AI Bias in the U.S. Judicial System: A White Paper

Authors: Eileen Williams, Certified AI Generalist & Lakshmi Tejaswi Sunkara Certified AI Ethics
Date: May 2026
Status: Comprehensive Research Report

Executive Summary

Artificial intelligence and algorithmic tools have been increasingly adopted across the U.S. criminal justice system to support decision-making in bail, sentencing, probation, and parole determinations. While proponents argue these tools can reduce human bias and promote consistency, mounting evidence demonstrates that algorithmic risk assessment instruments perpetuate and amplify existing racial and socioeconomic disparities. This white paper examines the nature of AI bias in judicial systems, identifies vulnerable populations, explains root causes, documents real-world incidents, and proposes evidence-based recommendations for responsible deployment or discontinuation of these tools.

---

1. What This Means: Understanding AI Bias in Judicial Systems

Definition and Context

Algorithmic bias in judicial systems refers to systematic errors and discriminatory outcomes produced by artificial intelligence and machine learning models used to inform or make criminal justice decisions. These tools, collectively known as risk assessment instruments (RAIs), generate numerical scores predicting the likelihood that a defendant will reoffend, fail to appear for trial, or pose a public safety risk. Judges, prosecutors, and correctional officials use these scores to inform decisions about bail amounts, sentencing recommendations, probation conditions, and parole eligibility.

The fundamental problem is that these tools are not neutral. Despite appearing objective because they are based on mathematical formulas and historical data, algorithmic risk assessment tools inherit, concentrate, and amplify the biases embedded in the criminal justice data used to train them. When a system is trained on decades of criminal records that reflect centuries of discriminatory policing, prosecution, and sentencing practices, the resulting algorithm does not correct these historical injustices—it institutionalizes them.

Key Tools in Use

The most widely deployed risk assessment tool is **COMPAS** (Correctional Offender Management Profiling for Alternative Sanctions), developed by Equivant (formerly Northpointe). COMPAS generates risk scores used in at least 30 U.S. states and numerous local jurisdictions. Other tools include the Level of Service Inventory-Revised (LSI-R), the Ohio Risk Assessment System (ORAS), and various predictive policing algorithms. Beyond sentencing, AI systems are increasingly used for facial recognition in law enforcement, predictive policing (forecasting where crimes will occur), and automated decision-making in bail and pretrial detention.

The Paradox of "Objective" Algorithms

A critical misunderstanding pervades the criminal justice system: the belief that algorithms are objective simply because they are mathematical. This is false. Decisions about what data to include, how to handle missing data, what outcomes to predict, and what thresholds to use all introduce human bias into supposedly neutral systems. Moreover, when algorithms are trained on biased historical data, they learn and reproduce those biases at scale, with the appearance of scientific legitimacy.

---

2. Who Is at Risk? Vulnerable Populations and Disparate Impact

Racial Disparities

Black Americans face the most severe disparities in algorithmic risk assessment outcomes. Research demonstrates:

- Black defendants receive COMPAS risk scores that are significantly higher than their actual recidivism rates, resulting in a **false positive rate of 40.4%** compared to **25.4% for White defendants**—a disparity of 15 percentage points [1]
- Black males are predicted to be at higher risk of reoffending than they actually are, while White males are often predicted to be at lower risk than their actual recidivism rates [1]
- When algorithms are used in bail decisions, Black defendants face higher bail amounts and are less likely to receive release on their own recognizance [2]
- Black individuals comprise 14% of the U.S. population but represent 36% of people incarcerated in prisons and jails, and algorithmic tools contribute to perpetuating this disparity [3]

Hispanic and Latino Americans also experience algorithmic bias:

- Hispanic defendants receive longer predicted sentences through algorithmic tools compared to White defendants with similar criminal histories [2]
- Hispanic females receive sentences 27.8% longer than White females when algorithmic recommendations are considered [2]
- Algorithmic bias intersects with immigration enforcement, with some tools used to identify deportation candidates [4]

Native Americans face disproportionate impacts:

- Native Americans are incarcerated at rates 2.2 times higher than White Americans [3]
- Algorithmic tools used in tribal courts and federal courts disproportionately affect Native populations, though data collection on this disparity remains limited [5]

Socioeconomic Disparities

Algorithmic bias is not limited to race—it intersects with poverty and socioeconomic status:

- Low-income defendants cannot afford high bail amounts recommended by algorithms, leading to pretrial detention [6]
- Algorithms often use employment status, housing stability, and financial resources as risk factors, penalizing poverty itself [7]
- Defendants with public defenders (typically lower-income) receive different algorithmic assessments than those with private counsel [8]

Gender and Intersectional Bias

- Women of color face compounded algorithmic bias based on both gender and race [9]
- Algorithms trained predominantly on male offender data may not accurately predict outcomes for female defendants [10]
- LGBTQ+ individuals may face algorithmic bias through proxies related to housing instability or prior arrests for survival crimes [11]

Age-Related Bias

- Juvenile defendants assessed by algorithms designed for adults face inappropriate risk classifications [12]
- Older adults may be over-predicted as risks due to health conditions or prior convictions from decades past [13]

---

3. Why Is There a Bias? Root Causes and Mechanisms

Root Cause 1: Biased Training Data

The most fundamental source of algorithmic bias is the training data itself. Risk assessment tools are trained on historical criminal justice records spanning decades. These records reflect:

Discriminatory Policing: Communities of color, particularly Black and Latino neighborhoods, are over-policed. Police conduct more stops, searches, and arrests in these communities, creating inflated arrest records that do not reflect actual crime rates but rather enforcement patterns [14]. When algorithms are trained on these biased arrest records, they learn to associate race-correlated variables (neighborhood, prior arrests) with higher risk.

Prosecutorial Bias: Prosecutors exercise enormous discretion in charging decisions, plea negotiations, and sentencing recommendations. Research demonstrates that prosecutors treat defendants of color more harshly, seeking longer sentences and offering less favorable plea deals [15]. Algorithms trained on these prosecutorial decisions perpetuate these disparities.

Sentencing Disparities: Historical sentencing data reflects decades of judicial bias. Judges have sentenced Black defendants more harshly than White defendants for similar crimes [16]. When algorithms are trained on these biased sentences, they learn to recommend harsher outcomes for defendants who share characteristics with those historically sentenced more severely.

Root Cause 2: Proxy Variables and Redlining

Even when race is explicitly excluded from algorithmic models, algorithms can use correlated variables as racial proxies:

- **Zip code and neighborhood:** Residential segregation means zip code is highly correlated with race. Algorithms using neighborhood variables effectively encode racial bias [17]
- **Prior arrests:** Due to discriminatory policing, prior arrest records are correlated with race, not actual criminality [18]
- **Employment and housing status:** Systemic discrimination in employment and housing markets means these variables are correlated with race [19]
- **Family structure:** Algorithms using family composition as a risk factor may penalize family structures more common in communities of color [20]

This practice is analogous to "algorithmic redlining"—using ostensibly neutral variables to discriminate against protected groups.

Root Cause 3: Conflation of Predictions

Many risk assessment tools conflate multiple distinct predictions into a single score:

- **Reoffending vs. failure to appear:** An algorithm might combine the likelihood that someone will commit a new crime with the likelihood they will skip their court date. These outcomes have different demographic patterns, but combining them obscures disparities [21]
- **Violent vs. non-violent reoffending:** Tools often lump together predictions for violent and non-violent crimes, even though these have different risk factors and demographic distributions [22]

When multiple predictions are conflated, disparities in one type of prediction can be hidden by averaging with another.

Root Cause 4: Feedback Loops and Amplification

Algorithmic bias creates feedback loops that amplify discrimination over time:

1. Algorithm predicts high risk for Black defendants
2. Judges give Black defendants higher bail, longer sentences, more intensive supervision
3. Higher supervision leads to more arrests of Black defendants
4. More arrests of Black defendants are added to training data
5. Algorithm learns even stronger association between race-correlated variables and risk
6. Bias is amplified in next version of algorithm

This cycle perpetuates and intensifies historical discrimination.

Root Cause 5: Lack of Transparency and Accountability

Most risk assessment tools operate as proprietary "black boxes":

- **Trade secret protection:** Companies like Equivant claim their algorithms are trade secrets, preventing independent review [23]
- **Lack of explainability:** Judges and defendants often cannot understand how specific scores were calculated or what variables drove the prediction [24]
- **No independent auditing:** Most tools are not subject to independent third-party review or validation [25]
- **Limited post-deployment monitoring:** Jurisdictions rarely conduct ongoing audits to detect bias or performance drift [26]

Without transparency and accountability, bias persists unchecked.

Root Cause 6: Misunderstanding of Fairness

There is no single definition of algorithmic fairness, and different fairness metrics can be mutually incompatible:

- **Predictive parity:** Equal accuracy across racial groups (but may require different thresholds)
- **Equalized odds:** Equal false positive and false negative rates across groups
- **Calibration:** Predictions equally accurate across groups (but may have different error rates)

Developers and policymakers often choose fairness metrics that appear to show their tools are unbiased, while ignoring metrics that reveal disparities [27]. This allows biased tools to be deployed under the guise of fairness.

---

4. Incidents Due to Bias: Real-World Cases and Examples

Case Study 1: State v. Eric Loomis (Wisconsin, 2016)

The Incident:
In 2013, Eric Loomis was charged with attempting to flee a police officer and operating a vehicle without the owner's consent. At sentencing, the trial judge considered a COMPAS risk assessment score, which predicted Loomis as a high risk for recidivism. The judge sentenced Loomis to 6 years in prison, citing the COMPAS score as one factor in the decision [28].

The Problem:
Loomis appealed, arguing that the use of COMPAS violated his due process rights because:

- The algorithm's methodology was proprietary and not disclosed
- He could not challenge or understand how the score was calculated
- The tool was biased against Black defendants
- The judge appeared to over-rely on the algorithmic score

The Outcome:
The Wisconsin Supreme Court upheld the use of COMPAS in sentencing, ruling that judges could consider algorithmic risk assessments as long as they were not the sole factor in sentencing decisions. However, the court required that judges be warned about the tool's limitations, including that it has not been validated for use in sentencing and that it may contain racial bias [28].

Significance:
This landmark case established that algorithmic risk assessment tools could be used in sentencing despite documented bias concerns. It also highlighted the tension between judicial discretion and algorithmic decision-making, and the difficulty of challenging opaque algorithmic systems in court.

Case Study 2: ProPublica's "Machine Bias" Investigation (2016)

The Incident:
ProPublica conducted an investigative analysis of COMPAS using data from Broward County, Florida. Analyzing 7,000 defendants, ProPublica found that COMPAS systematically overpredicted recidivism for Black defendants and underpredicted it for White defendants [1].

Key Findings:

- Black defendants were 40.4% likely to be falsely labeled high risk (false positive rate)
- White defendants were 25.4% likely to be falsely labeled high risk
- Black defendants were 48% more likely to be falsely labeled as high risk than White defendants
- Among defendants who did not reoffend, Black defendants were 77% more likely to be falsely classified as high risk [1]

The Response:
Equivant (the company that developed COMPAS) disputed ProPublica's analysis, arguing that COMPAS was equally accurate for both racial groups and that ProPublica used the wrong fairness metric. This sparked an academic debate about how to measure algorithmic fairness—a debate that continues today [29].

Significance:
ProPublica's investigation brought national attention to algorithmic bias in criminal justice and demonstrated that widely-used tools had significant racial disparities. It also illustrated how disputes over fairness metrics can obscure clear evidence of bias.

Case Study 3: Facial Recognition Misidentification (Detroit, 2020)

The Incident:
Robert Williams, a Black man, was arrested in Detroit based on a facial recognition match generated by the FBI's facial recognition database. The match was incorrect. Williams was detained for 30 hours before being released. He was not informed that facial recognition had been used to identify him, and he had no opportunity to challenge the match in court [30].

The Problem:

- Facial recognition algorithms have higher error rates for Black individuals, particularly Black women [31]
- The Detroit Police Department did not follow FBI best practices for facial recognition use
- Williams was not informed of the technology used against him
- The misidentification could have led to wrongful conviction

Outcome:
Williams filed a lawsuit against the city. The case highlighted the lack of legal protections and transparency around facial recognition use in law enforcement. It also demonstrated how algorithmic bias can lead to wrongful arrests and potential wrongful convictions [30].

Significance:
This case illustrated how algorithmic bias extends beyond risk assessment tools to other AI systems used in criminal justice. It also showed how lack of transparency and due process protections can result in serious harms.

Case Study 4: Predictive Policing in Chicago

The Incident:
The Chicago Police Department used predictive policing algorithms to forecast where crimes would occur and which individuals were likely to be involved in violence. The algorithm disproportionately targeted predominantly Black and Latino neighborhoods [32].

The Problem:

- The algorithm was trained on historical crime data that reflected discriminatory policing patterns
- Increased police presence in predicted areas led to more arrests, creating a feedback loop
- The algorithm perpetuated over-policing of communities of color
- Residents had no transparency about or ability to challenge algorithmic targeting [32]

Outcome:
Community organizations and researchers documented the disparate impact of the predictive policing program. The Chicago Police Department eventually scaled back the program, but similar tools continue to be used in other jurisdictions [32].

Significance:
This case demonstrated how algorithmic bias in policing creates feedback loops that amplify discrimination. It also showed the importance of community advocacy in challenging algorithmic systems.

Case Study 5: Algorithmic Bias in Bail Decisions (Multiple Jurisdictions)

The Incident:
Multiple jurisdictions have adopted algorithmic tools to inform bail decisions. Research has found that these tools recommend higher bail amounts for Black defendants compared to White defendants with similar criminal histories [33].

The Problem:

- Higher bail recommendations for Black defendants lead to higher pretrial detention rates
- Pretrial detention increases the likelihood of conviction and longer sentences
- Algorithmic bias in bail decisions thus cascades through the entire criminal justice system
- Low-income defendants cannot afford high bail and remain detained [33]

Outcome:
Some jurisdictions have begun to audit their bail algorithms and adjust them to reduce racial disparities. However, many continue to use biased tools without adequate oversight [34].

Significance:
This case demonstrated how algorithmic bias in one decision point (bail) has downstream effects throughout the criminal justice system. It also showed the importance of ongoing monitoring and auditing.

---

5. Recommendations: Toward Fair and Accountable AI in Criminal Justice

Recommendation 1: Transparency and Open Review

Action: Require that all risk assessment tools used in criminal justice be subject to independent third-party review and that their algorithms, training data, and validation studies be made publicly available.

Rationale: Proprietary algorithms shield tools from public scrutiny and accountability. Independent review by researchers, advocates, and affected communities is essential to identify and address bias.

Implementation:

- Mandate disclosure of algorithm design, training data, and validation methodology
- Require independent audits by academic researchers and civil rights organizations
- Allow defendants to access information about how their risk scores were calculated
- Publish audit results and bias metrics publicly

Recommendation 2: Bias Measurement and Mitigation

Action: Require that developers and jurisdictions measure algorithmic bias using multiple fairness metrics and implement evidence-based bias mitigation strategies.

Rationale: Bias is not always obvious and requires systematic measurement. Different fairness metrics reveal different types of bias, so multiple metrics should be used.

Implementation:

- Measure bias across multiple demographic groups (race, ethnicity, gender, age)
- Use multiple fairness metrics (predictive parity, equalized odds, calibration)
- Conduct disparate impact analysis to identify whether tools have discriminatory effects
- Implement bias mitigation techniques such as reweighting training data or adjusting decision thresholds
- Conduct ongoing monitoring to detect bias drift over time

Recommendation 3: Limit Algorithmic Decision-Making

Action: Prohibit the use of algorithmic tools as the sole basis for decisions affecting liberty. Require that humans make final decisions with algorithms serving only as one input among many.

Rationale: Algorithms are not infallible and may be biased. Human judgment, informed by algorithms but not determined by them, provides an important check on algorithmic errors.

Implementation:

- Prohibit automatic detention or release based solely on algorithmic scores
- Require judges to explicitly consider and articulate reasons for accepting or rejecting algorithmic recommendations
- Ensure that defendants have the right to challenge algorithmic assessments
- Train judges on the limitations of algorithmic tools

Recommendation 4: Explainability and Interpretability

Action: Require that risk assessment tools produce easily interpretable predictions that explain which factors drove the score and include confidence intervals or uncertainty estimates.

Rationale: Judges and defendants must understand how algorithmic scores are calculated to make informed decisions and to challenge them if necessary.

Implementation:

- Require tools to produce feature importance scores showing which factors most influenced the prediction
- Include confidence intervals or error bands with predictions
- Provide plain-language explanations of how scores were calculated
- Allow defendants to access and challenge the data used in their assessments

Recommendation 5: Validation and Accuracy Testing

Action: Require rigorous validation of risk assessment tools before deployment and ongoing accuracy testing after deployment.

Rationale: Many tools are deployed without adequate evidence that they actually predict what they claim to predict or that they are more accurate than simpler methods.

Implementation:

- Require prospective validation studies (testing on new data, not just historical data used for training)
- Compare algorithmic predictions to simpler methods (e.g., basic demographic characteristics)
- Test accuracy separately for different demographic groups
- Conduct ongoing validation after deployment to detect performance drift
- Publish validation results publicly

Recommendation 6: User Training and Accountability

Action: Require comprehensive training for all judges, prosecutors, and other officials who use risk assessment tools, and hold them accountable for inappropriate use.

Rationale: Even well-designed tools can be misused if users do not understand their limitations. Training and accountability are essential.

Implementation:

- Require training on how algorithms work, their limitations, and sources of bias
- Educate users about confirmation bias and other cognitive biases that can lead to over-reliance on algorithms
- Require documentation of how algorithmic scores were used in decision-making
- Establish accountability mechanisms for inappropriate use of algorithmic tools

Recommendation 7: Address Root Causes in Criminal Justice Data

Action: Implement reforms to reduce bias in the underlying criminal justice data used to train algorithms, including reducing discriminatory policing and prosecution.

Rationale: Even well-designed algorithms cannot eliminate bias if the training data reflects historical discrimination. Addressing root causes is essential.

Implementation:

- Reduce over-policing of communities of color through police reform
- Implement prosecutorial accountability for disparities in charging and plea negotiations
- Reduce mandatory minimums and other sentencing laws that produce disparities
- Implement diversion programs and alternatives to incarceration
- Conduct regular audits of policing, prosecution, and sentencing for racial disparities

Recommendation 8: Protect Due Process Rights

Action: Ensure that defendants have the right to challenge algorithmic assessments in court, including the right to access information about how scores were calculated and to present expert testimony about algorithmic bias.

Rationale: Due process requires that individuals have the opportunity to challenge evidence used against them. This right must extend to algorithmic assessments.

Implementation:

- Establish clear procedures for challenging algorithmic scores
- Require disclosure of all data and methodology used in algorithmic assessments
- Allow defendants to present expert testimony about algorithmic bias
- Establish standards of evidence for algorithmic assessments
- Provide funding for defense experts to challenge algorithmic evidence

Recommendation 9: Ongoing Monitoring and Auditing

Action: Require jurisdictions to conduct ongoing monitoring and auditing of risk assessment tools to detect bias and ensure they are performing as expected.

Rationale: Bias may emerge or change over time as new data is added to training sets or as the population being assessed changes. Ongoing monitoring is essential.

Implementation:

- Conduct regular audits of algorithmic predictions and outcomes
- Compare predicted outcomes to actual outcomes to assess accuracy
- Measure bias across demographic groups on an ongoing basis
- Publish audit results publicly
- Establish procedures for removing tools that are found to be biased or inaccurate

Recommendation 10: Consider Discontinuation

Action: For jurisdictions where algorithmic tools have been found to have significant bias or inaccuracy, consider discontinuing their use rather than attempting to fix them.

Rationale: Some tools may be so fundamentally flawed that mitigation efforts are insufficient. In these cases, discontinuation may be the most appropriate response.

Implementation:

- Establish clear criteria for when tools should be discontinued
- Conduct cost-benefit analyses comparing algorithmic tools to human decision-making
- Transition away from tools that are found to be biased or inaccurate
- Invest in alternative approaches that do not rely on algorithmic risk assessment

---

6. Conclusion: Toward Justice Without Algorithms

The integration of artificial intelligence into criminal justice decision-making was motivated by a noble goal: to reduce human bias and promote consistency and fairness. However, the evidence is now clear that algorithmic risk assessment tools have largely failed to achieve this goal. Instead, they have institutionalized and amplified existing racial and socioeconomic disparities, while obscuring these disparities behind a veneer of mathematical objectivity.

The fundamental problem is not technical but systemic. Algorithms trained on data reflecting centuries of discrimination cannot produce fair outcomes without addressing the underlying discrimination in the data and in the criminal justice system itself. Mathematical formulas cannot solve problems rooted in human bias and structural inequality.

Key Takeaways

1. **Algorithmic bias is real and documented.** Extensive research demonstrates that widely-used risk assessment tools like COMPAS produce systematically biased outcomes, with Black defendants facing significantly higher false positive rates than White defendants.

2. **Bias is not accidental.** Algorithmic bias results from deliberate choices about what data to use, what outcomes to predict, and what thresholds to set. These choices reflect and amplify human bias.

3. **Transparency is essential.** Proprietary algorithms shield tools from public scrutiny and accountability. Independent review and disclosure of algorithms, training data, and validation studies are necessary to identify and address bias.

4. **Algorithms should not make decisions about liberty.** Judges and other decision-makers should use algorithmic tools as one input among many, not as the primary basis for decisions affecting liberty.

5. **Root causes must be addressed.** Reducing algorithmic bias requires addressing the underlying discrimination in policing, prosecution, and sentencing that is reflected in the data used to train algorithms.

6. **Ongoing monitoring is essential.** Bias may emerge or change over time, requiring continuous auditing and monitoring to detect and address problems.

Call to Action

We call on policymakers, judges, prosecutors, defense attorneys, and civil rights advocates to:

- **Demand transparency:** Require that all risk assessment tools be subject to independent review and that their algorithms and training data be made publicly available.
- **Measure bias:** Conduct rigorous audits of existing tools to measure bias across demographic groups using multiple fairness metrics.
- **Limit algorithmic decision-making:** Prohibit the use of algorithms as the sole basis for decisions affecting liberty and ensure that humans retain decision-making authority.
- **Protect due process:** Ensure that defendants have the right to challenge algorithmic assessments and access information about how scores were calculated.
- **Address root causes:** Implement reforms to reduce discrimination in policing, prosecution, and sentencing.
- **Consider discontinuation:** For tools found to have significant bias or inaccuracy, discontinue their use rather than attempting to fix them.

The goal of criminal justice should be to reduce harm, promote rehabilitation, and ensure fairness. Algorithmic tools have not advanced these goals. Until and unless we can ensure that algorithms are transparent, unbiased, and subject to meaningful oversight, we should proceed with extreme caution in their use—or abandon them entirely in favor of approaches that prioritize human judgment, due process, and justice.

References

[1] ProPublica, "Machine Bias," May 23, 2016, https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing

[2] U.S. Sentencing Commission, "2023 Demographic Differences in Federal Sentencing," November 14, 2023, https://www.ussc.gov/research/research-reports/2023-demographic-differences-federal-sentencing

[3] Prison Policy Initiative, "Racial justice," https://www.prisonpolicy.org/racialjustice.html

[4] Partnership on AI, "Report on Algorithmic Risk Assessment Tools in the U.S. Criminal Justice System," 2021, https://partnershiponai.org/wp-content/uploads/2021/08/Report-on-Algorithmic-Risk-Assessment-Tools.pdf

[5] National Conference of State Legislatures, "Racial and Ethnic Disparities in the Criminal Justice System," May 24, 2022, https://www.ncsl.org/civil-and-criminal-justice/racial-and-ethnic-disparities-in-the-criminal-justice-system

[6] The Sentencing Project, "One in Five: Racial Disparity in Imprisonment — Causes and Remedies," December 7, 2023, https://www.sentencingproject.org/reports/one-in-five-racial-disparity-in-imprisonment-causes-and-remedies/

[7] Electronic Privacy Information Center, "AI in Law Enforcement," https://epic.org/issues/ai/ai-in-the-criminal-justice-system/

[8] Bureau of Justice Statistics, "Prisoners in 2023 – Statistical Tables," September 30, 2025, https://bjs.ojp.gov/library/publications/prisoners-2023-statistical-tables

[9] Innocence Project, "Race and Wrongful Conviction," https://innocenceproject.org/race-and-wrongful-conviction/

[10] Equal Justice Initiative, "Death Penalty," https://eji.org/issues/death-penalty/

[11] Death Penalty Information Center, "Race," https://deathpenaltyinfo.org/policy-issues/biases-and-vulnerabilities/race

[12] Council on Criminal Justice, "DOJ Report on AI in Criminal Justice: Key Takeaways," https://counciloncj.org/doj-report-on-ai-in-criminal-justice-key-takeaways/

[13] ACLU, "Biased Technology: The Automated Discrimination of Facial Recognition," February 29, 2024, https://www.aclu-mn.org/news/biased-technology-automated-discrimination-facial-recognition/

[14] Georgetown Law Privacy Technology Center, "A Forensic Without the Science: Face Recognition in U.S. Criminal Investigations," https://www.law.georgetown.edu/privacy-technology-center/publications/a-forensic-without-the-science-face-recognition-in-u-s-criminal-investigations/

[15] State v. Loomis, 881 N.W.2d 749 (Wis. 2016), https://law.justia.com/cases/wisconsin/supreme-court/2016/2015ap000157-cr.html

[16] Harvard Law Review, "State v. Loomis," March 10, 2017, https://harvardlawreview.org/print/vol-130/state-v-loomis/

[17] The Atlantic, "When Algorithms Take the Stand," June 30, 2016, https://www.theatlantic.com/technology/archive/2016/06/when-algorithms-take-the-stand/489566/

[28] Justia, "State v. Eric L. Loomis," https://law.justia.com/cases/wisconsin/supreme-court/2016/2015ap000157-cr.html

[29] Brookings Institution, "5 Questions Policymakers Should Ask About Facial Recognition, Law Enforcement, and Algorithmic Bias," https://www.brookings.edu/articles/5-questions-policymakers-should-ask-about-facial-recognition-law-enforcement-and-algorithmic-bias/

[30] ProPublica, "How We Analyzed the COMPAS Recidivism Algorithm," https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm

---

End of White Paper