Amato F, Ferraris AF, Gatt L, Giacalone M, Hagan MD, Molinari M, Steenhuis Q, Westermann H, 'Artificial Intelligence and Access to Justice: From Digital Assistance to Institutional Legitimacy' (2025) Special Issue 2 EJPLT.
DOI: https://doi.org/10.57230/EJPLT25SIJFAAFFLGMGMDHMMQSHW
Corresponding author. E-mail address: Marco.Giacalone@vub.be (M. Giacalone)
Pleinlaan no. 2, Brussels (BE)

EJPLT
ISSN 2704-8012

Artificial Intelligence and Access to Justice: From Digital Assistance to Institutional Legitimacy

Intelligenza artificiale e accesso alla giustizia: dall'assistenza digitale alla legittimazione istituzionale

Flora Amato
Associate Professor
Università degli Studi di Napoli Federico II

Andrea Filippo Ferraris
Joint Doctoral Researcher
Vrije Universiteit Brussel
University of Bologna

Lucilla Gatt
Full Professor
Director of Jean Monnet Centre of Excellence "ReCEPL4STAI"
Università degli Studi Suor Orsola Benincasa

Marco Giacalone
Research Professor
"Digitalisation and access to justice" (DIKE) research group
Vrije Universiteit Brussel

Margaret D. Hagan
DPhil, JD, MA
Stanford University

Marianna Molinari
Joint Doctoral Researcher
Vrije Universiteit Brussel
University of Bologna

Quinten Steenhuis
Senior Practitioner in Residence
Suffolk University Law School

Hannes Westermann
Assistant Professor
Maastricht Law and Tech Lab
Maastricht University

Abstract

This contribution provides an overview of the special issue of EJPLT, which gathers a selection of the best papers presented at the Workshop AI for Access to Justice, Dispute Resolution, and Data Access (AI4A2J@JURIX2025), focusing on the transformative role of artificial intelligence in the field of access to justice. It is argued, in fact, that AI is no longer conceived exclusively as a supporting tool available to justice systems, but rather as an environment in which legal meaning is progressively produced, interpreted, made accessible, and applied.

© The author(s) 2025, published by Suor Orsola Benincasa Università Editrice.
This contribution is licensed under a Creative Common Attribution 4.0 International Licence CC-BY-NC-ND, all the details on the license are available at:
https://creativecommons.org/licenses/by-nc-nd/4.0/

Abstract
Questo contributo presenta una panoramica del fascicolo speciale di EJPLT, che raccoglie una selezione dei migliori lavori presentati al Workshop AI for Access to Justice, Dispute Resolution, and Data Access (AI4A2J@JURIX2025), dedicato al ruolo trasformativo dell'intelligenza artificiale nell'ambito dell'accesso alla giustizia. Si afferma infatti che l'IA non si configura più esclusivamente come uno strumento di supporto a disposizione dei sistemi giudiziari, ma come un ambiente nel quale il significato giuridico viene progressivamente prodotto, interpretato, reso accessibile e applicato.

Keywords: Artificial Intelligence; digitalization; justice

Summary: 1. Background. – 2. Introduction. – 3. Preview. – 4. Conclusion.

1. Background.

The AIDA2J workshop (AI for Access to Justice, Dispute Resolution, and Data Access) was held in Turin as part of JURIX 2025. Since 1988, the Foundation for Legal Knowledge Based Systems – an association of law and computer science researchers from the Netherlands and Flanders – has organised the annual Legal Knowledge and Information Systems conference, the traditional forum for scholarly exchange within the European legal informatics community; the 2025 edition took place from 9 to 11 December.

Now in its fifth edition, the workshop grew out of a collaboration between the DIKE Research Group (DIgitalization and aCCEss to justice) at the Vrije Universiteit Brussel and the Jean Monnet Centre of Excellence ReCEPL4STAI (Research Centre in European Private Law for Sustainable and Trustworthy Artificial Intelligence) at the University of Naples Suor Orsola Benincasa, and brought together scholars and practitioners from institutions including Stanford University, Suffolk University, Maastricht University, the Vrije Universiteit Brussel, the University of Turin and the University of Naples Suor Orsola Benincasa.

The international character of the initiative was not limited to the affiliations of the speakers: it also shaped the method and the very object of the inquiry, which took the plurality of the legal traditions represented as its starting point. It was precisely from these differences that the discussion drew its depth and, in turn, its reach. The access to justice gap – which, according to World Justice Project estimates, affects some five billion people – is a global phenomenon, and it was treated as such. The contributions did not stop at the European legal space, or even at the Euro-Atlantic region, but extended the analysis to African and Latin American jurisdictions, where technological solutions have the greatest potential impact. Comparing experiences that have developed under profoundly heterogeneous institutional and infrastructural conditions gave the discussion a purchase beyond the single context in which each was elaborated, and yielded findings of value on a global scale.

Against this background, the workshop set out to assess whether, and how far, artificial intelligence can offer scalable responses to the gap, organising the inquiry along three lines. The first concerns access to legal data: the

construction of reusable and shareable datasets for research purposes through forms of international cooperation aimed at aggregating public data in machine-readable format. The second addresses the relationship between artificial intelligence and access to justice from the perspective of tool design – legal chatbots, rule-based systems, applications grounded in machine learning techniques – for self-represented litigants and legal aid programmes, with a view to resolving concrete disputes. The third concerns digital dispute resolution, surveying the state of the art in the use of AI in judicial proceedings and in alternative dispute resolution mechanisms (ADR/ODR), with the aim of reducing time to resolution, simplifying access to legal information and strengthening legal certainty.

These lines of inquiry took concrete expression in the papers presented, which ranged from the development of chatbots for citizen guidance to voice-interface services designed for disadvantaged settings, through to advanced ODR systems. One concern ran across all the contributions: the need to ensure the transparency, reliability and fairness of such systems – a necessary safeguard if automation is not to reproduce, or indeed amplify, the very inequalities it seeks to reduce.

2. Introduction.

This Special Issue originates from the selection of the best papers submitted to the AI for Access to Justice, Dispute Resolution, and Data Access workshop, held in the framework of JURIX 2025.

Its point of departure is a simple but increasingly difficult proposition: artificial intelligence is no longer merely an external instrument available to justice systems. It is becoming one of the environments in which legal meaning is produced, accessed, and acted upon. This shift matters because access to justice has never been only a question of physical entry into courts, nor even of the formal availability of legal information. It also concerns the conditions under which individuals can recognise a grievance as legally relevant, translate experience into institutional categories, understand the consequences of procedural choices, and decide whether to trust the actors and infrastructures that stand between them and a possible remedy.

Digitalisation, and especially generative AI, intensifies this problem. It may lower barriers of language, cost, distance and time. Yet it may also interpose new forms of opacity, standardisation and dependence. A person may receive an answer more quickly without being in a better position to understand why that answer is legally meaningful, whether it is reliable, which assumptions it embeds, or how it can be challenged. The central issue is therefore not simply whether AI can increase the supply of legal assistance. It is whether AI can alter the conditions under which people become capable of acting legally.

The six articles gathered in this Special Issue explore this ambivalence from different angles. They do not share a single jurisdiction, methodology or technological object. Their unity lies elsewhere: in the proposition that AI for access to justice must be understood as a chain of mediation. AI mediates legal language, legal-help work, legal knowledge, procedural pathways, settlement

expectations and community-based dispute resolution. At each point, it may bring the user closer to justice, or it may create a new layer of dependency between the user and the legal system. The difference turns on design, institutional setting and governance. This editorial, therefore, presents the contributions in an order that follows the justice journey from meaning to institutions.

3. Preview.

The opening article, by Andrea Filippo Ferraris and Marco Giacalone, provides the conceptual foundation for the volume. Its central move is to shift the debate from textual simplification to epistemic justice. The article argues that access to justice is not achieved merely by making legal language shorter, clearer, or more conversational. Users must also be able to recognise that their experience is legally relevant, connect it to institutional categories, understand available procedural options and evaluate the consequences of action or inaction. In this sense, access to justice includes epistemic accessibility: the capacity to situate oneself within the legal order.

This is a necessary correction to much of the current enthusiasm surrounding generative legal tools. Large language models can paraphrase statutes, generate explanations, structure procedural guidance and maintain adaptive dialogue. These capacities are not marginal. For many users, especially those without legal representation, the first encounter with the legal system may be less intimidating and easier to navigate. But linguistic fluency is not the same as legal empowerment. An explanation may be elegant yet incomplete; it may be accessible yet conceal its sources; it may be reassuring yet reproduce the dominant legal framing that made the user's experience difficult to express in the first place.

Ferraris and Giacalone capture this danger through the language of hermeneutical injustice and generative mediation. LLMs do not merely transmit legal information. They reformulate, prioritise and frame it. They therefore participate in the production of legal meaning. Where their outputs are probabilistic, opaque or insufficiently grounded in a jurisdictional context, they risk producing a new kind of epistemic opacity. The user receives language, but not necessarily the interpretive resources needed to understand the relationship between that language and a possible legal claim.

The design principles proposed in the article - situated linguistic adaptation, explainability and traceability, and participatory infrastructure - are accordingly more than technical recommendations. They are conditions of legal legitimacy. A system that supports access to justice must not merely speak to users in plain language. It must help them understand the legal significance of what is happening to them, while leaving visible the uncertainty, provenance and institutional limits of the guidance provided. This idea remains central to all the contributions that follow.

The second article, by Margaret D. Hagan, moves from interpretive agency to the practical architecture of legal help. Her Legal Help Task Taxonomy addresses a problem often underestimated in debates on AI and justice: before

deciding which legal tasks AI should perform, the field needs a shared vocabulary for describing what legal help actually entails. Legal aid organisations, court self-help centres, community justice workers, navigators, legal clinics and users themselves perform a wide range of activities. Yet these activities are often described through local, organisational or professional categories that do not travel easily across jurisdictions or service models.

Hagan's contribution is infrastructural in the strict sense. It does not introduce another tool; it provides a grammar for the field. By classifying tasks such as legal Q&A, issue-spotting, referral routing, action planning, intake, eligibility checks, document explanation, document preparation, negotiation, case management, administration and technology support, the taxonomy makes legal help visible as work. That visibility matters because AI risk is task-dependent. A system that explains a document raises different concerns from one that drafts a pleading, predicts an outcome, routes a vulnerable person away from a service, or shapes the tone of a negotiation.

The taxonomy also clarifies the institutional stakes of automation. Some tasks are informational or organisational. Others are relational, discretionary or trust-based. Some can be usefully supported by AI; others require professional judgment or human presence precisely because they involve vulnerability, credibility, strategy or responsibility. A credible agenda for AI in access to justice must therefore distinguish between automation, augmentation and delegation. Without this distinction, technological development risks treating legal help as an undifferentiated bundle of activities, when in fact its legitimacy depends on how particular tasks are allocated and supervised.

In the sequence of this Special Issue, Hagan's article follows naturally from the concern with epistemic justice. If the first contribution asks what users must be able to understand, the second asks what systems and intermediaries must be able to do. AI does not only mediate meaning; it also mediates labour. It reorganises workflows, relocates expertise, changes what is recorded and made measurable, and may alter the relationship between frontline workers and the communities they serve. The taxonomy makes these changes available for analysis rather than leaving them hidden behind the generic label of innovation.

The third article, by Simone Vagnoni, Lou-Anne Olivier, Coco Krumsick and Monica Palmirani, turns from legal-help tasks to legal-information infrastructure. CaseConnect addresses a different but closely related access-to-justice barrier: asymmetry in the ability to find and compare legal knowledge across jurisdictions. In human-rights litigation, environmental disputes and other transnational contexts, the problem is rarely the complete absence of legal information. It is the uneven capacity to locate, interpret and use relevant legal information across languages, legal traditions and institutional formats.

The article is important because it treats legal retrieval as a legal problem, not merely as an information-retrieval problem. Judicial decisions are not generic text. Their internal structure matters: facts, procedural history, legal reasoning, holdings and remedial consequences perform different functions. Across jurisdictions, these elements are expressed through different linguistic conventions and drafting cultures. A system that retrieves text nbased of

surface similarity may therefore produce results that appear relevant yet be legally misleading. Conversely, a system that respects legal structure can expand access to comparative knowledge without abandoning the contextual discipline of legal interpretation.

CaseConnect responds through a hybrid methodology. Legal experts identify structural patterns and sectioning conventions; automated processing then scales that analysis across larger datasets. This choice reflects a broader principle: in legal AI, human expertise should not appear only at the stage of after-the-fact validation. It should shape the architecture of the system from the beginning. The article also treats GDPR and EU AI Act requirements as design constraints rather than as an external compliance layer. Data protection, transparency, human oversight, and professional accountability become part of the infrastructure that enables access to legal knowledge.

Placed after Hagan's taxonomy, CaseConnect serves as a bridge. It shows that once legal work is made task-sensitive, the knowledge environment supporting those tasks must also become legally structured. Retrieval is not adjudication, however, and the article is careful not to collapse the two. Similarity is not relevance; relevance is not authority; and authority is not a substitute for argument. Human legal judgment remains necessary to determine whether a retrieved case is doctrinally, procedurally or jurisdictionally meaningful. In this respect, the article reinforces one of the volume’s recurring claims: AI can expand access to legal knowledge only if it does not conceal the interpretive work through which such knowledge becomes usable.

The fourth article, by Marianna Molinari and Marco Giacalone, brings the discussion into a concrete EU dispute-resolution setting: labour-law disputes concerning redundancy dismissals and the IDEA chatbot proposal. Its empirical basis is particularly valuable. Drawing on interviews and focus groups with legal professionals and policymakers in Belgium, Croatia, the Czech Republic, Estonia, Italy and Lithuania, the article identifies a legitimacy frontier that runs through contemporary debates on AI in justice. Stakeholders are generally receptive to bounded support functions, including legal research support, drafting assistance, anonymisation, translation, workflow management and procedural guidance. They become markedly more cautious where AI becomes outcome-facing and capable of shaping expectations, settlement behaviour or judicial reasoning.

This finding sharpens the distinction between assistance and displacement. The risk is not confined to systems that formally decide cases. A tool may have no binding legal force and still acquire practical authority. It may influence whether a worker settles, whether an employer litigates, whether a lawyer frames a claim in a particular way, or whether a judge's reasoning is subtly disciplined by a pattern presented as predictive knowledge. The legitimacy of AI in dispute resolution, therefore, cannot be measured only by asking whether the final decision remains human. It also depends on whether the system shapes the field of possible decisions before the human decision is made.

The IDEA chatbot is significant because it accepts this frontier rather than attempting to obscure it. It is framed not as an adjudicative predictor, but as a transparency-oriented instrument for procedural and doctrinal guidance in

redundancy-dismissal disputes. Its purpose is to reduce information asymmetries and make options more intelligible without claiming authority over the merits of the case. This is a more modest design ambition, but also a more legally serious one. It recognises that in labour justice, where disputes are fact-sensitive and socially consequential, human oversight is not a decorative safeguard. It is an institutional condition for preserving responsibility, contestability and individualised reasoning.

The fifth article, by Sergio Andres Escobar Repullo, deliberately moves to the edge of that legitimacy frontier. Focusing on environmental damage claims before Chile's Second Environmental Court, the article begins with an empirical diagnosis of procedural inefficiency. It identifies the conciliation phase as a major source of delay and limited effectiveness, reporting that proceedings reaching final judgment last, on average, 837 days, while conciliation lasts, on average, 267 days, nearly one-third of the total duration. The diagnosis is important because it links AI not to abstract modernisation, but to a concrete procedural bottleneck within an existing justice pathway.

On that basis, the article explores whether AI-based judicial-outcome prediction could support more efficient conciliation. The proposal is both attractive and difficult. If parties and judges had access to reliable estimates of likely outcomes, negotiations might become more informed, strategic uncertainty might be reduced, and settlement might occur earlier. Delay, after all, is not a neutral inconvenience. In many disputes, it is itself a denial of effective justice. Yet prediction in conciliation does not merely inform. It may anchor expectations, redistribute bargaining power and affect what parties perceive as a reasonable settlement. A prediction tool may therefore shape legal behaviour even when it does not decide the dispute.

This contribution is valuable precisely because it forces the access-to-justice argument into its hardest terrain. It would be too easy to say that outcome-facing AI should simply be avoided. It would be equally unsatisfactory to treat prediction as a neutral aid to efficiency. Escobar Repullo instead places the question where it belongs: within the procedural structure, legal culture and data environment of the jurisdiction concerned. The article's insistence on regionally trained models is crucial. Predictive AI cannot be imported as a generic product. It is built from local data, categories, judicial practices, and institutional assumptions. Without that embeddedness, prediction risks becoming a technologically sophisticated form of legal mistranslation.

Read after the IDEA article, the Chilean contribution does not abandon caution; it tests its limits. It asks whether carefully delimited predictive assistance may be justified when a concrete procedural dysfunction has been identified, and safeguards are designed for the institutional context. The answer is not given as a simple endorsement. Rather, the article invites a more precise inquiry into function, evidentiary basis, cultural fit, transparency, contestability and the risk of behavioural steering. That inquiry is indispensable if access-to-justice scholarship is to engage seriously with predictive technologies rather than merely denounce or celebrate them.

The final article, by Viviana Brun, Ronald Lenz and Simon Diaz Perez, brings the Special Issue to its most people-centred and socio-legally grounded setting. Their study of Hadin kai-bani fondo in Niger begins not with courts,

professional databases, or legal-service organisations, but with actual justice pathways in a legally plural and low-connectivity environment. This is an important closing movement. It reminds us that access to justice is not defined by the institutions lawyers most readily imagine, but by the places people actually go when they experience a land dispute, a family problem, a shared-resource conflict or a local grievance.

In Niger, customary and locally recognised authorities play a central role in dispute resolution, particularly in land and shared natural-resource matters. The design challenge is therefore not how to replace these actors with formal digital systems, but how to support them through reliable, nationally validated guidance. Hadin kai-bani fondo does this through a free phone service and a retrieval-augmented generation architecture restricted to a curated Niger-specific corpus. The system is voice-first, retrieval-bounded and institutionally embedded. These features should not be read as technical details. They are legitimate features.

Voice-first design lowers literacy and connectivity barriers. Retrieval-bounded generation reduces the risk of hallucination and keeps answers within validated local content. Embedding the service within customary and administrative pathways respects the intermediaries whom users already trust. The article, therefore, offers a powerful corrective to court-centred and text-heavy models of legal AI. In many settings, the access-to-justice question is not how to bring users into digital courts, but how to improve the consistency, reliability and usability of the justice pathways they already use.

As a closing contribution, the Niger article re-humanises the discussion. It shows that AI for access to justice is not necessarily most advanced when it is most technically complex. It is most meaningful when it is aligned with people's needs, languages, infrastructures, trusted intermediaries and legal pluralism. Its model of augmentation rather than displacement captures the normative centre of the Special Issue: AI should strengthen the human and institutional ecology of justice, not bypass it.

4. Conclusion.

What emerges across the six contributions is less a theory of automation than a theory of mediated legality. Each system examined in this issue stands between a person and a legal possibility: understanding a claim, identifying a task, retrieving a precedent, navigating a procedure, estimating a settlement position, or obtaining guidance from a trusted local authority. At each point, the same questions return in a different form. Does the system make legal meaning more intelligible, or merely more fluent? Does it clarify legal-help work, or conceal labour? Does it retrieve legal knowledge, or flatten legal relevance? Does it guide users, or steer them? Does it support institutional actors or displace accountability? Does it adapt to legal culture, or impose generic assumptions?

The conditions of legitimacy are accordingly not reducible to technical accuracy. They include the provenance of legal information, the visibility of uncertainty, the preservation of professional and institutional responsibility,

the possibility of challenge, and the fit between a system's design and the social setting in which people actually seek justice. These conditions are not secondary safeguards to be added once innovation has already taken place. They are constitutive of legality in the digital transformation of justice.

For policymakers, the implication is that AI governance in justice must be function-sensitive. The relevant question is not whether AI, as a broad category, is permitted or prohibited. The relevant question is which function is being performed, by whom, for whom, with what data, under which institutional responsibility and with what possibility of correction or contestation. A document explainer, a referral router, a cross-jurisdictional retrieval tool, a procedural chatbot, a prediction aid and a voice-first customary justice service do not raise the same risks. They should not be governed as if they did.

For technologists, the implication is equally direct. Legal AI must be built with a legal structure in mind. Legal meaning depends on jurisdiction, procedure, document architecture, institutional role, evidentiary context, language and social practice. Generic systems are therefore insufficient. Retrieval must preserve context; generation must be bounded; interfaces must disclose limits; datasets must be scrutinised; and evaluation must include not only accuracy but also comprehension, agency, trust, fairness, vulnerability, and institutional effects.

For scholars, the volume points to a research agenda that cannot be contained within a single discipline. Doctrinal law alone cannot explain how users understand legal systems. Technical performance alone cannot measure justice. Socio-legal studies, legal linguistics, human-computer interaction, computational law, empirical justice research, comparative law and AI governance must be brought together. The articles in this Special Issue show how such work can be done without losing normative clarity.

Ultimately, the promise of AI for access to justice is not that machines will resolve legal problems on behalf of people. It is that carefully designed systems may help people recognise their problems, understand their options, access relevant knowledge, engage trusted intermediaries and pursue fair outcomes with greater confidence. The danger is that AI becomes a new architecture of opacity, presenting itself as assistance while quietly reshaping legal agency and institutional authority.

The future of digital justice will therefore be judged not by the sophistication of its models, but by the quality of the legal relationships they mediate. If AI can make rights more intelligible, procedures more navigable, institutions more responsive and outcomes more contestable, it may strengthen access to justice. If it obscures uncertainty, standardises discretion, weakens accountability or displaces human judgment, it will reproduce the very barriers it claims to solve. The contributions gathered in this Special Issue offer a rich and rigorous map for pursuing the first path while guarding against the second.