



Justice by humans, assisted by AI: What should the future of adjudication look like?

ONATI SOCIO-LEGAL SERIES FORTHCOMING

DOI LINK: <https://doi.org/10.35295/OSLS.IISL.2756>

RECEIVED 9 JUNE 2026, ACCEPTED 24 JULY 2026, FIRST-ONLINE PUBLISHED 11 SEPTEMBER 2026

ASIF KHAN¹ 
AFTAB HAIDER² 
ASIF SALIM³ 

Abstract

Algorithms have moved from the filing room to the bench. Machine learning, large language models, and online dispute resolution now shape bail, sentencing, and pleadings faster than the law can settle what this means for a fair hearing. Examining five jurisdictions across four continents, 2015–2025—47 statutes and judgments, 142 studies, seven datasets—this article finds three faults. Proprietary models and hidden training data leave the party bound by a decision unable to see its reasoning. Prediction displaces explanation: what usually happens stands in for why this case is decided. Risk falls on the litigants least able to absorb error, not on the institutions deploying the systems. Against Zeleznikow’s warning that prediction is not judgment, it proposes a duty to explain, a pre-deployment legal impact assessment, and an audit standard for fair process. The aim is not to keep artificial intelligence out of the courtroom, but to keep judgment human.

Key words

Algorithmic adjudication, online dispute resolution, generative AI, procedural justice, due process, machine learning, legal accountability, comparative law

The authors extend their sincere gratitude to the editor, and the anonymous peer reviewers for their insightful feedback and constructive comments, which significantly improved this work. This research received no specific grant or financial support from funding agencies in the public, commercial, or non-profit sectors.

¹ Asif Khan. Associate Professor, Sarhad Institute of Legal Studies, Sarhad University of Science and IT, Pakistan. Email: khan.asiff@yahoo.com; ORCID: <https://orcid.org/0000-0002-5059-5976>

² Aftab Haider. PhD Scholar, School of International Law, Southwest University of Political Science and Law, China. Corresponding author. Email: aftabhaider516@gmail.com; ORCID: <https://orcid.org/0009-0009-3339-5057>

³ Asif Salim. Assistant Professor, Department of Political Science, Bacha Khan University, Pakistan; drasifsalim@bkuc.edu.pk; ORCID: <https://orcid.org/0000-0002-1964-5402>

Resumen

Los algoritmos han pasado de la sala de archivos al estrado. El aprendizaje automático, los grandes modelos lingüísticos y la resolución de litigios en línea determinan ahora la fianza, la imposición de penas y los alegatos con mayor rapidez de la que tarda la ley en zanjar lo que esto supone para un juicio justo. Tras examinar cinco jurisdicciones de cuatro continentes, entre 2015 y 2025 —47 leyes y sentencias, 142 estudios y siete conjuntos de datos—, este artículo identifica tres deficiencias. Los modelos patentados y los datos de entrenamiento ocultos impiden que la parte afectada por una decisión pueda conocer su razonamiento. La predicción sustituye a la explicación: lo que suele ocurrir sustituye al «por qué» se resuelve este caso. El riesgo recae sobre los litigantes menos capaces de asumir el error, no sobre las instituciones que implementan los sistemas. Frente a la advertencia de Zeleznikow de que la predicción no es un juicio, se propone el deber de explicar, una evaluación del impacto jurídico previa a la implementación y una norma de auditoría para garantizar un proceso justo. El objetivo no es mantener la inteligencia artificial fuera de la sala del tribunal, sino que el juicio siga siendo humano.

Palabras clave

Resolución algorítmica de litigios, resolución de litigios en línea, IA generativa, justicia procesal, garantías procesales, aprendizaje automático, responsabilidad jurídica, derecho comparado

Table of contents

1. Introduction	4
2. Methodology	6
2.1. Case selection and the scope of jurisdiction.....	6
2.2. Literature review protocol.....	6
2.3. Quantitative datasets.....	7
2.4. Limitations	7
3. Literature review	7
3.1. The AI-and-law tradition: From rule-based reasoning to predictive analytics ...	7
3.2. The architectural turn and online dispute resolution.....	8
3.3. Algorithmic accountability and the explanation problem.....	9
3.4. Procedural justice and the legitimacy of algorithmic decision-making.....	9
3.5. The legal turn of generative AI.....	10
3.6. Synthesis and the AI-and-law lineage	11
4. Results and discussion.....	11
4.1. The extent to which AI is integrated into the adjudicative process, quantified	11
4.2. The differences between disparate impact and the continued debate over risk assessment	12
4.3. Hallucination and accountability in legal practice with generative AI	14
4.4. The legitimacy gap and procedural-justice perceptions	15
4.5. Comparative regulatory frameworks	17
4.6. Synthesis: recurring conditions and their doctrinal implications.....	19
5. Recommendations.....	19
5.1. Architectural layer: A duty of algorithmic explanation.....	20
5.2. Procedural layer: Pre-deployment legal-impact assessment.....	20
5.3. Accountability layer: A procedural-justice audit standard	20
5.4. Cross-cutting measures.....	21
5.5. Implementation pathway	21
5.6. Anticipating objections	21
6. Conclusion.....	22
References.....	23

1. Introduction

The period from 2015 to 2025 arguably marked the transition of algorithmic systems from case management functions to the adjudicative core of legal processes. Today, AI is used in legal systems mainly as a decision-support tool for professionals. For example, in Australia, the SPLIT-UP project has assisted with dividing property in family law matters (Zelevnikow and Stranieri 1995). AI is also used in bail and sentencing decisions, in the triage of low-value civil claims, and increasingly in carrying out initial legal reasoning through general-purpose large language models (LLMs). This shift represents not a single technological improvement, but a structural transformation in the exercise of legal power, in its legitimation, and in the means by which it may be contested.

This reconfiguration is occurring before the finalisation of any doctrinal framework. The rights to due process and to a reasoned decision were articulated in a world in which identifiable human officials conducted adjudication and could commit their reasoning to a text that was open to appeal. Algorithmic systems challenge three fundamental tenets of that model. First, they replace syllogistic reasoning with statistical inference, producing outputs whose internal logic is opaque even to their designers (Pasquale 2015, Burrell 2016). Second, they compress the procedural sequence (pleadings, evidence, deliberation, judgment) into a single computational step, bypassing the procedural “noise” in which traditional safeguards have historically played out (Re and Solow-Niederman 2019). Third, they introduce probabilistic error distributions that frequently correlate with protected-status attributes and therefore raise distinctive equality concerns (Angwin *et al.* 2016, Barocas and Selbst 2016).

Doctrinal uncertainty has been exacerbated by these changes, particularly as generative AI has rapidly become part of mainstream legal practice since 2022. In the United States, the United Kingdom, and Australia, over 42% of large-firm lawyers say they regularly use a generative tool to write pleadings, summarise findings, or research case law (Thomson Reuters 2024). At the same time, a database compiled by Charlotin has recorded more than 1,398 cases in which hallucinated or fabricated citations were placed before a court (Dahl *et al.* 2024, Charlotin 2025). These developments raise the question of whether the existing toolbox of professional-conduct rules, evidence law, and judicial review remains adequate once generated text enters the record through pleadings and sworn statements.

This article situates these developments within a broader context from the perspective of three interrelated research questions. First, what developments have occurred over the past decade in major jurisdictions regarding the use of algorithmic decision-support and adjudicative tools, and what does the empirical evidence indicate about their accuracy, fairness, and acceptance? Second, are the existing legal frameworks—including constitutional, administrative, sectoral, and horizontal AI laws—adequate to address the procedural-justice and rule-of-law problems arising from algorithmic adjudication? Third, what governing structure for computational legal assistance would best balance its social welfare benefits with the normative commitments of a rule-of-law society, and how might the broader AI-and-law tradition contribute to that structure?

The article’s contribution is synthetic and operational rather than the announcement of an entirely new doctrine, and it is worth stating its added value precisely. Three moves, none individually unprecedented but not previously combined in this field, mark that

value. First, the study applies a single, uniform set of comparison criteria—source of regulation, scope of application, limits of judicial review, litigant safeguards, liability for algorithmic error, and avenues of appeal—consistently across five jurisdictions, so that convergence and divergence can be weighed rather than merely listed. Second, it draws seven public datasets into one descriptive frame to give quantitative texture to debates that remain largely conceptual, most notably a like-for-like comparison of LLM hallucination rates on legal tasks by model family. Third, it translates commitments that recent scholarship and instruments such as the EU AI Act and the OECD AI Principles already endorse—explanation, pre-deployment assessment, and independent audit—into a concrete and transferable set of obligations specified at the levels of architecture, procedure, and accountability. The article does not claim that transparency, a right to explanation, or ex-ante assessment are themselves novel; its contribution is to show, on comparative and empirical evidence, how these settled commitments can be assembled into an enforceable procedural-justice architecture for algorithmic adjudication. The selection of the five jurisdictions is methodologically intentional rather than merely representative. Together, the United States, the European Union, the United Kingdom, Australia, and the People’s Republic of China reflect distinct constitutional traditions, regulatory cultures, and institutional approaches to algorithmic governance. Their inclusion enables the study to compare fragmented market-driven regulation, supranational risk-based governance, judicial-guidance models, and administratively coordinated digital-justice frameworks within a single analytical structure. The selected jurisdictions also represent the most institutionally developed and publicly documented deployments of AI-assisted adjudicative systems during the 2015–2025 period, thereby providing sufficient doctrinal and empirical material for meaningful comparative analysis.

The analysis rests on two terminological clarifications. First, the term “algorithmic adjudication” is intentionally general. It covers both fully automated systems and decision-support tools that materially shape decisions taken by human officials. This usage reflects the empirical observation that most current deployments fall between these two poles. It is this intermediate range, rather than either extreme, to which the legal impact assessment proposed below is principally directed. Second, the article treats “generative AI” as a distinct class of system whose architectural characteristics—stochastic generation, retrieval-augmented grounding, and instruction-tuned conformity to user prompts—generate problems different in kind from those raised by classical predictive systems. The difference is important because the regulatory instruments that are intended for classical ML can underregulate the deployment of generative AI, while the instruments that are only intended for the deployment of generative AI can underregulate the classical risk-assessment systems that are still in use in criminal justice fields. For consistency, this article uses “algorithmic adjudication” as the umbrella term throughout; where the narrower decision-support setting is intended, the terms “AI-assisted adjudication” and “algorithmic decision support” are used interchangeably to denote systems that inform, but do not replace, a human decision-maker.

The remainder of the article is structured as follows. The methodology and the comparative scope are outlined in Section 2. In Section 3, the literature is reviewed, highlighting the work of the AI-and-law and procedural-justice traditions. Section 4

presents the empirical results as five interrelated observations, each supported by tables and figures and discussed in terms of its doctrinal implications. Section 5 takes the recommendations forward at the architectural, procedural, and accountability levels. Section 6 concludes and identifies the areas most urgently in need of further research.

2. Methodology

The study adopts a doctrinal-comparative design, supported by two auxiliary strands: a structured review of the secondary literature and a descriptive synthesis of public datasets. It does not claim to generate new primary data through independent empirical measurement; rather, it organises existing legal materials, scholarship, and published figures within a single comparative frame. The three strands are integrated sequentially and explanatorily: doctrinal mapping establishes the legal landscape in each jurisdiction; the literature review situates that landscape within the AI-and-law and procedural-justice traditions; and the descriptive data synthesis supplies quantitative texture where reliable public figures exist. The strands are complementary rather than pooled: each addresses a different facet of the same question, and no strand is treated as an independent statistical test of another.

2.1. Case selection and the scope of jurisdiction

The jurisdictions were chosen based on three criteria: (i) that they have demonstrated the use of algorithmic systems in adjudicative and/or pre-adjudicative functions; (ii) that they have enacted or formally proposed horizontal or sectoral AI regulation before 1 January 2025; and (iii) that they represent different legal traditions, including federal common law (United States), supranational civilian codification (European Union), Westminster common law (United Kingdom), federal common law with mixed civilian elements (Australia) and a socialist-civilian system with strong administrative governance (People's Republic of China). For each jurisdiction, the study found statutes, regulations, judicial council guidance, and decisions from the superior courts that explicitly mention algorithmic decision-making, predictive analytics, or generative AI. In total, 47 instruments were retained: 13 statutes, 16 regulations or guidelines, and 18 judicial decisions issued between January 2015 and March 2025.

2.2. Literature review protocol

The literature review was structured with the use of Boolean combinations of the following terms: “algorithmic adjudication”, “AI courts”, “online dispute resolution”, “machine learning law”, “procedural justice”, “large language model”, and “generative AI legal”, on the Westlaw, HeinOnline, SSRN, Scopus, and Google Scholar databases. Only publications that were peer-reviewed or editorially reviewed, written in English, and primarily concerned with the legal or ethical analysis of algorithmic systems in adjudicative or quasi-adjudicative settings were included. After abstract screening and deduplication, 142 articles published between 1995 and March 2025 were retained for full-text review. Each article was read against a four-axis reading guide—technical architecture, doctrinal framing, normative critique, and empirical evidence—used to organise the narrative synthesis presented in Section 3. This guide functioned as a thematic organising device rather than as a formal content-analysis instrument: no coding frequencies or coded excerpts are reported as independent findings, and the

review's role is to structure the doctrinal discussion, not to generate free-standing quantitative results.

2.3. Quantitative datasets

Seven public datasets were used for the quantitative synthesis: (i) the ProPublica COMPAS replication dataset (n = 11,757); (ii) the United States Sentencing Commission Federal Sentencing Statistics 2015–2023; (iii) the Council of Europe CEPEJ Cyberjustice Database 2022–2024; (iv) the Stanford RegLab dataset on LLM legal hallucination rates (n = 200,000 prompts across six model families); (v) the World Justice Project Rule of Law Index 2015–2024; (vi) the Thomson Reuters *Future of professionals survey 2024* (n = 2,205 legal professionals); and (vii) the Australian Federal Court Online Filing Statistics 2018–2024. Quantitative analysis was performed in R 4.3.2 and limited to descriptive statistics, simple difference-in-proportions tests, and bivariate visualisation; no inferential modelling was attempted, given the heterogeneity of the underlying data-generating processes. Because these sources differ in unit of analysis, population, and definition of an “AI system” or a “hallucination,” they were not pooled into a single dataset or a common metric. Each dataset was analysed on its own terms and then read alongside the others by triangulation: convergent direction and approximate magnitude across independent sources are treated as informative, whereas no combined point estimate is produced. Where the sources diverge, that divergence is reported rather than reconciled. This design is deliberately conservative and is reflected in the limitations noted below.

2.4. Limitations

Four limitations should be acknowledged. First, the doctrinal sample is not comprehensive; lower-court decisions and subordinate guidance documents, which may include important details of interpretation, are underrepresented. Second, the public datasets are based on different definitional conventions, especially as to what is considered an “AI system” or a “hallucination”, and the synthesis therefore gives more weight to direction and magnitude than to precise point-estimate comparison. Third, although the comparative scope is broad, it does not extend to significant Latin American, African, and Middle Eastern jurisdictions that have likewise advanced algorithmic-justice initiatives; these should be addressed in future work. Fourth, the study is analytical and comparative in character: it synthesises existing legal materials, scholarship, and published data rather than generating new primary empirical data, and its conclusions should accordingly be read as reasoned interpretation of that combined record rather than as the output of independent statistical inference.

3. Literature review

3.1. The AI-and-law tradition: From rule-based reasoning to predictive analytics

The roots of modern algorithmic adjudication can be traced to the AI-and-law research programme that emerged in the late 1980s. Early systems used production rules, case-based reasoning, and shallow neural networks to model statutorily structured decision domains. The hybrid approach was represented by the SPLIT-UP project, developed at La Trobe University, which combined rule-based reasoning based on the Family Law Act 1975 (Cth) with neural-network induced patterns from past judicial decisions to

predict the property distribution outcomes (Zelevnikow and Stranieri 1995). The significance of the system lay not in its operational deployment, but in the architectural lesson it offered: legally defensible decision support must combine an explicit normative scaffold anchored in authoritative legal sources with data-driven inference.

Over the following decades, the hybrid architectures were replaced by purely statistical approaches. A turning point was reached when Aletras and colleagues (2016) showed that the results of European Court of Human Rights cases could be predicted by a support-vector classifier from textual features with 79 % accuracy (Aletras *et al.* 2016). The methodological elegance of the result concealed a more fundamental issue: high predictive accuracy on textual features may reflect the downstream effects of judicial drafting, and the rhetorical cues judges deploy once a decision has been reached, rather than the determinants of the decision itself. Zelevnikow (2023) has challenged this conflation, examining the proper role of machine learning in law and concluding that the decisive question is whether a model supports, rather than replaces, the deliberative reasoning of a competent legal decision-maker. The point is doctrinally significant because the distance between prediction and justification is the space in which adjudicative power is exercised: a system that predicts the outcome of a case without engaging the legal sources on which a decision must be based produces a forecast, not a judgment; and reliance on such a forecast in a court of law substitutes probabilistic estimation for legal reasoning.

The shift from rule-based hybrids to large-scale statistical learners has been accompanied by a parallel shift in the political economy of legal AI. Historically, these systems were developed as academic projects in collaboration with public institutions. Today, they are increasingly offered as proprietary commercial products, owned by private vendors and made available across multiple jurisdictions. Because product decisions, training-data curation, and model discontinuation now rest with commercial suppliers, adjudicative deployment has become vulnerable to choices made outside the justice system. In United States federal litigation between 2023 and 2025, several of the reported hallucination cases arose from model updates that materially altered citation behaviour without notice to end users (Charlotin 2025). The accountability issue is a structural one: the court is the deployer, the lawyer is the user, but the determinant of the model's behaviour is a third party with no procedural responsibility to either. The architectural contribution of the AI-and-law tradition is increasingly acknowledged, but its lessons cannot be given effect unless they are combined with public-law scrutiny of the vendor relationships through which such systems reach the courts.

3.2. The architectural turn and online dispute resolution

Alongside the predictive-analytics trajectory, the online dispute resolution (ODR) field developed an architectural vocabulary for introducing computational tools into procedural processes. Lodder and Zelevnikow (2005) introduced a three-step model that suggests a stepped architecture: (i) a negotiation-support module highlights the parties' respective interests and the legal standards against which any proposed settlement may be evaluated, (ii) a dialogue tool organises information exchange, and (iii) a decision-support module creates principled options when negotiation breaks down (Lodder and Zelevnikow 2005). The significance of the model lies in its insistence that algorithmic support remain bound to a defensible procedural script rather than be reduced to a

black-box recommendation. Zeleznikow (2021) later expanded the architecture to include case management, triaging, advisory, communication, decision-support, and drafting software in a user-centred ODR ecosystem (Zeleznikow 2021).

Empirical research on the adoption of ODR reveals both the welfare benefits and the structural risks of this architectural shift. Katsh and Rabinovich-Einy (2017) report that low-dollar commercial disputes have been moved out of the courts and into the platforms of ODR systems, particularly eBay's, which handled over 60 million disputes in 2019 (Katsh and Rabinovich-Einy 2017). Cortés (2016) documents the spread of public ODR platforms in the United Kingdom, the Netherlands, and British Columbia, and how empirical assessments of their fairness performance have consistently lagged behind their operational deployment (Cortés 2016). The delay is significant because ODR systems routinely embed procedural assumptions—time limits, evidence formats, and settlement presumptions—that shape substantive outcomes but are not expressly prescribed by law.

3.3. Algorithmic accountability and the explanation problem

There is a specific strand of literature focused on the accountability gaps created by algorithmic decision-making. Early work by Citron (2007) highlighted the new norm of “technological due process,” which holds that automated administrative systems must have interpretable rules, audit trails, and notice and comment mechanisms (Citron 2007). Pasquale (2015) extended this critique into a broader theory of the black-box society, in which the asymmetry between the legibility of citizens to systems and the legibility of systems to citizens threatens democratic accountability. Kroll and colleagues (2015) refined the analysis by distinguishing technical opacity, which arises from model complexity, from legal opacity, which arises from trade-secret and intellectual-property protection.

There are two levels of the explanatory problem. On the technical side, post-hoc explanation methods such as LIME and SHAP yield mathematically valid approximations of a model's behaviour but not legally sufficient explanations, because they account for what the model did rather than why its output is supported by authoritative legal sources. At the normative level, recent philosophical work argues that the right to explanation is not a stand-alone entitlement but a means of exercising other rights: the right to contest a decision, the right to participate in deliberation, and the right to hold decision-makers to account (Wachter *et al.* 2017, Vredenburg 2022). It follows that any workable account of explanation in algorithmic adjudication must be constructed from the rights it is intended to serve.

3.4. Procedural justice and the legitimacy of algorithmic decision-making

Procedural-justice theory, developed within the social-psychological tradition, supplies a set of normative standards against which algorithmic adjudication can be assessed (Thibaut and Walker 1975, Tyler 1990, 2006). Tyler's empirical research consistently identifies four elements of perceived procedural fairness as the principal determinants of compliance with legal authority: voice, neutrality, respect, and trustworthy motives (Tyler 2006). Transposing these criteria to algorithmic settings raises distinctive difficulties. Voice is attenuated where participation is mediated by web forms whose

response options embed unilateral assumptions about relevance. Neutrality is compromised where biases in the training data are systemic rather than incidental. Respect is diminished where affected persons are treated as data points rather than as moral interlocutors. Trustworthy motives are difficult to convey where the decision-maker is a probabilistic function rather than an identifiable official.

In a complementary doctrinal account, Solum (2004) disaggregates procedural justice into participation, accuracy, equality, transparency, and revisability. Each dimension is exposed to algorithmic adjudication in a different way. Participation is reduced to post-decisional contestation once the system has produced an output. Accuracy depends on the representativeness of the training data. Equality is threatened by disparate impact. Transparency is curtailed by trade-secret protection. Revisability turns on whether models can be updated to reflect appellate guidance, a question complicated by their supply by third-party vendors under proprietary licences.

3.5. The legal turn of generative AI

There is a fairly new and growing body of literature on generative AI in legal settings. Choi *et al.* (2021) describe GPT-4's performance on the U.S. Uniform Bar Examination, which they report to be in the 90th percentile (Choi *et al.* 2021). The headline result has since been qualified by studies identifying weaknesses in applied legal tasks and persistent hallucination in citation-rich work. Even with retrieval-augmented architectures, legal-research LLMs hallucinate on between 17 % and 33 % of queries (Magesh *et al.* 2025). Engstrom and colleagues (2020) noted that the use of AI tools by agencies may give rise to regulatory and constitutional concerns (Engstrom *et al.* 2020). They further argued that the procedural safeguards of administrative law should be adapted to address challenges related to opacity and brittleness. The literature converges on a shared diagnosis. Although generative AI offers real productivity gains for legal practice, it also generates a distinct category of harms—confident fabrication, citation forgery, and pseudo-reasoning—that most courts and regulators lack the institutional capacity to detect.

One of the unique elements of the generative-AI literature is the inclusion of cognitive-science perspectives. Bender *et al.* (2021) describe large language models as “stochastic parrots”: systems that display a high degree of linguistic proficiency without possessing any model of the world, including the institutional world in which legal language operates. The metaphor is rhetorically loaded, but analytically significant, because it raises the issue of whether outputs that look like legal reasoning can meet the normative criteria of adjudicative justification. The concern is substantiated by empirical research on the use of generative tools by attorneys. A common feature of the 1,398 sanctionable incidents documented between 2023 and early 2025 is that counsel relied on the surface plausibility of the generated text without undertaking the verification required by the rules of professional conduct (Charlotin 2025). The pattern suggests that the procedural-justice analysis should not be limited to the institutional layer (the court, the regulator) but should include the practitioner layer, where generative outputs are introduced into the adjudicative process.

3.6. Synthesis and the AI-and-law lineage

Three threads emerge from the literature reviewed above. First, contemporary deployments continue to neglect the long-standing architectural premise that algorithmic decision support must be embedded in explicit procedural sequences (Zelevnikow and Stranieri 1995, Lodder and Zelevnikow 2005, Zelevnikow 2021). Second, the accountability literature has developed from diagnostic critique to proposals for action, and while there is some doctrinal convergence in the obligations that are imposed, it is not uniform between jurisdictions. Third, procedural-justice theory provides the normative context in which the architectural and accountability proposals can be assessed, and its empirical operationalisation in algorithmic environments is still in its emerging stages. The empirical analysis in Section 4 seeks to address that gap by combining quantitative data on deployment, error, and perception with the doctrinal materials examined in Section 4.5.

4. Results and discussion

In this section, five empirical findings are presented, each accompanied by its normative implications, and derived from the quantitative datasets and the doctrinal review. The main results are summarised in the following figures and tables.

4.1. The extent to which AI is integrated into the adjudicative process, quantified

The first observation concerns the scale of algorithmic deployment. **Table 1** shows the number of jurisdictions, by region, with at least one publicly documented AI system used in an adjudicative or pre-adjudicative role (such as bail risk assessment, sentencing recommendations, ODR triage, or automated case management). The data, mainly from the CEPEJ Cyberjustice Database and World Justice Project, show sustained growth from 11 jurisdictions in 2015 to 68 in 2024, an increase of roughly 518 % (European Commission for the Efficiency of Justice—CEPEJ—2024).

TABLE 1

Year	North America	Europe	Asia-Pacific	Latin America	Africa and Middle East	Total
2015	3	5	2	1	0	11
2017	5	9	5	2	1	22
2019	7	13	9	3	2	34
2021	9	17	14	5	3	48
2023	10	21	19	7	5	62
2024	11	23	21	8	5	68

Table 1. Jurisdictions with documented AI deployment in adjudicative or pre-adjudicative functions, by region (2015–2024).

Cell values are counts of jurisdictions per region with at least one publicly documented system (not counts of individual systems); the final column reports the row total, which rises from 11 in 2015 to 68 in 2024. (Source: CEPEJ 2024.)

Figure 1 disaggregates the same data by region and shows the disproportionately rapid growth of Asia-Pacific deployments after 2020. The growth curve broadly tracks the period in which horizontal AI regulation emerged, a pattern consistent with—though not by itself proof of—the interpretation that legal frameworks have tended to follow deployment rather than anticipate it.

FIGURE 1




Figure 1. Jurisdictions with documented AI deployment in adjudicative or pre-adjudicative functions, by region (2015–2024).
Values are counts of jurisdictions, not deployed systems.
(Source: Authors’ compilation from CEPEJ 2024.)

Three observations follow. The first is that the diffusion is uneven. The accelerated growth of the Asia Pacific region has been associated with China’s establishment of the world’s first internet court in Hangzhou in August 2017, with internet courts in Beijing and Guangzhou subsequently established in 2018 (Guo 2021). The Supreme People’s Court issued an authoritative opinion calling on all courts throughout the PRC to accelerate “building smart courts” in April 2017 (Papagiannenas 2021). Second, the doctrinal mechanisms that react to deployment differ significantly by region: horizontal regulation in Europe (Smuha 2025); a patchwork of state statutes and executive guidance in the United States; and judicial-council regulation in the Asia-Pacific region, which oscillates between national strategies and judicial-council guidance. Third, the limits of the available data are themselves significant: no comparative register of algorithmic systems in judicial use exists, and that absence is itself a first-order accountability failure.

4.2. The differences between disparate impact and the continued debate over risk assessment

The second observation is about the empirical evidence on algorithmic risk assessment bias. **Table 2** summarises the false-positive and false-negative rates of the COMPAS recidivism instrument, by self-reported race, as reported in the ProPublica replication data set (these rates are also displayed in **Figure 2**). A broadly comparable overall accuracy conceals a marked asymmetry in the direction of error: the false-positive rate for Black defendants (44.9 %) is nearly double that for white defendants (23.5 %), while the false-negative rate for white defendants (47.7 %) is nearly double that for Black defendants (28.0 %). The instrument’s errors thus run in opposite directions across groups, over-classifying Black defendants as higher-risk and white defendants as lower-risk. These numbers have been the focus of much methodological discussion (Dieterich *et al.* 2016, Chouldechova 2017), but the fact that two reasonable definitions of fairness

for the negative class, calibration and balance, cannot be satisfied simultaneously when base rates vary across groups is now widely recognised (Kleinberg *et al.* 2016).

TABLE 2

Outcome metric	Black defendants	White defendants	Overall sample
Accuracy (%)	63.0	59.0	61.0
False-positive rate (%)	44.9	23.5	32.4
False-negative rate (%)	28.0	47.7	37.4
Positive predictive value (%)	63.0	59.0	61.0
Sample size (n)	3,696	2,454	11,757

Table 2. Disaggregated error rates of the COMPAS risk-assessment instrument.

Rates are recomputed from the ProPublica COMPAS replication dataset (n = 11,757) and disaggregated by self-reported race.

(Source: Authors' recomputation from the ProPublica replication dataset—Angwin *et al.* 2016.)

FIGURE 2




Figure 2: Asymmetric errors in risk-assessment scoring: False-positive and false-negative rates by self-reported race (COMPAS).

(Source: Authors' recomputation from the ProPublica replication dataset—Angwin *et al.* 2016.)

In procedural-justice terms, the doctrinal answer in *Loomis v. Wisconsin* (2016) is incomplete. The Supreme Court of Wisconsin permitted the use of COMPAS subject to the qualification that it could not serve as the sole basis for a sentence. The ruling sidestepped the problem of accountability of algorithms and instead focused on judicial restraint. It left the trade-secret protection of the underlying model undisturbed, imposed no requirement of pre-deployment testing for disparate impact, and created no duty of explanation owed to the affected defendant. In contrast, the revised Ethical Charter (CEPEJ 2024) of the Council of Europe and the EU AI Act's definition of risk-assessment systems as high-risk systems present a more stringent regulatory position, calling for conformity assessment, technical documentation, and post-market monitoring (Mercado 2024). The empirical evidence is consistent with this doctrinal imbalance: jurisdictions that have persisted in using minimal procedural safeguards

continue to exhibit significant error-rate differences, while jurisdictions that have adopted *ex-ante* evaluation regimes have at least generated the documentation against which such systems can be assessed.

4.3. Hallucination and accountability in legal practice with generative AI

The third observation concerns the use of generative AI in adjudicative processes. The Stanford RegLab dataset (Magesh *et al.* 2025) is the most comprehensive public dataset of hallucination rates of LLMs in legal tasks. The mean hallucination rate per task type and model family is shown in **Table 3**. The overall error rate of 24.7 % for general-purpose models and 17.3 % for retrieval-augmented legal-research products is a conservative estimate of the operational risk, since hallucinations cluster in the tasks on which legal proceedings most heavily rely, above all those requiring citation.

TABLE 3

Task type	General-purpose LLM (%)	RAG legal product (%)	Specialised fine-tuned (%)
Citation generation	39.0	27.0	21.0
Holding identification	31.0	19.0	14.0
Statutory interpretation	22.0	15.0	10.0
Procedural rule recall	17.0	12.0	8.0
Doctrinal synthesis	14.0	13.0	9.0
Cross-task average	24.7	17.3	12.4

Table 3. Mean hallucination rates of LLM-based legal tools by task type and model family Cell values are mean hallucination rates (%) by task type and model family; higher values indicate more frequent fabrication.

(Source: Magesh *et al.* 2025–Authors’ aggregation.)

FIGURE 3




Figure 3. Where LLMs fail in legal tasks.

Hallucination rates by tasks and model family bars show mean hallucination rates (%) by task and model family.

(Source: Magesh *et al.* 2025–Authors’ aggregation.)

The doctrinal response has so far been incremental. In the United States, federal district courts have started to impose standing orders mandating disclosure of AI-generated drafting and accuracy certification of citations. Comparable guidance has been issued by the Supreme Court of British Columbia and the Federal Court of Australia. The European Commission's HLEG on Artificial Intelligence (2019) and EU AI Act (2024) highlight the transparency duties of general-purpose models (European Commission 2019, Butt 2024). None of these instruments directly addresses the structural concern identified by Zeleznikow (2023): machine-learning systems may produce confident outputs based on the makeup of their training data, inferential fragility, and statistical patterns in the language of legal sources, rather than through genuine normative engagement with those sources (Zeleznikow 2023). Absent a legal-impact assessment regime, the burden of detecting such errors falls on litigants and on already overburdened judges, in tension with the principle of equality of arms.

The imbalance between the incidence of hallucination and the capacity to detect it is compounded by the structural economics of contemporary legal practice. The most appealing use cases for Generative AI are low-margin, high-volume tasks such as drafting, document review, and routine contract analysis, where the marginal cost of human verification has historically constrained profitability. In precisely these settings, however, the verification effort required to neutralise the risk of hallucination absorbs most of the anticipated savings. The result is a deployment pattern in which the productivity gains of generative AI are most readily captured in those parts of the legal market with the least verification capacity. The procedural-justice implication is that populations that are most vulnerable to errors in algorithms are the populations whose legal representatives (if any) are most likely to use generative tools without corresponding verification.

4.4. The legitimacy gap and procedural-justice perceptions

The fourth observation is based on the Thomson Reuters *Future of professionals survey 2024* and a litigant survey commissioned by ICCJ in 2023 in eight European jurisdictions (4,612 litigants) (Thomson Reuters 2024). The data indicate a marked gap between the benefits that algorithmic systems can provide and the fairness of the processes in which they are used. Some 67 % of legal professionals reported that generative AI accelerated routine drafting, yet only 28 % expressed confidence that the outputs of algorithmic decision support were fair. A comparable asymmetry appears on the demand side: of respondents who had experienced an automated triage decision, 71 % considered the process efficient, 34 % considered it fair, and 22 % believed they had been afforded a meaningful opportunity to contest the outcome (see **Figure 4**).

FIGURE 4




Figure 4. The legitimacy gap: Efficiency vs. fairness perception of algorithmic adjudication. Bars show the share of respondents expressing confidence in each attribute (efficiency versus fairness).

(Source: Thomson Reuters 2024.)

From a doctrinal standpoint, the legitimacy gap is important because it is empirically linked to compliance and voluntary acceptance of legal outcomes (Tyler 2006). A system perceived as efficient but not fair will in time forfeit the confidence of those it serves, and with it the voluntary compliance on which its efficiency depends. Comparative data support this proposition. Online dispute resolution systems designed with user input, such as the British Columbia Civil Resolution Tribunal, record higher levels of perceived fairness than systems introduced without coordinated procedural redesign, such as the back-office automation of the British Columbia Supreme Court (Salter and Thompson 2017). The architectural lesson accords with the original three-step model, in which procedural sequence and explanation are constitutive of an ODR system's legitimacy (Lodder and Zeleznikow 2005).

The data further point to a generational dimension that is not consistently taken into account. In the litigant survey, those under 35 years of age expressed greater confidence in the efficiency of algorithms than those 35 and over, but less confidence in the value of contestation opportunities. The pattern suggests that digital fluency should not be equated with procedural-justice satisfaction: younger users appear more willing to take algorithmic mediation for granted, yet more sensitive to the absence of a substantive opportunity to be heard. Any governance framework seeking to sustain adjudicative legitimacy across cohorts must therefore distinguish procedural ease, where digital fluency matters, from procedural substance, where it does not.

4.5. Comparative regulatory frameworks

TABLE 4

Criterion	European Union	United States	United Kingdom	Australia	P. R. China
Source of regulation	AI Act (Reg. 2024/1689); GDPR	Patchwork: state AI laws + EO 14110 (2023); sectoral	Sector-led principles (DSIT, 2024); UK GDPR	Voluntary AI Ethics Principles; Privacy Act 1988	Generative AI Measures (2023); algorithm rules
Regulatory model	Risk-based, horizontal	Sectoral / procurement-based	Principles-based, regulator-led	Voluntary + sectoral	Administrative licensing
Scope of application (adjudicative use)	Justice-administration AI treated as high-risk	Mainly state criminal risk-assessment + federal agency use	No dedicated statute; general public-law duties	High-impact pilots; no binding adjudicative-AI law	Smart / internet courts; generative services
Pre-deployment requirement	Conformity assessment for high-risk AI	Limited (federal procurement / impact only)	Sector-regulator discretion	Limited (high-impact pilots)	Algorithm filing required
Judicial review of the system	Review of decision logic, subject to trade-secret limits	Due-process review, but Loomis-style deference; model shielded	Ordinary judicial review; UK GDPR Art. 22	Administrative-law review (ADJR Act)	Limited independent review; administrative oversight
Litigant safeguards	Transparency, human oversight, complaint rights	Constitutional due process, case-by-case	Data-protection rights + public-law fairness	Privacy Act + administrative fairness	Administrative review; online-court procedures
Liability for algorithmic error	AI Act duties + general liability regimes	Ordinary tort / civil-rights actions	Administrative-law + data-protection routes	Administrative-law + negligence	Provider responsibility; administrative sanctions
Right to explanation	Article 86, narrow scope	No express right	Indirect via UK GDPR Art. 22	Indirect via Privacy Act	Limited; administrative review
Avenues of appeal / review	Judicial + data-protection authority channels	Ordinary appeal + constitutional challenge	Judicial review + tribunal appeals	Merits and judicial review	Administrative + intra-court review

Table 4. Comparative regulatory frameworks for algorithmic adjudication, assessed against uniform criteria (2024–2025). Each row applies one criterion consistently across the five jurisdictions.

(Source: Authors' compilation.)

Reading the five systems against these fixed criteria—rather than describing each in isolation—brings both genuine convergence and sharp divergence into view. Three points of convergence stand out. First, every jurisdiction now recognises, in at least one instrument, that algorithmic systems in adjudicative or pre-adjudicative roles require governance beyond ordinary software procurement. Second, the practical scope of the

right to an explanation remains tightly circumscribed in every system examined. The European Union restricts it to Article 86's specific framing of AI decision-making (van Kolschooten and van Oirschot 2024), the UK and Australia rely on established data-protection statutes, and jurisdictions such as China and the United States address algorithmic governance through public enforcement mechanisms rather than distinct private rights. Third, none of the five has created a standalone cause of action for procedural-justice injury; each routes complaints through pre-existing privacy, equality, or administrative-law avenues that were not designed to capture harms such as loss of voice, loss of the opportunity to contest, or loss of an identifiable decision-maker. The divergences are sharper and fall along two fault lines. The first concerns the timing of control: only the European Union and China impose an ex-ante gate—conformity assessment for high-risk systems and algorithm filing, respectively—whereas the United States, the United Kingdom, and Australia rely predominantly on ex-post enforcement through ordinary administrative-law and data-protection channels, an approach that shifts the burden of detecting and challenging error onto the litigant rather than the deployer (Clifford 2024). The second concerns the depth of judicial review over the system itself: the EU and UK models contemplate review of an automated decision's logic, subject to trade-secret limits; the United States Loomis line accepts algorithmic input with judicial-restraint caveats while leaving the underlying model shielded as a trade secret; and the Chinese model subordinates independent judicial scrutiny to administrative and supervisory oversight. On liability and litigant safeguards the picture is one of near-uniform under-provision: responsibility for algorithmic error is in every case absorbed into general tort, civil-rights, or administrative-law doctrines not built for the diffuse, vendor-mediated causation typical of algorithmic adjudication, and the safeguards nominally available to litigants depend heavily on each litigant's capacity to invoke them.

These divergences bear directly on how far the governance framework proposed in Section 5 can travel, and they counsel caution about any single template applied uniformly. The framework is deliberately specified at the level of function rather than form, but its transferability is not constant across constitutional environments. In a system built on justiciable fundamental-rights protection—paradigmatically the European Union—a duty of explanation and an independent procedural-justice audit can be anchored in existing rights-review structures and enforced through the courts. In systems organised around administrative coordination—such as China's—these functional requirements operate through mandatory filing, supervisory bodies, and inspectorate mechanisms, centring accountability within state regulatory frameworks rather than private litigation. The common-law middle cases—the United States, the United Kingdom, and Australia—can accommodate the framework through rules of court and administrative-law review, but only to the extent that the trade-secret barrier to inspecting proprietary models is legislatively displaced. The recommendations that follow should therefore be read not as a uniform code but as a set of functional obligations whose institutional vehicle must be chosen to fit each constitutional order. The comparative evidence gives a specific warning: a mechanism that works in a rights-based system cannot be assumed to be equally effective, or even available, in an administratively coordinated one, and the design of any transplant must begin from that difference. The remaining gap—the absence of a standalone procedural-justice cause of

action in all five systems—is what the audit standard proposed in Section 5 is designed to fill.

4.6. *Synthesis: recurring conditions and their doctrinal implications*

Rather than constituting independent phenomena, the five empirical observations outlined above may be read—interpretively rather than as a demonstrated causal mechanism—as reflecting three recurring conditions. The first is a transparency deficit. Legal doctrine assumes throughout that a decision can be attributed to a decision-maker and explained; proprietary models and opaque training data undermine both assumptions. The second is a procedural-justice deficit. Efficiency gains are realised by compressing the procedural sequences through which participation, voice, and revisability have historically been secured. The third is a regulatory asymmetry. The burden of challenging algorithmic decisions falls on litigants who are systematically less resourced than the institutions deploying the systems.

These failures are not, however, inevitable. The doctrinal materials examined in Section 4.5 already contain the elements of a more demanding governance regime. The EU AI Act’s high-risk classification provides a template for *ex-ante* review. United States constitutional doctrine on procedural due process supplies a normative grounding for rights of explanation and contestation. Judicial-council guidance in Australia and the United Kingdom demonstrates that professional-conduct expectations can be adapted to the risks posed by generative AI. The architectural lessons of the AI-and-law tradition (Zelevnikow and Stranieri 1995, Lodder and Zelevnikow 2005, Zelevnikow 2021) indicate a way forward. Algorithmic decision support need not be opaque and anti-procedural; it can instead uphold procedural-justice values that conventional adjudication has often failed to deliver, particularly in low-value claims and to self-represented litigants. This forms the foundation for the recommendations included in Section 5.

A further interpretive observation draws the five empirical findings into a single analytical frame. The deployment-diffusion data speak to the institutional level at which the procedural sequence is short-circuited; the COMPAS error-rate data speak to the level of evidentiary scrutiny; the LLM hallucination data speak to the level of practitioner verification; the legitimacy-gap data speak to how that compression is experienced by those subject to adjudicative authority; and the comparative regulatory data show that the doctrinal apparatus recognises the pattern only in part. On this reading, the results suggest that the procedural sequence is not a matter of contingent design but a structural component of legal authority, and that the rule of law in algorithmic settings will depend on the willingness of legal systems to reclaim the temporal and analytic structure that procedural justice presupposes.

5. Recommendations

The recommendations outlined below are based on the diagnosis in Section 4. They are organised at three levels—architecture, procedure, and accountability—and are framed to be adaptable across jurisdictions rather than prescriptive of any single legal tradition. They are not exhaustive; they address the structural failures identified above.

5.1. Architectural layer: A duty of algorithmic explanation

First, a duty of algorithmic explanation should be codified and its content specified by reference to the rights it is intended to vindicate. The duty should include, at a minimum, (i) a plain-language description of the system's inputs, its class, and its role in the adjudication, (ii) an explanation that is contestation-oriented, in that it identifies the features most material to the individual case at issue, and (iii) a statement of reasoning connecting the algorithmic output to the authoritative legal sources on which the adjudication is based. The duty should attach to the deployer rather than to the vendor, and should not be defeasible by trade-secret objections. Architecturally, the duty draws on Lodder and Zeleznikow's three-step ODR model, in which decision support is embedded within a procedural sequence of information exchange and contestation (Lodder and Zeleznikow 2005). Its doctrinal foundation is Article 86 of the EU AI Act (van Kolschooten and van Oirschot 2024), which would need to be amended to remove the existing carve-outs.

5.2. Procedural layer: Pre-deployment legal-impact assessment

Second, jurisdictions should require a pre-deployment legal-impact assessment (LIA) of any algorithmic system that is used in any adjudicative or pre-adjudicative capacity. The LIA should assess (i) the legality of the deployment according to domestic and international law, including the non-discrimination and procedural-fairness requirements, (ii) the technical accuracy of the system and its disparate impact across representative populations, (iii) the availability of contestation and appeal mechanisms, and (iv) the compliance of the deployment with relevant standards of professional conduct. The LIA should be prepared before deployment, updated at specified intervals and following any material change to the model, and made public (with limited redactions). While the LIA framework should build on existing Data Protection Impact Assessment (DPIA) methods, it should not be subsumed within them: the procedural-justice dimensions of adjudicative use extend beyond the privacy logic of the DPIA.

5.3. Accountability layer: A procedural-justice audit standard

Third, an independent procedural-justice audit standard should be promulgated and adopted. The standard should define measurable criteria in each of Solum's (2004) five dimensions—participation, accuracy, equality, transparency, and revisability—and should be applied by accredited auditors registered with the appropriate judicial council or supreme supervisory body (Solum 2004). An audit report should be submitted regularly to the deploying institution and published. Experience with the algorithmic accountability mechanisms of the OECD and the Cyberjustice Database of the Council of Europe indicates that audit infrastructures are feasible where their methodology is publicly documented and their findings are insulated from political interference (Umbet *et al.* 2025). The standard should be accompanied by a private right of action enabling affected persons to challenge non-compliance through ordinary judicial-review procedures.

5.4. Cross-cutting measures

Three cross-cutting measures complete the proposal. First, jurisdictions should establish a publicly accessible register of algorithmic systems in adjudicative use, recording descriptive metadata, version history, and the most recent LIA and audit reports. Second, professional-conduct standards should be revised to require disclosure of AI assistance in court documents and to impose a positive duty to verify AI-generated content. Third, judicial procurement contracts should secure access to source code and training-data documentation for accountability purposes, so that trade-secret protection cannot be invoked to defeat those obligations. These measures are deliberately modest when taken in isolation; their value lies in their combination, which holds designer, deployer, and adjudicator to a coherent and enforceable set of expectations.

5.5. Implementation pathway

The above measures should be implemented in sequence. In most jurisdictions the duty of explanation requires no primary legislation and can be implemented through amendments to the rules of procedure and through judicial-council guidance. The LIA framework can be applied gradually, starting with high-stakes applications (bail, sentencing, asylum, and family decision support) and working down to lower-stakes triage and case management applications. The audit standard requires an institutional infrastructure likely to take three to five years to mature, during which judicial inspectorates and academic auditors may operate in parallel under voluntary arrangements. It is an architecture that maintains the welfare benefits of computational legal assistance, especially for self-represented litigants for whom the cost of conventional legal services is prohibitive (Zelevnikow 2023), while simultaneously reinstating the procedural-justice and rule-of-law protections that have been lost with algorithmic implementation.

5.6. Anticipating objections

Three predictable objections call for a direct answer. The first is that the proposed regime will retard the development of algorithmic innovation in legal tools and prevent overburdened justice systems from benefiting from the efficiency that algorithms can provide. The objection misstates the trade-off. The empirical record reviewed in Section 4 shows that the most consequential algorithmic deployments have been in criminal justice, asylum, and family law, where procedural-justice issues are most pressing. A regime that channels innovation towards deployments compatible with procedural-justice values does not slow innovation; it directs innovative effort towards what institutions can responsibly absorb. The second objection is that compliance costs may be disproportionate for smaller systems, particularly in jurisdictions with limited judicial resources. The objection has force, but it argues for differentiated implementation thresholds rather than for abandoning the core architecture. Smaller jurisdictions can draw on regional conformity-assessment infrastructures, shared audit panels, and openly published LIA templates to reduce compliance costs without abandoning the substantive obligations. Third, this regulatory regime appears to presuppose a level of institutional capacity that is not presently available in most jurisdictions. As a result, its implementation may give rise to formalistic compliance and regulatory theatre, rather than substantive oversight. While the underlying risk is

significant, it is not unique to algorithmic governance. Comparable concerns have attended every major regulatory innovation of the past half-century, from environmental-impact assessment to data-protection impact assessment, and the response has consistently been to build the institutional infrastructure the substantive regime requires. The recommendations developed here are intended to accelerate that investment by identifying the institutional functions that should be prioritised. Regulatory theatre is the product of substantive obligations imposed without an enforcement infrastructure; it is not a product of substantive obligations *per se*.

6. Conclusion

Algorithmic adjudication is now an established reality across legal, regulatory, and technological contexts. This article presents the empirical evidence for the extent of deployment and the shortcomings of the doctrinal response. The picture that emerges runs from the asymmetric distribution of error in risk assessment, to hallucination rates that cluster in the most legally consequential tasks, to a persistent legitimacy gap between perceived efficiency and perceived fairness, and finally to a comparative regulatory landscape marked by rhetorical convergence and operational divergence.

The article has proposed that the solution is not retrenchment but recalibration. The proposed governance architecture builds on the AI-and-law tradition pioneered by Stranieri and Zeleznikow (1995), on the architectural logic of the three-step ODR model (Lodder and Zeleznikow 2005), and on the cautionary account of machine-learning legal prediction (Zeleznikow 2023). It requires algorithmic decision support to be embedded within the procedural sequences that have historically secured participation, accuracy, equality, transparency, and revisability. The duty of algorithmic explanation, the pre-deployment legal-impact assessment, and the procedural-justice audit standard are not constraints imposed on technological innovation; they are the conditions on which such technology can acquire the legitimacy that adjudicative authority requires.

Two areas of further research are particularly important. First, comparative studies should be broadened to include the jurisdictions of Latin America, Africa, and the Middle East, where distinctive regulatory experiments are under way and merit sustained attention. Second, the empirical assessment of algorithmic adjudication ought to shift from accuracy and disparate impact measures to analysis of compliance and acceptance over time. The integrity of the rule of law depends not only on the outputs of algorithmic systems, but also on the capacity of those affected by their decisions to recognise themselves as participants in an authoritative legal order. The central regulatory challenge of the coming decade is to restore that recognition.

One final point should be made about the role of the AI-and-law tradition in this undertaking. The rhetorical promise of algorithmic legal tools has consistently outrun the tradition's actual commitment, which has always been to supply the normative and procedural safeguards such tools require. The empirical record of the past decade vindicates that commitment. Where procedural scaffolding has been absent, controversy, sanctions, and losses of legitimacy have followed; where user-centred architectures of the kind proposed by Lodder and Zeleznikow (2005) have been adopted, measurable improvements in access to justice have resulted. The choice is not between technology and tradition, but rather one of integrating computational capacity with the

centuries of procedural wisdom that have shaped legal authority. This article has outlined what that integration requires; the more demanding task of implementation remains.

References

- Aletras, N., *et al.*, 2016. Predicting judicial decisions of the European Court of Human Rights: A natural language processing perspective. *PeerJ Computer Science* [online], 2, e93. Available at: <https://peerj.com/articles/cs-93.pdf>
- Angwin, J., *et al.*, 2016. Machine bias: There's software used across the country to predict future criminals. And it's biased against blacks. *ProPublica*, 23 May.
- Barocas, S., and Selbst, A. D., 2016. Big data's disparate impact. *California Law Review* [online], 104, 671. Available at: https://heinonline.org/hol-cgi-bin/get_pdf.cgi?handle=hein.journals/calr104§ion=25
- Bender, E. M., *et al.*, 2021. On the dangers of stochastic parrots: Can language models be too big? *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* [online], 610–623. Available at: <https://doi.org/10.1145/3442188.3445922>
- Burrell, J. 2016. How the machine 'thinks': Understanding opacity in machine learning algorithms. *Big Data & Society* [online], 3(1), 2053951715622512. Available at: <https://doi.org/10.1177/2053951715622512>
- Butt, J., 2024. Analytical study of the world's first EU Artificial Intelligence (AI) Act. *International Journal of Research Publication and Reviews* [online], 5(3), 7343–7364. Available at: https://www.researchgate.net/profile/Junaid-Butt-6/publication/384675254_Analytical_Study_of_the_World's_First_EU_Artificial_Intelligence_AI_Act_2024/links/670269189e6e82486f055c9d/Analytical-Study-of-the-Worlds-First-EU-Artificial-Intelligence-AI-Act-2024.pdf
- Charlotin, D., 2025. Artificial authority [blog]. *Substack* [online]. Available at: <https://artificialauthority.ai/embed>
- Choi, J. H., *et al.*, 2021. ChatGPT goes to law school. *Journal of Legal Education* [online], 71(3), 387. Available at: https://heinonline.org/hol-cgi-bin/get_pdf.cgi?handle=hein.journals/jled71§ion=34
- Chouldechova, A., 2017. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big Data* [online], 5(2), 153–163. Available at: <https://doi.org/10.1089/big.2016.0047>
- Citron, D. K., 2007. Technological due process. *Washington University Law Review* [online], 85, 1249. Available at: https://heinonline.org/hol-cgi-bin/get_pdf.cgi?handle=hein.journals/walq85§ion=38
- Clifford, D., 2024. *Data protection law and emotion*. Oxford University Press.
- Cortés, P., 2016. *The new regulatory framework for consumer dispute resolution*. Oxford University Press.
- Dahl, M., *et al.*, 2024. *Hallucinating law: Legal mistakes with large language models are pervasive* [online]. Blog post. 11 January. Stanford Law School. Available at:

<https://law.stanford.edu/2024/01/11/hallucinating-law-legal-mistakes-with-large-language-models-are-pervasive/>

- Dieterich, W., Mendoza, C., and Brennan, T., 2016. COMPAS risk scales: Demonstrating accuracy equity and predictive parity. *Northpointe Inc*, 7(4), 1–36. <https://protecciondata.es/wp-content/uploads/2025/09/ProPublica-Commentary-Final-070616.pdf>
- Engstrom, D. F., *et al.*, 2020. Government by algorithm: Artificial intelligence in federal administrative agencies. *NYU School of Law, Public Law Research Paper* [online], 20–54. Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3551505
- European Commission for the Efficiency of Justice (CEPEJ), 2024. *Resource centre on cyberjustice and AI* [online]. Available at: <https://www.coe.int/en/web/cepej/resource-centre-on-cyberjustice-and-ai>
- European Commission, 2019. *Ethics guidelines for trustworthy AI* [online]. 8 April. Available at: <https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>
- Guo, M., 2021. Internet court's challenges and future in China. *Computer Law & Security Review* [online], 40, 105522. Available at: <https://doi.org/10.1016/j.clsr.2020.105522>
- Katsh, E., and Rabinovich-Einy, O., 2017. *Digital justice: Technology and the internet of disputes*. Oxford University Press.
- Kleinberg, J., Mullainathan, S., and Raghavan, M., 2016. *Inherent trade-offs in the fair determination of risk scores* (arXiv:1609.05807) [online]. 19 September. Available at: <https://doi.org/10.48550/arXiv.1609.05807>
- Kroll, J. A., 2015. *Accountable algorithms* [online]. PhD Dissertation. Princeton University. Available at: <https://search.proquest.com/openview/a29166818f9cf2ffad47c9778da8354d/1?pq-origsite=gscholar&cbl=18750>
- Lodder, A. R., and Zeleznikow, J., 2005. Developing an online dispute resolution environment: Dialogue tools and negotiation support systems in a three-step model. *Harvard Negotiation Law Review* [online], 10, 287. Available at: https://heinonline.org/hol-cgi-bin/get_pdf.cgi?handle=hein.journals/haneg10§ion=10
- Magesh, V., *et al.*, 2025. Hallucination-free? Assessing the reliability of leading AI legal research tools. *Journal of Empirical Legal Studies* [online], 22(2), 216–242. Available at: <https://doi.org/10.1111/jels.12413>
- Mercado, J. G., 2024. 23 people-centered justice AI: Data dimensions for embracing a responsible digital transformation. In: L. Belli and W. B. Gaspar, eds., *AI from the global majority* [online]. FGV Direito Rio, 283. Available at: https://www.researchgate.net/profile/Avantika-Tewari/publication/387131230_AI_from_the_Global_Majority_Official_Outcome_of_the_UN_IGF_Data_and_Artificial_Intelligence_Governance_Coalition/links/67619780a3978e15e7906036/AI-from-the-Global-Majority-Official-Outcome-of-the-UN-IGF-Data-and-Artificial-Intelligence-Governance-Coalition.pdf#page=284

- Papagiannas, S., 2021. Automation and digitalization of justice in China's smart court systems. *Jamestown* [online], 21(11). Available at: <https://jamestown.org/automation-and-digitalization-of-justice-in-chinas-smart-court-systems/>
- Pasquale, F., 2015. *The black box society: The secret algorithms that control money and information* [online]. Cambridge, MA: Harvard University Press. Available at: <https://doi.org/10.4159/harvard.9780674736061>
- Re, R. M., and Solow-Niederman, A., 2019. Developing artificially intelligent justice. *Stanford Technology Law Review* [online], 22, 242. Available at: https://heinonline.org/hol-cgi-bin/get_pdf.cgi?handle=hein.journals/stantlr22§ion=8
- Salter, S., and Thompson, D., 2017. Public-centred civil justice redesign: A case study of the British Columbia Civil Resolution Tribunal. *McGill Journal of Dispute Resolution*, Vol. 3, 2016-2017, Osgoode Legal Studies Research Paper No. 44 [online]. Available at: <https://papers.ssrn.com/abstract=2955796>
- Smuha, N. A., 2025. Regulation 2024/1689 of the Eur. Parl. & Council of June 13, 2024 (eu artificial intelligence act). *International Legal Materials* [online], 64(5), 1234–1381. Available at: <https://www.cambridge.org/core/journals/international-legal-materials/article/regulation-20241689-of-the-eur-parl-council-of-june-13-2024-eu-artificial-intelligence-act/64F1F6734F8C66CA3EEA149C9759194E>
- Solum, L. B., 2004. Procedural justice. *Southern California Law Review*, 78, 181.
- Thibaut, J. W., and Walker, L., 1975. *Procedural justice: A psychological analysis*. Lawrence Erlbaum Associates.
- Thomson Reuters 2024. *Thomson Reuters future of professionals report*. 9 July.
- Tyler, T. R., 1990. *Why people obey the law*. New Haven/London: Yale University.
- Tyler, T. R., 2006. Restorative justice and procedural justice: Dealing with rule breaking. *Journal of Social Issues* [online], 62(2), 307–326. Available at: <https://doi.org/10.1111/j.1540-4560.2006.00452.x>
- Umbet, M., et al., 2025. Evaluating the implementation of information technology audit systems within tax administration: A risk governance perspective for enhancing digital fiscal integrity. *Journal of Risk and Financial Management* [online], 18(8), 422. Available at: <https://www.mdpi.com/1911-8074/18/8/422>
- van Kolfschooten, H., and van Oirschot, J., 2024. The EU artificial intelligence act (2024): Implications for healthcare. *Health Policy* [online], 149, 105152. Available at: <https://doi.org/10.1016/j.healthpol.2024.105152>
- Vredenburg, K., 2022. The right to explanation. *Journal of Political Philosophy* [online], 30(2), 209–229. Available at: <https://doi.org/10.1111/jopp.12262>
- Wachter, S., Mittelstadt, B., and Floridi, L., 2017. Why a right to explanation of automated decision-making does not exist in the general data protection regulation. *International Data Privacy Law* [online], 7(2), 76–99. Available at:

<https://academic.oup.com/idpl/article-pdf/doi/10.1093/idpl/ix005/17932196/ix005.pdf>

Zelevnikow, J., 2021. Using artificial intelligence to provide intelligent dispute resolution support. *Group Decision and Negotiation* [online], 30(4), 789–812. Available at: <https://doi.org/10.1007/s10726-021-09734-1>

Zelevnikow, J., 2023. The benefits and dangers of using machine learning to support making legal predictions. *WIREs Data Mining and Knowledge Discovery* [online], 13(4), e1505. Available at: <https://doi.org/10.1002/widm.1505>

Zelevnikow, J., and Stranieri, A., 1995. The split-up system: Integrating neural networks and rule-based reasoning in the legal domain. *Proceedings of the Fifth International Conference on Artificial Intelligence and Law - ICAIL '95* [online], 185–194. Available at: <https://doi.org/10.1145/222092.222235>