

# INTERNATIONAL JOURNAL FOR LEGAL RESEARCH AND ANALYSIS



Open Access, Refereed Journal Multi Disciplinary  
Peer Reviewed

[www.ijlra.com](http://www.ijlra.com)

## DISCLAIMER

No part of this publication may be reproduced, stored, transmitted, or distributed in any form or by any means, whether electronic, mechanical, photocopying, recording, or otherwise, without prior written permission of the Managing Editor of the *International Journal for Legal Research & Analysis (IJLRA)*.

The views, opinions, interpretations, and conclusions expressed in the articles published in this journal are solely those of the respective authors. They do not necessarily reflect the views of the Editorial Board, Editors, Reviewers, Advisors, or the Publisher of IJLRA.

Although every reasonable effort has been made to ensure the accuracy, authenticity, and proper citation of the content published in this journal, neither the Editorial Board nor IJLRA shall be held liable or responsible, in any manner whatsoever, for any loss, damage, or consequence arising from the use, reliance upon, or interpretation of the information contained in this publication.

The content published herein is intended solely for academic and informational purposes and shall not be construed as legal advice or professional opinion.

**Copyright © International Journal for Legal Research & Analysis.  
All rights reserved.**

## ABOUT US

The *International Journal for Legal Research & Analysis (IJLRA)* (ISSN: 2582-6433) is a peer-reviewed, academic, online journal published on a monthly basis. The journal aims to provide a comprehensive and interactive platform for the publication of original and high-quality legal research.

IJLRA publishes Short Articles, Long Articles, Research Papers, Case Comments, Book Reviews, Essays, and interdisciplinary studies in the field of law and allied disciplines. The journal seeks to promote critical analysis and informed discourse on contemporary legal, social, and policy issues.

The primary objective of IJLRA is to enhance academic engagement and scholarly dialogue among law students, researchers, academicians, legal professionals, and members of the Bar and Bench. The journal endeavours to establish itself as a credible and widely cited academic publication through the publication of original, well-researched, and analytically sound contributions.

IJLRA welcomes submissions from all branches of law, provided the work is original, unpublished, and submitted in accordance with the prescribed submission guidelines. All manuscripts are subject to a rigorous peer-review process to ensure academic quality, originality, and relevance.

Through its publications, the *International Journal for Legal Research & Analysis* aspires to contribute meaningfully to legal scholarship and the development of law as an instrument of justice and social progress.

## ***PUBLICATION ETHICS, COPYRIGHT & AUTHOR RESPONSIBILITY STATEMENT***

The *International Journal for Legal Research and Analysis (IJLRA)* is committed to upholding the highest standards of publication ethics and academic integrity. All manuscripts submitted to the journal must be original, unpublished, and free from plagiarism, data fabrication, falsification, or any form of unethical research or publication practice. Authors are solely responsible for the accuracy, originality, legality, and ethical compliance of their work and must ensure that all sources are properly cited and that necessary permissions for any third-party copyrighted material have been duly obtained prior to submission. Copyright in all published articles vests with IJLRA, unless otherwise expressly stated, and authors grant the journal the irrevocable right to publish, reproduce, distribute, and archive their work in print and electronic formats. The views and opinions expressed in the articles are those of the authors alone and do not reflect the views of the Editors, Editorial Board, Reviewers, or Publisher. IJLRA shall not be liable for any loss, damage, claim, or legal consequence arising from the use, reliance upon, or interpretation of the content published. By submitting a manuscript, the author(s) agree to fully indemnify and hold harmless the journal, its Editor-in-Chief, Editors, Editorial Board, Reviewers, Advisors, Publisher, and Management against any claims, liabilities, or legal proceedings arising out of plagiarism, copyright infringement, defamation, breach of confidentiality, or violation of third-party rights. The journal reserves the absolute right to reject, withdraw, retract, or remove any manuscript or published article in case of ethical or legal violations, without incurring any liability.

# **ARTIFICIAL INTELLIGENCE IN COURTROOMS: LEGAL AND ETHICAL CHALLENGES IN JUDICIAL DECISION-MAKING**

AUTHORED BY - ZAIN RAFIQ & DR. PRAMOD KUMAR

## **Abstract**

*The integration of artificial intelligence (AI) into judicial systems—ranging from risk-assessment instruments used in bail and sentencing to natural-language-processing tools that assist legal research—has accelerated sharply over the past decade. Proponents argue that AI can reduce case backlogs, standardize outcomes, and curb certain forms of human bias; critics counter that opaque, proprietary algorithms risk entrenching discrimination, undermining due process, and diffusing accountability away from identifiable human decision-makers. This paper undertakes a structured review of the legal and ethical challenges posed by AI in courtrooms. It examines algorithmic risk-assessment tools such as COMPAS and the landmark case *State v. Loomis* (2016); analyzes due-process concerns surrounding transparency, explainability, and the right to a human decision-maker; surveys comparative regulatory responses in the United States, the European Union, and India; and synthesizes recommendations for responsible adoption. The analysis finds that while AI can offer efficiency gains, its courtroom use without robust human oversight, contestability mechanisms, and bias auditing poses material risks to fundamental fair-trial guarantees. The paper concludes that a calibrated, rights-centered governance framework—rather than outright prohibition or unregulated adoption—offers the most defensible path forward.*

**Keywords:** *artificial intelligence, judicial decision-making, algorithmic bias, due process, COMPAS, EU AI Act, legal technology, explainable AI*

## 1. Introduction

Courts have always relied on tools to manage information—from paper dockets to electronic case-management systems. What distinguishes the current moment is the shift from AI as a passive record-keeper to AI as an active participant in adjudicative reasoning: scoring a defendant's likelihood of reoffending, flagging relevant precedent, drafting portions of judicial opinions, or, in a small number of pilot jurisdictions, proposing dispositions for low-value disputes. This shift raises the stakes considerably, because judicial decisions implicate liberty, property, and fundamental rights in ways that most administrative AI applications do not.

This paper examines the legal and ethical terrain created by this shift. It is organized around four questions. First, where is AI currently being used in courts, and what forms does that use take? Second, what specific due-process and constitutional concerns arise when algorithmic outputs inform judicial decisions? Third, how have different legal systems—common law and civil law, common-law and statute-driven—responded through regulation and case law? Fourth, what design and governance principles would allow courts to capture efficiency benefits while safeguarding fair-trial rights?

The paper draws on publicly documented case studies (notably the COMPAS recidivism-scoring controversy), primary legal sources (including the Wisconsin Supreme Court's decision in *State v. Loomis* and the European Union's Artificial Intelligence Act), and the broader academic literature on algorithmic fairness and accountability. Because empirical, court-verified data on AI's real-world impact remains limited and unevenly reported, the analysis is necessarily qualitative in places and flags where claims are contested or unverified.<sup>1</sup>

## 2. Current Applications of AI in Judicial Systems

AI tools currently used in or around courtrooms can be grouped into four broad categories, summarized in Table 1.

### 2.1 Risk-Assessment and Sentencing-Support Tools

Actuarial risk-assessment instruments such as COMPAS (Correctional Offender Management Profiling for Alternative Sanctions), the Public Safety Assessment (PSA), and various state-level recidivism tools generate numerical scores intended to predict the likelihood that a defendant will reoffend or fail to appear. These scores are used at multiple stages—pretrial bail determinations, sentencing, and parole—and are typically presented to judges as one input

---

<sup>1</sup> COMPAS 2021 Policy

among several.

### 2.2 Legal Research and Drafting Assistance

Large-language-model-based tools increasingly assist judges' clerks and court staff with case-law research, document summarization, and first-draft generation of routine orders. These uses are generally framed as decision-support rather than decision-making, but the line between the two can blur in practice, particularly when time-pressured staff defer heavily to AI-generated summaries.

### 2.3 Online Dispute Resolution and Small-Claims Automation

Several jurisdictions have piloted automated or semi-automated resolution of low-value civil disputes (such as parking fines or small consumer claims), where parties submit documentation through a digital platform and receive an algorithmically generated proposed outcome, typically subject to appeal to a human judge. Estonia's Ministry of Justice was widely reported in 2019 media coverage to be developing such a "robot judge" pilot; notably, the Ministry officially stated in 2022 that no such project existed in the form reported, illustrating how readily media narratives about "AI judges" can outpace verified deployment. This gap between hype and documented reality is itself a recurring theme in the literature and a caution for researchers citing such examples.

### 2.4 Prosecutorial and Administrative Decision-Support

Beyond the bench, prosecutorial offices have adopted AI tools for case prioritization and evidence analysis. The Prometea system, deployed by the Public Prosecutor's Office of the City of Buenos Aires beginning in 2017, is among the most frequently cited examples of AI assisting in drafting routine prosecutorial decisions and case triage, illustrating administrative rather than adjudicative use.

**Table 1. Categories of AI Applications in Judicial Systems**

Application Type	Example Tool / Jurisdiction	Decision-Making Role
Recidivism risk scoring	COMPAS (USA, multiple states)	Informs bail/sentencing; human judge retains final say
Legal research & drafting	AI research assistants (various courts)	Decision-support for clerks and staff

Small-claims automation	Online dispute resolution pilots (e.g., UK Money Claim Online; reported Estonia pilot)	Proposed outcome, appealable to a human judge
Prosecutorial triage	Prometea (Buenos Aires)	Drafts routine decisions; supports case prioritization

### 3. Legal Challenges: Due Process and Constitutional Dimensions

The introduction of algorithmic scoring into decisions that affect liberty implicates several interlocking due-process guarantees: the right to be heard, the right to confront and contest adverse evidence, the right to an individualized determination, and the right to a reasoned, reviewable decision. The following subsections examine how these guarantees come under strain.

#### 3.1 The State v. Loomis Precedent

The clearest judicial treatment of these issues remains the Wisconsin Supreme Court's 2016 decision in *State v. Loomis*. Eric Loomis challenged his sentence on the ground that the court's consideration of his COMPAS risk score violated due process because the proprietary algorithm's methodology was undisclosed, preventing him from challenging its scientific validity, and because the tool used group-based statistical data to make an individualized determination about him. The Wisconsin Supreme Court upheld the use of COMPAS, reasoning that the score was one of several factors considered, that the sentencing court was not required to rely on it, and that the defendant had access to the questionnaire responses underlying his own score even without access to the underlying algorithm. The court did, however, require that any presentence report using COMPAS include a written advisement of the tool's limitations. The U.S. Supreme Court denied certiorari, leaving *Loomis* as the leading—and largely unchallenged—precedent on algorithmic risk assessment in American sentencing.

Legal scholars have criticized *Loomis* for permitting continued reliance on a tool whose internal weighting remains a trade secret, arguing that the "one factor among many" framing understates how heavily anchoring effects can influence even well-intentioned judges.

#### 3.2 The Black-Box Problem and the Right to Contest Evidence

Most commercial risk-assessment and legal-AI tools are proprietary, with vendors asserting

trade-secret protection over their scoring methodologies. This creates tension with the principle, rooted in due process, that a party should be able to understand and contest the basis of an adverse decision. Even where the underlying questionnaire is disclosed, the relative weighting of factors, the training data, and the statistical model itself typically remain undisclosed, making meaningful contestation difficult.

### **3.3 Individualized Justice versus Group-Based Prediction**

A second, more conceptual challenge concerns the tension between actuarial, group-based prediction and the common-law commitment to individualized sentencing. Risk scores are generated by comparing a defendant's characteristics to those of a historical reference population; the resulting score describes a statistical tendency of people "like" the defendant rather than a finding of fact about the defendant. Critics argue this sits uneasily with constitutional and common-law traditions that treat each defendant's circumstances as requiring particularized assessment.

### **3.4 Automation Bias and the Erosion of Judicial Independence**

Empirical research on human-automation interaction, primarily from aviation and clinical-decision-support contexts, documents a robust tendency toward automation bias: a propensity for human decision-makers to over-trust algorithmic recommendations, particularly under time pressure or heavy caseloads. Applied to courts, this raises the concern that risk scores presented as neutral, data-driven outputs may carry more practical weight than their formal status as "one factor among many" suggests, subtly displacing the independent judgment that due process presumes.

## **4. Ethical Challenges: Bias, Fairness, and Accountability**

### **4.1 The COMPAS Controversy: A Case Study in Algorithmic Bias**

In 2016, ProPublica journalists Julia Angwin and colleagues published "Machine Bias," an investigative analysis of approximately 7,000 COMPAS risk scores assigned to defendants in Broward County, Florida, cross-referenced against actual two-year reoffending outcomes. The investigation found a clear racial disparity in the algorithm's error patterns, illustrated in Figure 1.

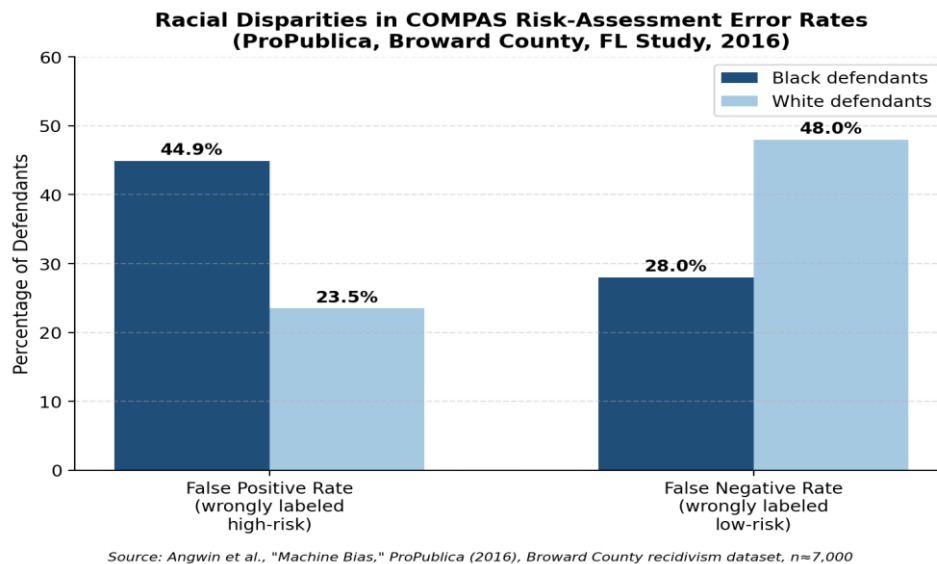


Figure 1. Racial disparities in COMPAS false-positive and false-negative rates (ProPublica, 2016).

As Figure 1 shows, Black defendants who did not subsequently reoffend were nearly twice as likely as similarly situated white defendants to be wrongly flagged as high-risk (a false positive). Conversely, white defendants who did go on to reoffend were considerably more likely than Black defendants to have been wrongly scored as low-risk (a false negative). Northpointe (the algorithm's developer, later renamed Equivant) disputed ProPublica's framing, noting that COMPAS was approximately equally calibrated across racial groups—meaning that defendants who received the same score had similar actual reoffending rates regardless of race. This dispute exposed a now well-known mathematical reality in the fairness literature: when base rates of an outcome differ across groups (as historical arrest and reoffending data do, for reasons that are themselves entangled with policing and socioeconomic patterns), it is mathematically impossible for an algorithm to simultaneously satisfy both equal calibration and equal error rates across groups. The COMPAS debate is therefore not merely a dispute about one vendor's software but an illustration of a structural trade-off inherent to any predictive tool trained on historically skewed data.

#### 4.2 Sources of Bias in Judicial AI

- Historical training data: models trained on past arrest, charging, or sentencing records can encode the biases embedded in decades of policing and prosecutorial practice.
- Proxy variables: factors such as zip code, employment history, or family criminal history can function as statistical proxies for race or socioeconomic status even when race is formally excluded from the model.

- Feedback loops: when algorithmic outputs influence which populations are subject to closer scrutiny, the resulting data can reinforce the original bias in subsequent model iterations.
- Validation gaps: tools validated on one jurisdiction's population are frequently deployed in other jurisdictions without adequate revalidation.

### 4.3 Accountability and the Diffusion of Responsibility

When an AI-informed decision causes harm—for example, an inflated risk score contributing to an unwarranted detention—existing accountability structures struggle to assign responsibility cleanly. The vendor can point to the judge's discretion to disregard the score; the judge can point to the tool's presumed validation and the absence of disclosed methodology; and the court system can point to legislative or executive mandates requiring the tool's use. This diffusion of responsibility is itself an ethical harm, distinct from any single biased outcome, because it weakens the deterrent and corrective functions that accountability is meant to serve.

### 4.4 The Dignitary Dimension

Beyond statistical fairness, scholars in law and philosophy have raised a dignitary objection: that being judged, even in part, by a statistical comparison to other people's behavior, rather than solely on the facts of one's own conduct, may itself offend the principle that legal subjects are entitled to be treated as individuals rather than instances of a category. This objection persists even for a hypothetically perfectly calibrated, bias-free algorithm, and distinguishes ethical critique from purely technical bias-mitigation efforts.

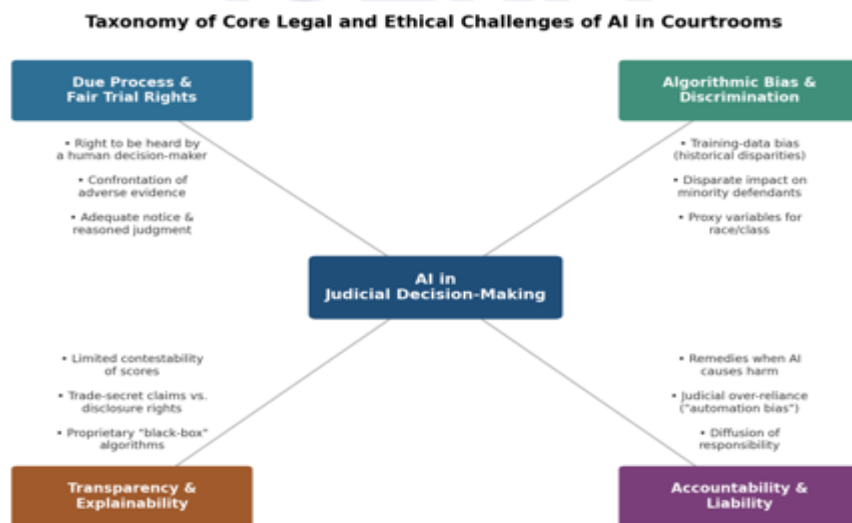


Figure 2. Taxonomy of core legal and ethical challenges of AI in courtrooms.

## 5. Comparative Regulatory Approaches

Jurisdictions have responded to AI's expansion into courts with markedly different regulatory philosophies, ranging from case-by-case judicial scrutiny to comprehensive ex-ante legislative classification.

### 5.1 United States: Case Law and Fragmented State Regulation

The United States lacks a comprehensive federal statute governing AI in courts. Regulation has instead developed through a patchwork of state-level legislation (such as requirements that risk-assessment vendors disclose validation studies), professional and judicial-conduct guidance, and case law such as *Loomis*. This produces significant variation in safeguards across states and leaves many questions—particularly around vendor transparency—unresolved at the federal level.<sup>2</sup>

### 5.2 European Union: Risk-Based Statutory Classification

The European Union has taken the most comprehensive legislative approach. The EU Artificial Intelligence Act (Regulation (EU) 2024/1689), which entered into force in 2024, establishes a tiered risk classification system. Under Article 6(2) and Annex III, AI systems intended to assist a judicial authority in researching and interpreting facts and law, or in applying the law to a concrete set of facts, are classified as high-risk, triggering obligations including risk-management systems, data-governance standards, technical documentation, human oversight, and accuracy and robustness requirements. Recital 61 of the Act explicitly states that AI tools may support but must not replace judicial decision-making power or judicial independence. Following the 2026 "AI Omnibus" political agreement, the compliance timeline for these high-risk obligations was extended, with Annex III obligations now scheduled to apply from December 2027 and Annex I obligations from August 2028, subject to formal adoption. This structured, ex-ante approach contrasts sharply with the U.S. model of after-the-fact judicial review.

### 5.3 India: Emerging Digital-Courts Infrastructure

India's approach has emphasized digital-infrastructure modernization—exemplified by the e-Courts Mission Mode Project and the Supreme Court's SUPACE (Supreme Court Portal for Assistance in Court's Efficiency) initiative—over comprehensive AI-specific legislation.

---

<sup>2</sup> US Policy 2025

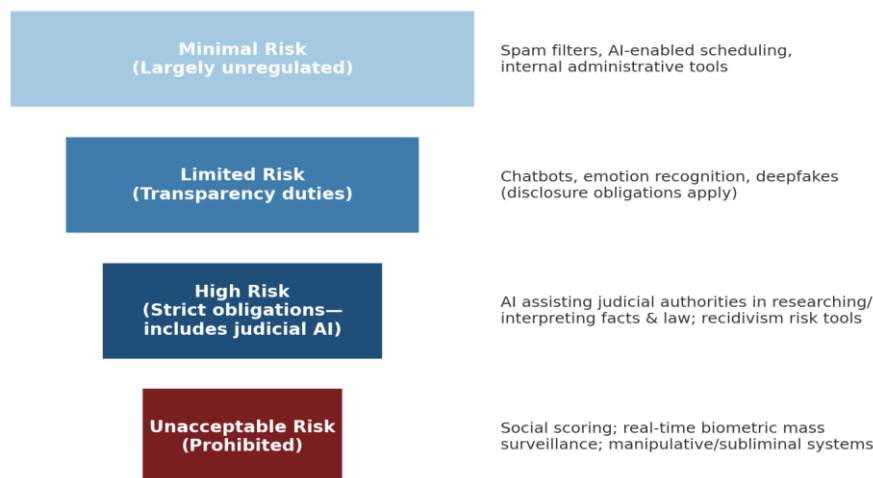
SUPACE is explicitly framed as a research and information-retrieval aid for judges rather than a decision-making tool, reflecting an official preference for keeping AI confined to a supportive role. As of mid-2026, India does not have a binding cross-sectoral AI statute comparable to the EU AI Act, though policy discussions on AI governance continue under the broader Digital India and data-protection law framework.

### 5.4 Comparative Summary

**Table 2. Comparative Regulatory Approaches to AI in Courts**

Jurisdiction	Primary Mechanism	Treatment of Judicial AI	Transparency Requirement
United States	Case law + fragmented state statutes	Permitted as advisory input; subject to due-process challenge	Limited; varies by state and vendor
European Union	EU AI Act (Reg. 2024/1689), Annex III	Classified as “high-risk”; human oversight mandated	Risk-management & technical documentation required
India	e-Courts policy + SUPACE guidelines	Confined to research/decision-support role by design	Administrative guidance; no binding AI-specific statute (as of 2026)

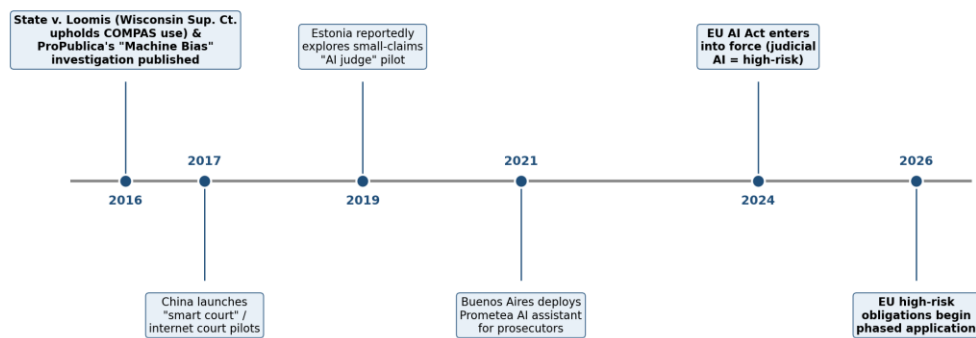
**EU AI Act: Risk-Based Classification Pyramid**  
**Judicial AI Tools Fall Under "High-Risk" (Annex III, Administration of Justice)**



Source: Regulation (EU) 2024/1689, Article 6 & Annex III; obligations phased in 2026-2028

Figure 3. EU AI Act risk-based classification pyramid; judicial AI falls within the “high-risk” tier.

Timeline: Key Milestones in AI and Judicial Decision-Making (2016-2026)



Note: The Estonia "AI judge" project was widely reported in 2019 media but officially denied by Estonia's Ministry of Justice in 2022, illustrating the gap between AI hype and verified deployment.

Figure 4. Timeline of key milestones in AI and judicial decision-making, 2016–2026.

## 6. Toward a Rights-Centered Governance Framework

Drawing on the legal and ethical challenges identified above, this paper proposes five governance principles for the responsible integration of AI into judicial processes.

### 6.1 Mandatory Algorithmic Transparency for Court-Facing Tools

Vendors supplying AI tools for use in or around courts should be required to disclose validation methodology, the categories of input variables used, and independently audited error and disparity rates, even where the underlying source code remains proprietary. Transparency of methodology and disparity metrics is distinguishable from, and does not require, disclosure of trade-secret source code.

### 6.2 Meaningful Human Oversight, Not Nominal Sign-Off

Human oversight requirements, as found in the EU AI Act, must be designed to counteract automation bias rather than merely formalize it. This implies training judicial officers on the statistical limitations of risk tools, structuring presentence reports to present uncertainty ranges rather than single scores, and building institutional mechanisms (such as periodic disparity audits) that create real friction against uncritical reliance.

### 6.3 A Right to Individualized Contestation

Defendants should have a clear, accessible procedural mechanism to contest an algorithmic score's relevance to their individual circumstances, separate from a general due-process challenge to the algorithm's existence. This could include the right to present individualized

mitigating evidence explicitly weighed against the score in the written record.

#### **6.4 Independent Pre-Deployment and Periodic Bias Auditing**

Tools should undergo independent bias and accuracy auditing before deployment and at regular intervals thereafter, with results made public in aggregate form, mirroring emerging practice in EU high-risk AI compliance and proposed U.S. state-level algorithmic accountability bills.

#### **6.5 Confining AI to Decision-Support, Not Decision-Replacement**

Across the jurisdictions surveyed, the most defensible uses of AI in courts are consistently those confined to a support role—research assistance, drafting aid, administrative triage—with the dispositive decision retained by an accountable human judicial officer. Recital 61 of the EU AI Act captures this principle directly: AI tools may support judicial decision-making power but should not replace it.

### **7. Conclusion**

Artificial intelligence is already embedded in judicial systems worldwide, most consequentially through risk-assessment tools that inform bail and sentencing decisions. The COMPAS controversy and the Loomis litigation demonstrate that these tools raise genuine, well-documented concerns about racial disparity, transparency, and the erosion of individualized judicial reasoning—concerns that are conceptual and structural, not merely the product of poorly engineered software. At the same time, comparative regulatory experience, particularly the EU's risk-tiered statutory framework, suggests that calibrated governance—mandating transparency, meaningful human oversight, contestability, and independent auditing—offers a more defensible path than either unregulated adoption or blanket prohibition. As AI capabilities continue to advance and judicial caseloads continue to grow, the central challenge for legal systems will be ensuring that efficiency gains do not come at the cost of the due-process guarantees that have historically defined what it means for a decision to be genuinely judicial.

### **References**

1. Angwin, J., Larson, J., Mattu, S., & Kirchner, L. (2016). *Machine Bias*. ProPublica.
2. Angwin, J., & Larson, J. (2016). *How We Analyzed the COMPAS Recidivism Algorithm*. ProPublica.

3. Chouldechova, A. (2017). *Fair prediction with disparate impact: A study of bias in recidivism prediction instruments*. *Big Data*, 5(2).
4. European Parliament and Council. (2024). *Regulation (EU) 2024/1689 (Artificial Intelligence Act), Article 6 and Annex III*.
5. Freshfields. (2026). *EU AI Act Unpacked #32: Draft Commission Guidelines on “High-Risk” AI*.
6. *State v. Loomis*, 881 N.W.2d 749 (Wis. 2016), cert. denied, 137 S. Ct. 2290 (2017).
7. Ulenaers, J. (2020). *The Impact of Artificial Intelligence on the Right to a Fair Trial: Towards a Robot Judge?* *Asian Journal of Law and Economics*, 11(2).
8. Estonian Ministry of Justice. (2022). *Public statement clarifying the status of the reported “AI judge” project*.
9. *International Journal for Court Administration*. (2024). *From Court Automation to e-Justice and Beyond in Europe*.
10. Wachter, S., Mittelstadt, B., & Russell, C. (2021). *Why fairness cannot be automated: Bridging the gap between EU non-discrimination law and AI*. *Computer Law & Security Review*, 41.

