose that the tool should have acceptable performance at least on instances sharing some number of features in common with the defendant’s instance. For PGS tools, Butler et al. [20] describes more than 20 distinct features that contribute to the performance and validity of these systems: the number of mixture components, the quantity of DNA, and so on. In this case, we would evaluate the tool on distributions of instances that match some number of the defendant’s features (e.g., the same number of mixture components and the same quantity of DNA). Formally, this corresponds to taking the distribution family to be all  $k$ -way marginals defined by the defendant’s features.

To implement this procedure, we start with the database of all instances described in the previous step. The defense then defines each possible distribution by first picking one or more features, for instance, the amount of DNA in the sample, that are deemed relevant. Then, we consider the defendant’s value or range of possible values for this feature, say 100 picograms, and take all instances in the database that approximately match this feature and have between, say, 95-105 picograms of DNA. This defines one particular distribution. The family of distributions then arises from repeating this process multiple times and also considering feature interactions (e.g., matching on both DNA quantity and the number of mixture contributors).

Determining which features are relevant, the number of features to consider jointly, and the number of distributions to include in the family ultimately depends on the context involving both the tool in question and the case at hand. These choices present a trade-off between requiring more specific validation of the evidentiary software and placing an unreasonably large burden on its validation that would preclude the use of all statistical software (e.g., if the family of distributions is too large or involves too specific distributions).

**5.2.3 Performance Check Selection.** The choice of a performance check relies on the defense choosing an evaluation metric that is appropriate for the particular tool and instance at hand. For instance, President’s Council of Advisors on Science and Technology [95] recommend using the inclusion error rate to evaluate a variety of forensic procedures. In general, the choice of a performance check should incorporate the notion of “general acceptance” borrowed from the legal standard, with specific attention to the possibility that a method will fall out of general acceptance. Concretely, the relevant legal and scientific research communities can keep track of what is considered acceptable for a task at hand. This empowers relevant research communities to remove old standards from general acceptance as acceptable performance evolves. This could happen, for example, by publishing results with improved accuracy or better trade-offs between false positives and negatives, which the defense can then point to when establishing the standard used in the check function. For instance, a software tool that recognizes handwritten digits with 80% accuracy no longer has general acceptance, since the task has long been solved with significantly higher accuracy [75].

If a suitable performance measure is unavailable, it is a good indication that the task or problem formulation themselves have not

gained general acceptance, or has fallen out of general acceptance. This, on its own, is reasonable cause for contestation of the output of the tools. Take for instance the hypothetical case of a prosecution team that wishes to introduce evidence based on a risk score that rates the “latent criminality” of a person based on facial images. As the task lacks a scientific basis, there is no performance measure for the task that has gained “general acceptance” by the relevant scientific communities.

On a more technical note, the performance check abstracts out issues of statistical significance. If the test distributions chosen by the defense correspond to a sub-population that is too small to estimate the performance with high confidence, then the check fails.

### 5.3 Relationship to Existing Lines of Work

Our proposal builds upon and benefits from recent lines of work in several different communities across law, forensic science, and machine learning.

*Statistical evidence literature.* There is a long line of work in the legal community discussing the use of statistical evidence in court [24, 34]. Closely related to our paper, Faigman et al. [43] discuss difficulties that arise when courts apply statistical statements to individuals. The authors call this step “group to individual” (G2i) inference and suggest standards of evidence for such inferences. At a high level, our concerns with statistical evidence are similar, but our perspective on the role of software differs markedly: Faigman et al. [43] consider risk scores from a risk assessment tool as sufficient to establish error rates for evidence applied to individuals. In contrast, we question the validity of such risk scores and propose adversarial testing of statistical software as a way to test the risk assessment tool in question.

Another line of work in the statistical evidence literature considers the subtleties involved in choosing an appropriate “reference class” when making probability estimates about an individual case. Technically, this corresponds to choosing which variables to the condition when fitting a statistical model. Allen and Pardo [12] argue the difficulties involved in choosing an appropriate reference class limit the value of statistical models in court. Cheng [23] suggests using ideas from model selection to address the reference class selection problem. While choosing a reference class bears some similarities to choosing our distribution family, these problems differ in an important aspect: The reference class problem is about how to build probability models, e.g., by appropriately selecting the features in a regression, whereas we are concerned with using the distribution family to validate already-fit models in software.

*Robust machine learning.* Our proposal benefits from close ties to recent lines of work in machine learning. The first is a close connection to a robustness notion that originated in work on algorithmic fairness, but has since found numerous other applications. Hebert-Johnson et al. [56] argue that existing fairness criteria are too weak insofar as they only hold in the population as a whole but not in rich families of sub-populations. This observation motivated the notion of multi-calibration and multi-accuracy that require a statistical model to achieve calibration and accuracy, respectively, in a given family of sub-populations. Subsequent work by Dwork

et al. [40] introduced a closely related notion called *outcome indistinguishability* with an emphasis on the perspective of reasoning about statistical models on individual instances. It is not difficult to show that our notion of robust adversarial testing generalizes multi-calibration and multi-accuracy at a formal level. But it also shares normative content. All notions aim to strengthen broad average measures of performance with stronger guarantees that come closer to providing meaningful guarantees for individual instances.

Fundamentally, robust adversarial testing requires a tool to perform well on relevant distributions related to the reference population. This requirement accounts for a central failure point of statistical machinery. An increasingly rich literature in the machine learning community has documented the failure of statistical models to generalize from one domain to another closely related domain [17, 71, 98, 101, 116, 127].

#### 5.4 Discussion and Limitations

Our definition establishes an important bridge between legal practice and machine learning research on distributional robustness. This connection provides a beneficial connection in both directions. On the one hand, criminal defense advocates can more effectively challenge tools using insights from lines of investigation that exhibit failure points of existing machine learning techniques. On the other hand, as machine learning experts work to make tools more robust to changes in distributional contexts, we expect that more well-designed evidentiary statistical software will hold up against defense challenges by our standard.

There remains a significant burden on the defense in at least two regards. First, our framework requires the tool to solve a well-defined statistical problem. If the problem formulation underlying the software tool is flawed, this on its own should become the primary point of contestation. Defense counsel must therefore be able to identify software that fails with regard to problem formulation. Difficulty in choosing a well-defined data universe, distribution family, or performance measure is a helpful indicator that a problem formulation may be invalid. These are, however, imperfect indicators that may miss other challenges in the problem formulation [9, 90]. For instance, scholars argue that risk assessment tools conflate different notions of risk. Namely, they can conflate missing a court appointment with absconding from the jurisdiction or violently harming another person before a trial can be held [9, 70, 90]. Moreover, in some cases the problem statement may lack a scientific basis entirely [110, 111].

Second, our proposal leaves open aspects of executing our definition in context. There may be disagreement on what input and performance measures are relevant. The defense may not have access to relevant test cases or relevant data may be unavailable. The lack of representative data in existing databases may itself indicate the tool has not been sufficiently validated for the case at hand. The right level of adversarialism is also important. A defense that is too adversarial might choose a set of guaranteed failure cases, in which case the prosecution would likely argue that this choice was unreasonable. Our framework does not resolve what happens when the prosecution contests the choices of the defense.

While these are important challenges, attempts to apply our framework would provide valuable transparency about where concretely the process of adversarial scrutiny ran into difficulty. These challenges also point to fruitful research directions. Additional research can help in identifying misapplications of statistics, as well as facilitate the choice of parameters in our definition.

## 6 CONFRONTING STRUCTURAL AND INSTITUTIONAL CHALLENGES

The use of evidentiary software without adequate scrutiny distorts the truth-seeking process and undermines fairness in the U.S. criminal legal system. In this work, we propose robust adversarial testing as a technical framework to empower criminal defense counsel to adversarially scrutinize evidence output by statistical software. This framework provides an approach to mitigate inadequate validity testing and cross-examination of forensic evidence. In doing so, we take a step towards formalizing an important aspect of audits of evidentiary statistical software.

Our framework can apply broadly throughout the life-cycle of a case, from plea negotiations to admissibility hearings to trial and beyond. It can be useful to the defense regardless of the varying legal standards that apply at each stage of the case, such as the more lenient threshold for admissibility or the more onerous standard of proof at trial. Further, while we focus on statistical software used as evidence, our approach could also apply to statistical software used for other criminal legal purposes. For instance, our framework could be used to scrutinize face matching software that police rely on during investigations, or to scrutinize risk assessment instruments that decision makers use to predict recidivism during pretrial bail hearings, sentencing, and parole.

While our work is focused on statistical software, we also believe that such frameworks could be provided for non-statistical software, such as hash-matching algorithms used to identify contraband files or malware tools used to conduct remote investigations of networked digital devices. We encourage academic research and relevant technical communities to investigate and formulate similar adversarial frameworks in such cases.

Notwithstanding the potential for this framework to have broad applicability, a host of structural and institutional barriers challenge the implementation of our framework in practice. Without simultaneously tackling these barriers, no framework can address the real needs of defense counsel or lead to long-lasting changes. Below, we highlight some of these challenges that currently create a barrier for proper evaluation of evidentiary software. We focus on two key axes relevant to our framework: (1) defense lack of access to software and data, and (2) defense lack of access to expert witnesses. We then discuss proposals for mitigating these barriers and highlight avenues for further consideration.

*Access to Software and Data.* Vendors of statistical software can severely limit access to the software, for instance by citing concerns around intellectual property or by using contract law to bar independent audits. Specifically, multiple vendors have asserted trade secret law to block criminal defense subpoenas and discovery requests for source code and other information about how their software works [62, 100, 107, 123]. Meanwhile, some vendors also refuse to grant research licenses to independent academic scientists

who wish to evaluate their tools, even while claiming to judges across the country that their tools are subject to peer review and hence should be admissible as evidence in court [58]. This practice runs directly contrary to scientific values of peer review and reproducibility, and presents one explanation for the dearth of independent peer-reviewed validation studies of evidentiary statistical software [95].

Even when vendors do offer defense counsel some access to aspects of their software, they may demand onerous terms and conditions that make the access largely ineffective for performing meaningful assessments. For instance, in one case involving a popular probabilistic genotyping software, TrueAllele [35], the vendor initially demanded that the defense team only access the source code using a read-only iPad and only use pen and paper to take notes. Only after extensive lawyering by a coalition of skilled attorneys from across the country—a luxury that most criminal defendants lack—was the defense team able to obtain more reasonable terms of access [67, 114].

We echo the prior arguments in Wexler [123] that, at minimum, judges should not extend trade secret law to block criminal defense access to relevant information about how evidentiary software works. Regardless of the precise information the defense is seeking to obtain—source code, validation studies, executable versions of software programs, etc.—vendors' intellectual property interests can be protected to the full extent reasonable by ordering disclosure subject to appropriate protective orders. Trade secret law offers no good reason to withhold such information from the defense.<sup>7</sup>

Broadly, vendor-imposed limits on defense access to evidentiary software reflect more widespread barriers to criminal defense access to data. For example, national DNA databases are available to law enforcement to search for evidence of guilt, but are often unavailable to defense counsel to search for evidence of innocence [46]. The same is true for fingerprint [47] and face matching databases [79]. Judges have interpreted federal law to make private sector digital communications content databases available to law enforcement but categorically barred to the defense [124]. We argue that these access limitations should be reversed; There are no just reasons to categorically deny defense counsel access to relevant data that they need to investigate forensic evidence in their case.

Without access to data to run evidentiary statistical software tools, defendants are at a disadvantage in both testing the software to expose flaws in the government's evidence of guilt, and in using the software to surface evidence of innocence. For adversarial scrutiny to succeed, defendants need access to relevant evidence, whether that means queries of forensic databases, software source code, training data, or executables to run and test the software.

*Access to Technical Experts.* Even if defense counsel manage to obtain meaningful access to software and data, they face another substantial barrier to performing adversarial scrutiny: access to experts. Defense counsel themselves are typically not trained to do the kind of technical work required to scrutinize statistical evidentiary software, and most defense offices lack the resources to support specialized attorneys or staff with these skills. Hence, most

defense counsel will need to depend on outside experts. Yet, while the Supreme Court has held, at minimum, that indigent defendants in capital cases have a due process right to state-funded psychiatric experts, [10] case law remains ambiguous as to whether indigent defendants have a right to expert assistance in other contexts [52].

Further, the structure of the forensic science software market tilts the pool of potential experts in favor of the prosecution, since law enforcement and prosecution entities are the primary paying customers—if not the developers themselves—and private developers build tools primarily to serve law enforcement needs [100, 107]. Specifically, employees of software vendors are ideally positioned to gain expertise as to their own products, and sometimes testify about the strengths of their tools on behalf of the prosecution.<sup>8</sup> However, these employees can hardly be expected to volunteer information about the weaknesses of their tools for use by the defense. To fill this gap, defense counsel may seek to rely on academic experts. But here again they face obstacles. Defense counsel have limited opportunities to identify the right academics and to provide them with sufficient incentives and training. In recent years, there have been efforts, such as the nonprofit, PDQuery [91], which seek to address this problem by providing a matching service to connect technical expert volunteers with public defenders. More innovations of this sort are needed at scale.

In sum, defense counsel now confront a growing number and diversity of evidentiary statistical software used in an increasing number of cases without access to the technical expertise they need to adequately evaluate these tools.

## A Way Forward

Beyond the reforms mentioned above—specifically, eliminating trade secret barriers for criminal defense access to relevant evidence, ensuring defense access to relevant data, and scaling existing efforts to match criminal defense counsel with volunteer academic experts—we offer three additional proposals to help mitigate the structural and institutional barriers to implementing our framework for adversarial scrutiny of evidentiary statistical software.

First, our adversarial scrutiny framework provides an opportunity for the computational academic community to contribute in ways beyond serving as experts in individual cases. Namely, academic scientists could create, sign, and publicly distribute generic affidavits that are ready to file in court and that attest to known issues from existing research and the current state-of-the-art for various evidentiary statistical software tools. The affidavits could support the selection of the input distribution and performance check in our framework using existing research and known limitations of the task that the software purports to perform. The affidavits could also do the challenging work of translating expert knowledge to make it accessible to generalist judges and lay jurors. In short, these affidavits could provide an off-the-shelf resource for defense counsel who are scrutinizing the software in any specific case.

Second, academic scientists could create off-the-shelf testing tools that would allow for checking for different cases in a more streamlined way, requiring less expertise and time from the defense.

<sup>7</sup>Lawmakers have begun taking important steps to limit vendors' ability to assert trade secrets to block criminal defense scrutiny. See, for instance, the Justice in Forensic Algorithms Act of 2021, H. R. 2438 (2021).

<sup>8</sup>See, e.g., <https://www.cybgcn.com/information/newsroom/2016/apr/files/B-Perlin-second-declaration.pdf>.

Rather than starting from scratch for each case, defense counsel could instead use an existing tool corresponding to a particular evidentiary software to enter the most relevant input to their case at hand and test for the tools' performance. An important component of such testing tools enabling adversarial scrutiny would also be comprehensive databases of input (e.g., DNA samples) so that the defense can easily select relevant subsets of the data to stress test the statistical software in question.

Importantly, our proposal for scientific communities to provide off-the-shelf affidavits and testing tools does not serve as a substitute for defense offices having the appropriate expertise and time to conduct adequate scrutiny, but aims to ease the burden on these offices by sharing information that is currently present in scientific communities.

Third, we advocate for government funding for public defender offices to hire specialized technical staff who can perform robust adversarial testing in-house. The Digital Forensics Unit (DFU) of The Legal Aid Society of New York City offers a ready model. The DFU integrates technologists and defense attorneys in a single unit to enable coordinated advocacy across computer science and law [115].

Prior commentators have emphasized the importance of independent audit and oversight bodies to help ensure the fair and accurate use of new AI and other software products in high stakes settings such as the criminal legal system. We agree that such independent audit and oversight bodies are essential to improve the quality of forensic evidence in the courts. This includes government oversight bodies that take a neutral role, such as the National Institute of Standards and Technology and the recently disbanded National Commission on Forensic Science [85].

Yet, oversight committees tasked with *neutral* evaluations of forensic science methods are not ideally situated to perform the kinds of stress-tests we envision for *adversarial* scrutiny of evidentiary statistical software. Rather, public defender offices provide a natural home for technical staff to perform adversarial scrutiny because these offices are already committed to zealous advocacy on behalf of the accused. We hope that placing technical auditors inside public defender offices will protect against industry capture. Further, by placing technical staff inside public defender offices, we directly empower defense lawyers and their clients, who often belong to the most marginalized communities that are disproportionately harmed by faulty forensics [109].

## ACKNOWLEDGMENTS

JM was supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE 1752814. The authors declare no other financial interests.

The authors would like to thank the following individuals for invaluable discussions and comments on previous drafts of this work: Pieter Abbeel, Nathan Adams, Steven M. Bellovin, Matthew Boulos, Alex Carney, Edward Cheng, A. Feder Cooper, Fernando Delgado, Brandon L. Garrett, Jonah Gelbach, Andreas Alexander Haupt, Jon Kleinberg, Logan Koepke, Beryl Lipton, Deborah Raji, Andrea Roth, Christopher Strong, Ramon Vilarino, Dana Yeo, Harlan Yu, as well as members of the Berkeley Algorithmic Fairness and Opacity Group (AFOG) and the Berkeley Equity and Access, in

Algorithms, Mechanisms, and Optimization (BEAAMO) research group. The authors would also like to thank the anonymous reviewers for their feedback. This work would not have been possible without insightful discussions with Richard Torres and Clinton Hughes from the Brooklyn Defenders and Queens Defenders offices, Elizabeth Vasquez from Legal Aid, as well as Nathan Adams and Dan Krane from Forensic Bioinformatics. The authors also received substantial reference support from Dean Rowan at the Berkeley Law Library.

## REFERENCES

- [1] 2020. Survey Shows STRmix Has Been Used in 220,000 Cases Worldwide. <https://www.strmix.com/news/survey-shows-strmix-has-been-used-in-220000-cases-worldwide/>. Accessed: 2022-01-21.
- [2] 2021. H.R.2438 - 117th Congress (2021-2022): Justice in Forensic Algorithms Act of 2021. Congress.gov, Library of Congress. <https://www.congress.gov/bills/117/congress-house-bill/2438/titles>.
- [3] Fed. R. Evid. 401.
- [4] Fed. R. Evid. 402.
- [5] Fed. R. Evid. 403.
- [6] Fed. R. Evid. 702.
- [7] Fed. R. Evid. 902.
- [8] U.S. Const. Amend. VI.
- [9] Rediet Abebe, Solon Barocas, Jon Kleinberg, Karen Levy, Manish Raghavan, and David G Robinson. 2020. Roles for computing in social change. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 252–260.
- [10] Ake v. Oklahoma, 470 U.S. 68. (1985).
- [11] Lauren E Alfonse, Amanda D Garrett, Desmond S Lun, Ken R Duffy, and Catherine M Grgicak. 2018. A large-scale dataset of single and mixed-source short tandem repeat profiles to inform human identification strategies: PROVEDIt. *Forensic Science International: Genetics* 32 (2018), 62–70.
- [12] Ronald J Allen and Michael S Pardo. 2007. The problematic value of mathematical models of evidence. *The Journal of Legal Studies* 36, 1 (2007), 107–140.
- [13] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2016. Machine bias. *ProPublica*, May 23, 2016 (2016), 139–159.
- [14] Steven M Bellovin, Matt Blaze, Susan Landau, and Brian Owsley. 2021. Seeking the Source: Criminal Defendants' Constitutional Right to Source Code. *Ohio State Technology Law Journal* 17 (2021), 1.
- [15] Aharon Ben-Tal, Laurent El Ghaoui, and Arkadi Nemirovski. 2009. *Robust Optimization*. Princeton University Press, Princeton and Oxford. <https://doi.org/10.1515/9781400831050>
- [16] David E Bernstein and Eric G Lasker. 2015. Defending Daubert: it's time to amend federal rule of evidence 702. *William & Mary Law Review* 57 (2015), 1.
- [17] Battista Biggio and Fabio Roli. 2018. Wild Patterns: Ten Years After the Rise of Adversarial Machine Learning. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. <https://doi.org/10.1145/3243734.3264418>
- [18] *Bullock v. New Mexico*, 564 U.S. 647. (2011).
- [19] Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*. PMLR, 77–91.
- [20] John M Butler, Hari Iyer, Rich Press, Melissa Taylor, Peter M Vallone, and Sheila Willis. 2021. DNA Mixture Interpretation: A NIST Scientific Foundation Review. *National Institutes of Standards and Technology* (2021).
- [21] Marc Canellas. 2021. Defending IEEE Software Standards in Federal Criminal Court. *Computer* 54, 6 (2021), 14–23.
- [22] Celebrite Inc. [n. d.]. Celebrite UFED. <https://celebrite.com/en/ufed>. Accessed: 2022-05-31.
- [23] Edward K. Cheng. 2009. A practical solution to the reference class problem. *Columbia Law Review* 109 (2009), 2081.
- [24] Edward K Cheng, David L Faigman, Jennifer L Mnookin, Erin E Murphy, Joseph Sanders, and Christopher Slobogin. 2021. *Modern scientific evidence: the law and science of expert testimony*. Thomson West.
- [25] Edward K Cheng and G Alexander Nunn. 2018. Beyond the witness: Bringing a process perspective to modern evidence law. *Texas Law Review* 97 (2018), 1077.
- [26] Edward K Cheng and Albert H Yoon. 2005. Does Frye or Daubert matter? A study of scientific admissibility standards. *Virginia Law Review* (2005), 471–513.
- [27] Christian Chessman. 2017. A "Source" of Error: Computer Code, Criminal Defendants, and the Constitution. *California Law Review* (2017), 179–228.
- [28] Alexandra Chouldechova. 2017. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big data* 5, 2 (2017), 153–163.
- [29] CMI Inc. [n. d.]. Intoxilyzer 8000. <https://www.alcoholtest.com/intoxilyzer-8000/>. Accessed: 2022-05-31.

- [30] Michael D Coble and Jo-Anne Bright. 2019. Probabilistic genotyping software: an overview. *Forensic Science International: Genetics* 38 (2019), 219–224.
- [31] A Feder Cooper, Benjamin Laufer, Emanuel Moss, and Helen Nissenbaum. 2022. Accountability in an Algorithmic Society: Relationality, Responsibility, and Robustness in Machine Learning. *arXiv preprint arXiv:2202.05338* (2022).
- [32] A Feder Cooper, Yucheng Lu, Jessica Forde, and Christopher M De Sa. 2021. Hyperparameter Optimization Is Deceiving Us, and How to Stop It. *Advances in Neural Information Processing Systems* 34 (2021).
- [33] Sam Corbett-Davies and Sharad Goel. 2018. The measure and mismeasure of fairness: A critical review of fair machine learning. *arXiv preprint arXiv:1808.00023* (2018).
- [34] National Research Council. 2011. *Reference Manual on Scientific Evidence: Third Edition*. The National Academies Press. <https://nap.nationalacademies.org/catalog/13163/reference-manual-on-scientific-evidence-third-edition>
- [35] Cybergentics. [n. d.]. TrueAllele. <https://www.cybgen.com>. Accessed: 2022-05-31.
- [36] Daubert v. Merrell Dow Pharmaceuticals, Inc., 509 U.S. 579. (1993).
- [37] Fernando Delgado, Solon Barocas, and Karen Levy. 2022. An Uncommon Task: Participatory Design in Legal AI. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW1 (2022), 1–23.
- [38] Dräger. [n. d.]. Dräger Alcotest 9510. [https://www.draeger.com/en-us\\_us/Products/Alcotest-9510](https://www.draeger.com/en-us_us/Products/Alcotest-9510). Accessed: 2022-05-31.
- [39] Julia Dressel and Hany Farid. 2018. The accuracy, fairness, and limits of predicting recidivism. *Science Advances* 4, 1 (2018), eaao5580.
- [40] Cynthia Dwork, Michael P Kim, Omer Reingold, Guy N Rothblum, and Gal Yona. 2021. Outcome indistinguishability. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. ACM, New York, NY, USA, 1095–1108.
- [41] Laurel Eckhouse, Kristian Lum, Cynthia Conti-Cook, and Julie Ciccolini. 2019. Layers of bias: A unified approach for understanding problems with risk assessment. *Criminal Justice and Behavior* 46, 2 (2019), 185–209.
- [42] David Freeman Engstrom and Jonah B Gelbach. 2020. Legal Tech, Civil Procedure, and the Future of Adversarialism. *University of Pennsylvania Law Review* 169 (2020), 1001.
- [43] David L. Faigman, John Monahan, and Christopher Slobogin. 2014. Group to Individual (G2i) Inference in Scientific Expert Testimony. *The University of Chicago Law Review* 81, 2 (2014), 417–480.
- [44] *Frye v. U.S.*, 315 F.2d 491. (1963).
- [45] Brandon L Garrett. 2016. Constitutional regulation of forensic evidence. *Washington & Lee Law Review* 73 (2016), 1147.
- [46] Brandon L Garrett. 2021. *Autopsy of a Crime Lab*. University of California Press, 189.
- [47] Brandon L Garrett. 2021. *Autopsy of a Crime Lab*. University of California Press, 192.
- [48] Brandon L Garrett and M Chris Fabricant. 2017. The Myth of the Reliability Test. *Fordham Law Review* 86 (2017), 1559.
- [49] Brandon L Garrett and Peter J Neufeld. 2009. Invalid forensic science testimony and wrongful convictions. *Virginia Law Review* (2009), 1–97.
- [50] Nancy Gertner. 2010. Commentary on the Need for a Research Culture in the Forensic Sciences. *UCLA Law Review* 58 (2010), 789.
- [51] Paul C Giannelli. 1991. Criminal discovery, scientific evidence, and DNA. *Vanderbilt Law Review* 44 (1991), 791.
- [52] Paul C Giannelli. 2003. Ake v. Oklahoma: The right to expert assistance in a post-Daubert, post-DNA world. *Cornell Law Review* 89 (2003), 1305.
- [53] Grayshift. [n. d.]. GrayKey Cell Phone Forensics Tool. <https://www.grayshift.com/graykey/>. Accessed: 2022-05-31.
- [54] Jennifer L Groscup, Steven D Penrod, Christina A Studebaker, Matthew T Huss, and Kevin M O’Neil. 2002. The effects of Daubert on the admissibility of expert testimony in state and federal criminal cases. *Psychology, Public Policy, and Law* 8, 4 (2002), 339.
- [55] Moritz Hardt and Benjamin Recht. 2021. *Patterns, predictions, and actions: A story about machine learning*. <https://mlstory.org>. arXiv:2102.05242 [cs.LG]
- [56] Ursula Hebert-Johnson, Michael Kim, Omer Reingold, and Guy Rothblum. 2018. Multicalibration: Calibration for the (Computationally-Identifiable) Masses. In *Proceedings of the 35th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 80)*, Jennifer Dy and Andreas Krause (Eds.). PMLR, 1939–1948.
- [57] Jim Hilbert. 2018. The disappointing history of science in the courtroom: Frye, Daubert, and the ongoing crisis of junk science in criminal trials. *Oklahoma Law Review* 71 (2018), 759.
- [58] Karen L Howard. 2021. Forensic Technology: Algorithms Strengthen Forensic Analysis, but Several Factors Can Affect Outcomes. GAO Report No. GAO-21-435SP. <https://www.gao.gov/assets/gao/21-435sp.pdf>
- [59] Aziz Z Huq. 2018. Racial equity in algorithmic criminal justice. *Duke LJ* 68 (2018), 1043.
- [60] Edward J Imwinkelried. 2016. Computer source code: a source of the growing controversy over the reliability of automated forensic techniques. *DePaul Law Review* 66 (2016), 97.
- [61] Institute of Environmental Science and Research Ltd. [n. d.]. STRMix. <https://www.strmix.com/strmix/>. Accessed: 2022-05-31.
- [62] Sonia Katyal and Charles Graves. 2021. From Trade Secrecy to Seclusion. *Georgetown Law Journal, Forthcoming* (2021).
- [63] David H Kaye. 2011. *The new Wigmore: a treatise on evidence: expert evidence*. Wolters Kluwer Law & Business.
- [64] David H Kaye. 2017. How Daubert and its progeny have failed criminalistics evidence and a few things the judiciary could do about it. *Fordham Law Review* 86 (2017), 1639.
- [65] Danielle Leah Kehl and Samuel Ari Kessler. 2017. Algorithms in the criminal justice system: Assessing the use of risk assessments in sentencing. (2017).
- [66] Lauren Kirchner. 2017. Thousands of Criminal Cases in New York Relied on Disputed DNA Testing Techniques. <https://www.propublica.org/article/thousands-of-criminal-cases-in-new-york-relied-on-disputed-dna-testing-techniques>. Accessed: 2022-04-17.
- [67] Lauren Kirchner. 2021. Powerful DNA Software Used in Hundreds of Criminal Cases Faces New Scrutiny. *The Markup* (Mar 2021). <https://themarkup.org/news/2021/03/09/powerful-dna-software-used-in-hundreds-of-criminal-cases-faces-new-scrutiny>
- [68] Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan. 2016. Inherent trade-offs in the fair determination of risk scores. *arXiv preprint arXiv:1609.05807* (2016).
- [69] Allison Koenecke, Andrew Nam, Emily Lake, Joe Nudell, Minnie Quartey, Zion Mengesha, Connor Toups, John R Rickford, Dan Jurafsky, and Sharad Goel. 2020. Racial disparities in automated speech recognition. *Proceedings of the National Academy of Sciences* 117, 14 (2020), 7684–7689.
- [70] John Logan Koepke and David G Robinson. 2018. Danger ahead: Risk assessment and the future of bail reform. *Washington Law Review*. 93 (2018), 1725.
- [71] Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiko Yasunaga, Richard Lanus Phillips, Irena Gao, Tony Lee, Etienne David, Ian Stavness, Wei Guo, Berton Earnshaw, Imran Haque, Sara M. Beery, Jure Leskovec, Anshul Kundaje, Emma Pierson, Sergey Levine, Chelsea Finn, and Percy Liang. 2021. WILDs: A Benchmark of in-the-Wild Distribution Shifts. In *International Conference on Machine Learning (ICML)*.
- [72] Joshua A Kroll, Joanna Huey, Solon Barocas, Edward W Felten, Joel R Reidenberg, David G Robinson, and Harlan Yu. 2017. Accountable algorithms. *University of Pennsylvania Law Review* 165 (2017), 633.
- [73] Katherine Kwong. 2017. The Algorithm says you did it: The use of Black Box Algorithms to analyze complex DNA evidence. *Harv. JL & Tech.* 31 (2017), 275.
- [74] Stephanie J Lacabra, Jeanna Matthews, and K Walsh. 2018. Opening the black box: Defendants’ rights to confront forensic software. *The Champion* 5 (2018).
- [75] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 11 (1998), 2278–2324.
- [76] Thomas Liao, Rohan Taori, Inioluwa Deborah Raji, and Ludwig Schmidt. 2021. Are We Learning Yet? A Meta Review of Evaluation Failures Across Machine Learning. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.
- [77] Lifeloc Technologies. [n. d.]. LX9 Breathalyzer. <https://lifeloc.com/lx9>. Accessed: 2022-05-31.
- [78] Magnet Forensics. [n. d.]. Magnet AXIOM. <https://www.magnetforensics.com/products/magnet-axiom/>. Accessed: 2022-05-31.
- [79] Aaron Mak. 2019 [Online]. Facing Facts: A case in Florida demonstrates the problems with using facial recognition to identify suspects in low-stakes crimes. *Slate* (Jan. 25, 2019 [Online]). <https://slate.com/technology/2019/01/facial-recognition-arrest-transparency-willie-allen-lynch.html>
- [80] Jeanna Matthews, Marzieh Babaeianjelodar, Stephen Lorenz, Abigail Matthews, Mariama Njie, Nathaniel Adams, Dan Krane, Jessica Goldthwaite, and Clinton Hughes. 2019. The right to confront your accusers: Opening the black box of forensic DNA software. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*. 321–327.
- [81] Jeanna Neeffe Matthews, Graham Northup, Isabella Grasso, Stephen Lorenz, Marzieh Babaeianjelodar, Hunter Bashaw, Sumona Mondal, Abigail Matthews, Mariama Njie, and Jessica Goldthwaite. 2020. When trusted black boxes don’t agree: Incentivizing iterative improvement and accountability in critical software systems. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*. 102–108.
- [82] Sandra G Mayson. 2018. Bias in, bias out. *Yale Law Journal* 128 (2018), 2218.
- [83] *Melendez-Diaz v. Massachusetts*, 557 U.S. 305. (2009).
- [84] Jennifer L Mnookin. 2014. Repeat Play Evidence: Jack Weinstein, Pedagogical Devices, Technology, and Evidence. *DePaul Law Review* 64 (2014), 571.
- [85] Jennifer L Mnookin. 2018. The uncertain future of forensic science. *Daedalus* 147, 4 (2018), 99–118.
- [86] Jennifer L Mnookin, Simon A Cole, Itiel E Dror, and Barry AJ Fisher. 2010. The need for a research culture in the forensic sciences. *UCLA Law Review* 58 (2010), 725.

- [87] Charles J Ogletree Jr. 1992. Beyond justifications: Seeking motivations to sustain public defenders. *Harvard Law Review* 106 (1992), 1239.
- [88] Osonde A Osoba and William Welsler IV. 2017. *An intelligence in our image: The risks of bias and errors in artificial intelligence*. RAND Corporation.
- [89] Frank Pasquale. 2015. The black box society. In *The Black Box Society*. Harvard University Press.
- [90] Samir Passi and Solon Barocas. 2019. Problem formulation and fairness. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. 39–48.
- [91] PDQuery. [n. d.]. <https://pdquery.org> Accessed: 2022-05-09.
- [92] *People v. Lopez*, 55 Cal.4th 569. (2012).
- [93] Mark W Perlin. 2018. When DNA is not a gold standard: failing to interpret mixture evidence. *The Champion* 42 (2018), 50–6.
- [94] Riana Pfefferkorn. 2019. “Deepfakes” in the Courtroom. *Boston University Public Interest Law Journal* 29 (2019), 245.
- [95] President’s Council of Advisors on Science and Technology. 2016. Forensic Science in Criminal Courts: Ensuring Scientific Validation of Feature-Comparison Methods. *President’s Council of Advisors on Science and Technology* (2016).
- [96] President’s Council of Advisors on Science and Technology. 2017. *Addendum to the PCAST Report on Forensic Science in Criminal Courts*. President’s Council of Advisors on Science and Technology.
- [97] Innocence Project. [n. d.]. Overturning Wrongful Convictions Involving Misapplied Forensics. <https://innocenceproject.org/overturning-wrongful-convictions-involving-flawed-forensics/>. Accessed: 2022-05-09.
- [98] Joaquin Quiñero-Candela, Masashi Sugiyama, Neil D Lawrence, and Anton Schwaighofer. 2009. *Dataset shift in machine learning*. MIT Press.
- [99] Inioluwa Deborah Raji and Joy Buolamwini. 2019. Actionable auditing: Investigating the impact of publicly naming biased performance results of commercial ai products. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*. 429–435.
- [100] Natalie Ram. 2017. Innovating criminal justice. *Northwestern University Law Review* 112 (2017), 659.
- [101] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. 2019. Do imagenet classifiers generalize to imagenet?. In *International Conference on Machine Learning*. PMLR, 5389–5400.
- [102] William Robinson. 2021. RFC Response: NIST Internal Report 8351-DRAFT DNA Mixture Interpretation: A NIST Scientific Foundation Review. <https://ieeusa.org/assets/public-policy/policy-log/2021/111821.pdf>
- [103] Andrea Roth. 2016. Machine testimony. *Yale Law Journal* 126 (2016), 1972.
- [104] Safety Dynamics Inc. [n. d.]. Safety Dynamics - Gunshot Detection - Gunshot Location System - Products. <http://www.safetydynamics.net/prods.html>. Accessed: 2022-05-31.
- [105] Stephen A Saltzburg, Michael M Martin, and Daniel J Capra. 2006. *Federal Rules of Evidence Manual*. Vol. 5. LEXIS Law Publishing.
- [106] ShotSpotter Inc. [n. d.]. ShotSpotter. <https://www.shotspotter.com/law-enforcement/gunshot-detection-technology/>. Accessed: 2022-05-31.
- [107] Eli Siems, Katherine J Strandburg, and Nicholas Vincent. 2022. Trade Secrecy and Innovation in Forensic Technology. *Hastings Law Journal* 73, 3 (2022).
- [108] William Singer and Kendra K Albert. 2020. Brief of Amicus Curiae Upturn, Inc. In Support of Movant-Appellant. <https://clinic.cyber.harvard.edu/files/2020/10/Upturn-Amicus-Brief-for-Filing-10-15-20-115-pm.pdf>. Accessed: 2022-01-21.
- [109] Maneka Sinha. 2021. Radically Reimagining Forensic Evidence. *Alabama Law Review, Forthcoming, U of Maryland Legal Studies Research Paper* 2021-10 (2021).
- [110] Mona Sloane. 2021. The Algorithmic Auditing Trap. *OneZero (blog)*, March 17 (2021).
- [111] Mona Sloane, Emanuel Moss, and Rumman Chowdhury. 2022. A Silicon Valley love triangle: Hiring algorithms, pseudo-science, and the quest for auditability. *Patterns* 3, 2 (2022), 100425.
- [112] Christopher Slobogin. 2021. *Just Algorithms: Using Science to Reduce Incarceration and Inform a Jurisprudence of Risk*. Cambridge University Press.
- [113] Abbe Smith. 1999. Defending Defending: The Case for Unmitigated Zeal on Behalf of People Who Do Terrible Things. *Hofstra Law Review* 28 (1999), 925.
- [114] *State v. Pickett*, 246 A. 3d 279. (2021).
- [115] The Legal Aid Society. [n. d.]. Digital Forensics Unit - The Legal Aid Society. <https://legalaidnyc.org/programs-projects-units/digital-forensics-unit/>. Accessed: 2022-05-09.
- [116] Antonio Torralba and Alexei A Efros. 2011. Unbiased look at dataset bias. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [117] Ultra Electronics Forensic Technology. [n. d.]. Correlation Engine. <https://www.ultra-forensictechnology.com/en/our-products/ballistic-identification/correlation-server>. Accessed: 2022-05-31.
- [118] Ultra Electronics Forensic Technology. [n. d.]. IBIS Matchpoint. <https://www.ultra-forensictechnology.com/en/our-products/ballistic-identification/matchpoint>. Accessed: 2022-05-31.
- [119] *U.S. v. Lizarraga-Tirado*, 789 F.3d 1107. (2015).
- [120] *U.S. v. Vayner*, 769 F.3d 125. (2014).
- [121] Abraham Wald. 1945. Statistical Decision Functions Which Minimize the Maximum Risk. *Annals of Mathematics* (1945).
- [122] Emily West and Vanessa Meterko. 2015. Innocence project: DNA exonerations, 1989-2014: review of data and findings from the first 25 years. *Albany Law Review* 79 (2015), 717.
- [123] Rebecca Wexler. 2018. Life, liberty, and trade secrets: Intellectual property in the criminal justice system. *Stanford Law Review* 70 (2018), 1343.
- [124] Rebecca Wexler. 2020. Privacy as Privilege: The Stored Communications Act and Internet Evidence. *Harvard Law Review* 134 (2020), 2721.
- [125] Rebecca Wexler. 2021. Privacy Asymmetries: Access to Data in Criminal Investigations. *UCLA Law Review* 68, 1 (2021).
- [126] *Williams v. Illinois*, 567 U.S. 50. (2012).
- [127] John R Zech, Marcus A Badgeley, Manway Liu, Anthony B Costa, Joseph J Titano, and Eric Karl Oermann. 2018. Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: a cross-sectional study. *PLoS medicine* 15, 11 (2018).

## A EXAMPLE EVIDENTIARY SOFTWARE

In Table 1, we present examples of different categories of evidentiary software as well as examples of known products for each such category.

## B THE RULES OF EVIDENCE THAT APPLY TO SOFTWARE OUTPUT

Introducing software output as evidence against a criminal defendant at trial can implicate several legal rules that control the admissibility of evidence. This section presents an explanation of those rules for a general audience.<sup>9</sup> Notably, the overall thrust of the rules favors lenient admissibility, meaning more evidence will be admitted and thus seen by the jury rather than excluded by the judge and kept away from the jury. This overall thrust reflects a policy preference for trusting juries to weigh evidence in reaching their ultimate judgment in the case, rather than empowering the judge to weigh the evidence in advance and keep information secret from the jury merely because the judge believes is unworthy of consideration. The lenient thrust of the rules also reflects a policy choice to rely on the adversarial process of cross-examination by lawyers during trial to educate the jury about the strengths and weaknesses of each piece of admitted evidence to enable the jury to perform that weighing function. While those two policy preferences operate in the background of most if not all of the rules, there are exceptions that increase the judge's gatekeeping role for information that a jury might be less equipped to understand.

### B.1 Relevance and Prejudice

At the most general level, the rules for relevance and prejudice apply to software output. Like all other evidence, the judge must exclude software output if they are irrelevant to the case [3, 4]. Additionally, the judge has the discretion to exclude the output if their probative value would be "substantially outweighed" by a risk of unfair prejudice, confusing or misleading the jury, or wasting time [5]. Resource disparities between the parties are a potential source of prejudice. For instance, consider demonstrative models that simulate a party's theory of the case. A flashy CGI simulation depicting how an automobile crash could have occurred had the defendant been at fault might have a substantial persuasive effect on the jury. Meanwhile, the defendant might lack the resources in money, time, or expertise to produce an alternative CGI simulation depicting how the same crash could have occurred had the plaintiff been at fault. To correct the risk of one-sided persuasive force in such circumstances, Law Professor and U.C.L.A. Law School Dean Jennifer Mnookin has argued that parties should have a right to test their adversaries' demonstrative models with alternative input and present the results to the jury as a simulation of an alternative theory of the case [84]. To date, however, the vast majority of courts have neither required any such equalizing measures nor relied on resource disparities to exclude demonstratives from evidence as unfairly prejudicial.

<sup>9</sup>We focus on the Federal Rules of Evidence, which are substantially similar to the evidence rules for most state jurisdictions. A significant number of states have an alternate Frye rule for the admissibility of scientific evidence that differs from the standard in federal court. Accordingly, Frye is also discussed in the expert evidence section.

### B.2 Authentication

Beyond the relevance and prejudice rules, introducing software output as evidence can also implicate the authentication rules. Fundamentally, the authentication rules require a party to introduce a physical object as evidence (rather than oral testimony from a witness) to show that the object is what the party claims it is—in other words, that it has not been forged or tampered with or misidentified. For instance, is this the weapon or bag of white powder that police found at the crime scene? Is this the defendant's signature on the confession or voice in the recording? Does this photograph accurately depict the crime scene?

Applying the authentication rules to software output can be complex, but the key questions are analogous. Is the data that a Celebrite machine extracted from a tablet identical to the data that the tablet contained before it was connected to the Celebrite machine [7]? Is the content on this Facebook page authored by the person whose name and photograph the page displays [120]? Does a tack generated by a Google automated process when a human enters longitude and latitude coordinates into google maps accurately depict which side of the U.S.-Mexican border those coordinates lie [119]?

Importantly, the showing of authenticity that is required to pass the admissibility stage is minimal. The proffering party need only show sufficient evidence for a reasonable jury to conclude that the object is authentic. The rationale behind the low admissibility threshold is that, once the object is admitted into evidence, the opposing party will have an opportunity to challenge its authenticity before the jury. The jury will then get to decide whether they believe that the object is what its proponent claims it is. If the jury decides it is a fake or has been mislabeled, then the jury presumably will not rely on it to reach judgment in the case [94].

Sometimes, a party introducing physical evidence can authenticate it using testimony from an ordinary witness who applies a mode of everyday reasoning that is familiar to the jury. For example, a party seeking to introduce sufficient evidence to support a reasonable jury's finding that a recorded voice is the defendant's voice might call the defendant's co-worker to testify. Assume the co-worker worked closely with the defendant for many years and has developed personal knowledge of what the defendant's voice sounds like. The co-worker can testify that, in their "lay witness" opinion, the recorded voice is the defendant's. There is nothing mysterious or unfamiliar about the co-worker's mode of reasoning. If the jurors believe that the co-worker is telling the truth, then they can draw on their own everyday experience to assess the reliability of the co-worker's opinion.

Other times, however, no such lay witness is available. In that case, the party seeking to introduce the recording might hire a forensic analyst to authenticate it. The analyst might, in turn, deploy a software program to analyze and compare the audio characteristics of the recorded voice and the defendant's voice and generate a statistical likelihood that the two voices match. Based on the software output, the analyst can testify that, in their "expert witness" opinion, the recorded voice is the defendant's. But this time, the analyst's mode of reasoning is highly technical and unfamiliar to jurors drawing on their everyday experiences. It is far more challenging for the jury to properly assess the reliability of the

Category	Products
Breath Testing	Alcotest 9510 [38], Intoxilyzer 8000 [29], LX9 Breathalyzer [77]
Digital Device Extraction Forensics	Cellebrite Universal Forensic Extraction Device (UFED) [22], GrayKey [53], Magnet AXIOM [78]
Environment Audio Detection	SENTRI (Sensor Enabled Neural Threat Recognition and Identification) [104], ShotSpotter [106]
Probabilistic Genotyping	FST (Forensic Statistical Tool) [66], STRMix [61], TrueAllele [35]
Toolmark Analysis	Correlation Engine [117], Matchpoint [118]

**Table 1: Example categories of evidentiary software and a sample of known products for each category.**

testimony. As a result, in these types of circumstances, the legal rules impose an additional gatekeeping responsibility on the judge to make sure that the expert’s analysis meets a minimum level of reliability before admitting the evidence to the jury.

### B.3 Expert Evidence Rules

When experts testify based on technical or scientific modes of reasoning that are unfamiliar to everyday jurors, the party seeking to introduce evidence has to make an additional showing that the expert’s testimony is based on reliable methods that have been reliably applied. There are two leading standards across federal and state jurisdictions: Daubert and Frye [26]. Daubert gives judges flexibility in how they choose to make the reliability determination, but encourages them to consider the following series of factors: “whether the theory or technique in question can be (and has been) tested, whether it has been subjected to peer review and publication, its known or potential error rate and the existence and maintenance of standards controlling its operation and whether it has attracted widespread acceptance within a relevant scientific community [36].” The test in Frye jurisdictions focuses on one factor: whether a method is “sufficiently established to have gained general acceptance in the particular field in which it belongs [44].”

The Daubert and Frye standards have garnered widespread and longstanding debate. A 2009 report by the National Academy of Sciences and a 2016 report by the President’s Council of Advisors on Science and Technology both raised the alarm that judges in criminal cases are admitting substantial amounts of expert testimony based on methods that have not been adequately validated and have no known or knowable error rates. For instance, the PCAST report identified probabilistic genotyping software systems, or software that purports to calculate the likelihood that a defendant’s DNA is present in a complex or degraded mixture of DNA evidence found at a crime scene, as a forensic method that needs more robust validation to determine reliability.

Beyond these critiques from the scientific community, known structural disparities also bias the adversarial process in which judges make their Daubert or Frye admissibility determinations. For example, prosecutors are generally the party seeking to introduce a novel forensic method. In most courts, prosecutors have more resources and access to expertise than do criminal defense counsel. This puts them at an advantage in persuading a judge that the novel method or technique is reliable. Moreover, prosecutors can elect to introduce the method in a case with facts favoring admissibility, encouraging judges to find the method valid in the initial hearing. Subsequent courts often rely on the initial judge’s assessment rather than conducting another foundational reliability

hearing. For software systems specifically, vendors may assert intellectual property rights to withhold methodological information about how the system works, such as source code. Or analysts may assert privacy interests in keeping certain data secret. Those types of assertions raise the burden on defense counsel seeking to access information to conduct an adversarial scrutiny of the methods to identify potential flaws or weaknesses for the judge. If the burden to defeat an intellectual property or privacy claim is too high, defense counsel may be blocked entirely from accessing information necessary to conduct adversarial review.

### B.4 Hearsay and the Confrontation Clause

Finally, other Evidence Law rules that one might expect to control the admissibility of software output do not, under current doctrine, apply. This includes the hearsay doctrine, which restricts the admission of assertions by an out-of-court declarant introduced to prove the truth of the matter asserted. The policy rationale behind this rule is that we want witnesses to come to court and make their assertions from the stand, where they can be placed under oath, observed by the jury, and subjected to cross-examination. Courts have generally held that software systems do not make assertions, and therefore the hearsay rules do not apply to require the production of the developer, designer, or any other person who could be subjected to cross-examination about how the system works. Similarly, courts have generally determined that a criminal defendant’s Sixth Amendment right to confront and cross-examine their accusers does not apply to raw data generated by a machine. To be sure, the Confrontation Clause may require the prosecution to produce the analyst who operates the machine for cross-examination by the defense. But most courts addressing this issue have declined to require the prosecution to produce the software developer—or indeed anyone with a detailed understanding of how the system works—for cross-examination by the defense.