

Working Paper presented at the

Peer-to-Peer Financial Systems 2026 Workshop

June 2026

Judicial Empathy and Machine Assistance: Legal, Ethical and Governance Constraints on AI in Adjudication

Alan Midgley

AG Midgley Ltd

Alastair Moore

University College London

Dominic Courage

-

Flavia Kenyon

33 Chancery Lane

Geoffrey Goodell

University College London

Karl Reimer

University of Zurich

Michael Mainelli

Z/Yen

Nick Davies

IC4DTi

Rasto Koval

Cequence.io

Ryland Ash

Watson Farley & Williams

Stacy Sinclair

Fenwick Elliott

Stephen Dowling

TrialView

Powered by



P2P Financial Systems



Judicial Empathy and Machine Assistance: Legal, Ethical and Governance Constraints on AI in Adjudication

Alan Midgley, AG Midgley Ltd
Dr Alastair Moore, UCL
Dominic Courage
Flavia Kenyon, 33 Chancery Lane
Dr Geoffrey Goodell, UCL
Karl Reimer, University of Zurich
The Rt. Hon. Lord Neuberger of Abbotsbury
Professor Michael Mainelli, Z/Yen
Nick Davies, IC4DTi
Rasto Koval, Cequence.io
Ryland Ash, Watson Farley & Williams
Dr Stacy Sinclair, Fenwick Elliott
Stephen Dowling SC, TrialView

March 2026

Abstract

This paper examines the legal position of empathy in judicial decision-making and the extent to which artificial intelligence (AI) systems may legitimately assist (or ever replace) human adjudicators. It is prompted by recent practical experimentation with generative AI in the production of judicial reasons and by a stakeholder roundtable held in London in July 2025 (Appendix A). Drawing on legal doctrine, interdisciplinary scholarship on emotion and judgment, and emerging AI governance regimes, the paper argues that empathy is not a freestanding legal test but is nonetheless legally consequential as an element of procedural justice: it supports the human qualities of fair hearing, equal treatment, and reasoned decision-making that underpin legitimacy. AI can support forms of procedural empathy (for example, improving comprehension, consistency, accessibility of reasons, and identification of missing factual context), but current law and constitutional principle require that adjudicative authority remain human, accountable, and publicly answerable. The paper proposes a governance and design framework for empathy-supporting judicial AI that is compatible with open justice, contestability, and auditability, and outlines a research agenda for evaluating whether such tools improve outcomes for court users without undermining independence or impartiality.

Keywords: judicial empathy; procedural justice; adjudication; generative AI; automated decision-making; explainability; auditability; EU AI Act; Article 6 ECHR; UK GDPR; ISO/IEC 42001.

1 Introduction

Courts and tribunals are under sustained pressure to deliver timely, intelligible and accessible justice in the face of growing caseloads, increasing litigant-in-person participation, and rising complexity of evidence. In parallel, the public release of large language model (LLM) systems has accelerated the uptake of generative AI across legal practice, including early judicial experimentation in drafting and summarising decisions. This convergence has reignited two long-running debates: first, what role, if any, empathy should play in judicial decision-making; and second, whether (and under what constraints) a machine may be involved in the adjudicative process.

In common-law discourse, “judicial empathy” is frequently invoked as either a virtue (enabling understanding of the lived consequences of legal rules) or a vice (risking partiality, sentimentality, or bias). The term often functions as a proxy for deeper questions about the rule of law: Is judging a technical exercise in applying rules to facts, or an interpretive practice that necessarily involves moral judgment and a situated understanding of human experience? Scholars such as Bandes have argued that empathy, defined as perspective-taking, is an essential judicial attribute while remaining distinct from outcome-driven sympathy Bandes [2]. Others, including critics in the post-2008 “anti-empathy” turn, caution that empathy can be selective, politically instrumentalised, and in tension with ideals of impartiality West [19].

AI complicates this debate in two ways. First, it offers tools that can assist with tasks historically performed by judges and court staff (for example, summarising case papers, extracting issues, or generating plain-language explanations). Second, it raises a more disruptive possibility: that systems could influence or determine the outcome of a dispute, whether directly (automated adjudication) or indirectly (decision support that becomes *de facto* determinative). Each possibility triggers questions of legality (fair trial and reasons), legitimacy (public trust and open justice), and governance (accountability for error or bias).

This paper is informed by a July 2025 stakeholder roundtable on “AI and judicial empathy” held at the International Dispute Resolution Centre, London, and captured in an anonymised summary (Appendix A) Midgley and Sinclair [15]. The workshop is treated as a practice-oriented starting point rather than as an academic authority. The paper’s aim is to connect those practitioner insights to the best available scholarship and to emerging governance instruments relevant to judicial AI.

The central claim developed below is that empathy is not a freestanding legal criterion that can be mechanically optimised. Instead, in lawful adjudication empathy matters primarily as a component of procedural justice: the ability to understand parties’ circumstances, treat them with dignity, and give reasons that show the parties were heard. AI can support this procedural dimension, but only under a governance model that preserves human responsibility, contestability, and transparency.

The analysis proceeds through: (i) a conceptual clarification of empathy and

its relationship to impartiality; (ii) a doctrinal account of the legal “position” of empathy through the lenses of fair hearing, reasoned decision-making, and equality; (iii) an examination of current judicial guidance and early case-law on AI use in decisions; (iv) an assessment of whether AI can simulate or support empathy and what risks follow; and (v) a practical framework for courts and policymakers.

2 Workshop insights and research approach

The July 2025 roundtable brought together judges, academics, technologists and practitioners to discuss the relationship between empathy, legitimacy, and machine assistance in adjudication. The anonymised summary (Appendix A) highlights a broadly shared position: current AI systems cannot “feel” empathy, but may be trained to recognise contextual cues and respond in ways that appear empathetic; nevertheless, participants favoured a model in which AI assists rather than decides, with strong emphasis on transparency, auditability and accountability Midgley and Sinclair [15].

This paper uses the roundtable summary as qualitative input for identifying the most salient questions, risks and governance concerns expressed by stakeholders. It then undertakes a doctrinal and literature-based analysis to (a) test whether those concerns map onto legal requirements and (b) identify where existing legal frameworks are under-specified for judicial AI.

Methodologically, the paper combines:

- Doctrinal legal analysis (domestic public law principles, fair trial norms, data protection and emerging AI regulation).
- Interdisciplinary literature review (legal theory on empathy and impartiality; psychology of empathy; human-computer interaction).
- Case study analysis of early judicial practice around AI-assisted drafting and of comparative experiences with algorithmic decision support.
- Governance synthesis (standards, risk management frameworks and institutional assurance mechanisms relevant to courts).

The scope is intentionally focused. The paper does not attempt to resolve the full philosophy-of-law debate about whether judging is ultimately rule-bound or discretionary. Nor does it evaluate a particular AI product. Instead it aims to develop a legally grounded framework that courts can use to evaluate any proposed AI tool purporting to support judicial reasoning or communication, including tools marketed as “empathetic”.

3 What is judicial empathy? Conceptual and normative accounts

3.1 Definitional clarifications

Empathy is a contested term. In psychology it is often broken into at least two components: (i) cognitive empathy (the ability to understand another person’s perspective, intentions or emotions), and (ii) affective empathy (an emotional resonance or “feeling with” another). Some accounts also distinguish compassion as a prosocial motivational state (a desire to alleviate suffering) rather than mere resonance. Legal debates often blur these distinctions, sometimes using “empathy” to mean compassion for disadvantaged litigants, and at other times to mean perspective-taking across all parties.

For judicial purposes, the most legally tractable notion is cognitive empathy: the disciplined capacity to grasp the practical stakes of a dispute, the significance of evidence for the parties, and the social context within which legal rules operate. This can include an awareness of how trauma, disability, language barriers, or power imbalances may shape a person’s ability to give coherent testimony or participate effectively. Importantly, this does not entail adopting any party’s interests as one’s own.

3.2 Empathy, impartiality and the rule of law

The relationship between empathy and impartiality is frequently misunderstood. Critics often equate empathy with partiality: the fear is that by “identifying” with one party a judge will distort legal analysis. However, empathy-as-perspective-taking can be framed as a tool for impartiality rather than its enemy. If judges are inevitably situated, with their own experiences and blind spots, then awareness of other perspectives may counteract selective empathy and implicit bias. Bandes, for example, argues that empathy helps judges to understand the stakes for all litigants but does not resolve which claims should prevail Bandes [2].

The normative question therefore becomes: what kind of empathy is compatible with the rule of law? A common answer is that empathy is legitimate when it operates at the level of understanding and communication but illegitimate when it becomes a substitute for legal reasons. In this sense, empathy belongs to the craft of judging (how evidence is heard, how reasons are expressed) rather than as a separate decision rule.

3.3 Procedural justice as a legal lens for empathy

A productive way to operationalise empathy for legal and technological purposes is through procedural justice. A large empirical literature associates perceived legitimacy of legal authorities with process-based factors: whether people have “voice”, are treated with respect, perceive neutrality, and trust decision-makers’

motives Tyler [17]. These factors map closely onto what court users often describe as “being heard” and “being treated as a person” – concepts that are close cousins of empathy even when the term is not used.

Procedural justice does not require that decision-makers share parties’ emotions. It requires that decision-makers provide a process that is intelligible and respectful. A judge may be entirely correct on the law yet still fail in legitimacy terms if reasons are dismissive, opaque, or inattentive to a party’s arguments. Empathy, reconceptualised as disciplined perspective-taking and respectful engagement, can therefore be seen as part of the institutional function of courts: securing acceptance of outcomes in a pluralistic society.

4 The legal position of empathy in adjudication

4.1 No standalone doctrine of empathy, but multiple legal touchpoints

There is no general rule in common-law systems that a judge must “be empathetic” as an express legal duty. Yet empathy is legally relevant in at least three ways. First, fair trial and natural justice principles require genuine consideration of parties’ arguments and the giving of reasons. Second, equality and non-discrimination norms require courts to avoid stereotyping and to make reasonable procedural adjustments so that parties can participate effectively. Third, the legitimacy of adjudication depends on public confidence, which is influenced by the perceived humanity and respectfulness of the process.

4.2 Fair hearing and reasoned decisions

The duty to provide reasons is a cornerstone of lawful adjudication. Reasons demonstrate that the decision-maker has understood and evaluated the parties’ arguments, and they enable meaningful appeal and public scrutiny. Human rights jurisprudence under Article 6 of the European Convention on Human Rights has repeatedly linked fair trial with reason-giving. European Court of Human Rights [8], while recognising that the extent of reasons required varies with the nature of the decision. The practical implication is that a “human” element of judging is partly expressed through reasons that engage with what mattered to the parties.

4.3 Equal treatment and dignitarian elements

In England and Wales, the judiciary’s Equal Treatment Bench Book (ETBB) illustrates how courts operationalise dignity, participation and fairness in practice. Judicial College (England and Wales) [13]. It provides guidance on communicating with and about diverse court users, and on avoiding assumptions that can distort credibility assessments or lead to exclusion. While not law, such guidance helps to give content to fairness obligations. In this practical sense,

empathy is embedded in the everyday judicial craft of ensuring that parties can participate and are treated with respect.

4.4 The boundary condition: impartiality and bias

Any account of empathy must confront the legal requirement of impartiality. Legal systems demand not only actual impartiality but also the appearance of impartiality. Empathy becomes legally problematic when it translates into asymmetrical engagement that signals prejudgment or unequal valuation of parties. Accordingly, the most defensible conception of judicial empathy is symmetrical: perspective-taking and respectful engagement extended to all parties, combined with awareness of how structural disadvantage or vulnerability may affect participation. Empathy in this sense does not dictate outcomes; it shapes the process by which outcomes are reached and justified.

5 Machine involvement in judging: current practice and legal constraints

5.1 A spectrum of AI involvement

Discussions about “AI judges” often obscure the range of possible system roles. For analytical purposes, AI involvement can be placed on a spectrum:

1. Administrative support (e.g., transcription, document management, search and retrieval, translation).
2. Decision support (e.g., summarising evidence, identifying issues, suggesting comparable cases or ranges).
3. Drafting assistance (e.g., generating a first draft of reasons or plain-language explanations for human editing).
4. Outcome recommendation (e.g., generating a proposed disposition or sentence recommendation).
5. Automated decision-making (e.g., system issues a binding judgment with no meaningful human involvement).

The legal and ethical risk increases sharply as tools move from support and drafting towards outcome recommendation and automation. The roundtable summary endorses this spectrum-based thinking, emphasising that written-only, procedural matters may be more amenable to carefully bounded AI assistance than credibility-dependent, substantive determinations Midgley and Sinclair [15].

5.2 Judicial guidance on generative AI

Judicial systems have begun to issue guidance that recognises practical use while highlighting risks. For example, guidance for judicial office holders in England and Wales has emphasised that judges remain personally responsible for any work produced with AI assistance, that outputs may be inaccurate or biased, and that confidential information should not be entered into public AI tools without appropriate safeguards Courts and Tribunals Judiciary (England and Wales) [5].

Such guidance is not merely prudential; it reflects legal constraints on courts. If AI outputs are unreliable, unreviewable, or undisclosed, they can undermine reason-giving, the parties’ ability to challenge, and open justice.

5.3 Case study: *VP Evans & Ors v HMRC* (2025)

A significant recent illustration is *VP Evans & Ors v The Commissioners for HMRC* UK First-tier Tribunal (Tax Chamber) [18]. In a postscript, the judge disclosed the use of AI in producing the decision, described the tool as a “discrete case management application” and explained constraints on its use (for example, no use for legal research; verification of outputs; and suitability limited to decisions “on the papers” rather than credibility-based hearings). The postscript frames transparency and bounded use as essential to maintaining confidence.

This decision is important less because it resolves a doctrinal question than because it models a disclosure-based norm: if AI is used in a way that could affect reasons, parties should know, and a judge should remain accountable for verification. It also suggests a practical dividing line between contexts: when the judicial task turns on credibility, demeanour or the evaluation of contested oral evidence, the stakes for empathy and human judgment rise.

5.4 Data protection and solely automated decisions

Data protection law provides an additional constraint on automation. Article 22 of the UK GDPR recognises a right not to be subject to a decision “based solely on automated processing” that produces legal effects or similarly significant effects, subject to limited exceptions and safeguards European Union [9]. Although courts acting in a judicial capacity occupy a distinct constitutional position and may benefit from specific exemptions or oversight arrangements, the underlying principle resonates for adjudication: where decisions have serious legal consequences, meaningful human involvement and contestability matter.

5.5 Emerging AI-specific regulation: EU AI Act and beyond

The EU Artificial Intelligence Act treats certain AI uses in the administration of justice and democratic processes as high-risk, including AI systems intended

to assist a judicial authority in researching and interpreting facts and law or applying law to facts European Union [10]. This reflects the recognition that such systems can significantly affect the right to an effective remedy and a fair trial. Parallel developments exist internationally, including the Council of Europe’s Framework Convention on AI and human rights, democracy and the rule of law, which emphasises transparency, accountability and remedies Council of Europe [4].

5.6 Comparative cautionary tales: algorithmic decision support

While fully automated judging is rare, decision support systems have already been contested. In the United States, the use of proprietary risk assessment tools in sentencing (including COMPAS) has raised due process concerns about opacity, contestability and disparate impact; *State v Loomis* is widely discussed in this regard Wisconsin Supreme Court [20]. In Europe, judicial and administrative cases challenging algorithmic systems (for example, the Dutch SyRI judgment) illustrate how opacity and insufficient safeguards can render algorithmic governance incompatible with human rights standards District Court of The Hague [6]. These experiences underscore that “assistive” systems can still shape outcomes and must therefore be governed as if they matter.

6 Can AI be trained for judicial empathy? Capabilities, limits and risks

6.1 Apparent empathy and the simulation problem

Current generative AI systems can produce language that readers rate as empathetic, even in sensitive contexts. A widely cited example compares large language model responses with clinician responses and finds that lay readers often rate AI outputs as more empathetic Ayers et al. [1]. This creates a simulation problem: apparent empathy can be generated without human understanding or moral responsibility. In judicial settings, the risk is that surface-level empathy masks gaps in reasoning, accountability, or substantive fairness.

6.2 Bias, representativeness and over-reliance

AI systems inherit the biases of their training data and can reproduce existing inequalities, sometimes at scale. Critics of algorithmic governance warn that predictive systems can entrench disadvantage under a veneer of objectivity Eubanks [7]. Over-reliance on automated outputs can also create “automation bias” where human decision-makers defer to system recommendations, even when they are flawed Citron [3]. In judicial contexts, this could undermine impartiality or lead to unexamined disparities.

6.3 Reason-giving, explainability and contestability

The legal system demands reasons that can be scrutinised and contested. Black-box systems challenge this expectation. The literature on accountable algorithms emphasises that transparency and contestability are not optional add-ons but essential safeguards for legitimacy Kroll et al. [14]. For judicial AI, this means that any assistance must be explainable to both judges and parties, and that outputs must be verifiable and challengeable.

7 Governance and assurance: from black-box concerns to auditability

Transparency, contestability and accountability are central to lawful AI use. Emerging governance instruments, including ISO/IEC 42001 and ISO/IEC 23894, emphasise management systems, risk controls, and auditability for AI deployments International Organization for Standardization [12, 11]. Similarly, the NIST AI Risk Management Framework provides structured approaches to mapping, measuring, and managing AI risk National Institute of Standards and Technology [16].

For courts, these frameworks suggest a shift from ad hoc experimentation to structured assurance: documented use cases, risk assessments, human oversight protocols, and records of AI inputs and outputs. These controls must be aligned with open justice and the duty to give reasons, ensuring that parties can understand how AI-assisted material was used (if at all) in the decision-making process.

8 A practical framework for empathy-supporting judicial AI

A workable framework for empathy-supporting judicial AI should be anchored in four principles.

1. **Human responsibility:** A judge remains accountable for the decision and must verify all AI-assisted content.
2. **Procedural transparency:** Parties should know when AI has been used in material aspects of drafting or reasoning, and should have the opportunity to respond.
3. **Contestability and auditability:** AI outputs should be logged and auditable, with clear records of data sources and reasoning pathways.
4. **Proportionality of use:** The level of AI involvement should match the stakes and the nature of the task, with higher-risk uses subject to stricter controls.

Operationally, this framework requires courts to define permissible use cases, deploy AI systems within secure environments, and provide training for judges and staff. It also requires a governance structure that can adapt to evolving regulatory frameworks (including the EU AI Act and UK data protection law) and standards-based assurance regimes European Union [10, 9].

9 Conclusion

Empathy in judging is best understood not as a licence to decide by feeling, but as the disciplined ability to understand and communicate with those subject to legal authority. In that procedural sense, empathy is interwoven with legal requirements: fair hearing, equal treatment, and reasoned decisions.

AI systems can, in limited and carefully governed ways, support these aims by improving comprehension, consistency and accessibility. Yet the legitimacy of adjudication depends on human accountability and on a transparent record of reasoning. Tools that become outcome-determinative, operate opaquely, or distribute empathy unevenly risk undermining both legality and trust. The task for legal technology is therefore not to build an “AI judge” but to design and govern empathy-supporting systems that enhance procedural justice while preserving the constitutional function of courts.

References

- [1] John W. Ayers et al. Comparing physician and artificial intelligence chatbot responses to patient questions posted to a public social media forum. *JAMA Internal Medicine*, 2023.
- [2] Susan A. Bandes. Empathetic judging and the rule of law. *Cardozo Law Review de novo*, page 133, 2009.
- [3] Danielle K. Citron. Technological due process. *Washington University Law Review*, 85:1249, 2008.
- [4] Council of Europe. Framework convention on artificial intelligence and human rights, democracy and the rule of law, 2024. Opened for signature 5 September 2024.
- [5] Courts and Tribunals Judiciary (England and Wales). Artificial intelligence (ai) – judicial guidance, 2025. First published December 2023; updated October 2025.
- [6] District Court of The Hague. District court of the hague, ecli:nl:rb-dha:2020:1878 (syri), 2020. 5 February 2020.
- [7] Virginia Eubanks. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. St Martin’s Press, 2018.

- [8] European Court of Human Rights. Ruiz torija v spain, app no 18390/91, 1994. European Court of Human Rights, 9 December 1994.
- [9] European Union. Regulation (eu) 2016/679 (general data protection regulation), as retained and amended in uk law (uk gdpr), article 22, 2016.
- [10] European Union. Regulation (eu) 2024/1689 laying down harmonised rules on artificial intelligence (artificial intelligence act), 2024. OJ L, 12 July 2024.
- [11] International Organization for Standardization. Iso/iec 23894:2023 information technology — artificial intelligence — guidance on risk management, 2023.
- [12] International Organization for Standardization. Iso/iec 42001:2023 information technology — artificial intelligence — management system, 2023.
- [13] Judicial College (England and Wales). Equal treatment bench book, 2025. Current online edition.
- [14] Joshua A. Kroll et al. Accountable algorithms. *University of Pennsylvania Law Review*, 165:633, 2017.
- [15] Alan Midgley and Stacy Sinclair. Roundtable: Law, ethics and ai: Can, or should, ai be trained for judicial empathy?, 2025. International Dispute Resolution Centre, London, 17 July 2025.
- [16] National Institute of Standards and Technology. Ai risk management framework (ai rmf 1.0), 2023. January 2023.
- [17] Tom R. Tyler. Procedural justice and the courts. *Court Review*, 2007.
- [18] UK First-tier Tribunal (Tax Chamber). Vp evans & ors v the commissioners for hmrc [2025] ukftt 1112 (tc), 2025. First-tier Tribunal (Tax Chamber), decision released 14 October 2025.
- [19] Robin West. The anti-empathic turn, 2017. Selected work on empathy in adjudication.
- [20] Wisconsin Supreme Court. State v loomis, 881 n.w.2d 749, 2016. Wisconsin Supreme Court.

A Workshop

Roundtable: Can, or should, AI be trained for judicial empathy?

17 July 2025, International Dispute Resolution Centre, London, UK

This Roundtable, held on 17 July 2025 at the IDRC, was curated pro bono for the public good by Alan Midgley with the support of the IDRC and all the speakers listed below.

Speakers and attendees

- *Chair:* Alan Midgley, AG Midgley Ltd • *Host:* Damian Hickman, IDRC
- The Rt. Hon. Sir Geoffrey Vos, MR • Andrew Churchill, Technology Strategy
- The Rt. Hon. Lord Neuberger of Abbotsbury • Ben Haslam, Superset Technologies
- Professor Michael Mainelli, Z/Yen Group Limited • Professor Christopher Clack, UCL
- *Nick Davies, iC4DTi* • Dante Small
- *Dr. Stacy Sinclair, Fenwick Elliott* • Julia Mackenzie, University of Cambridge
- *Karl Reimer, University of Zurich* • Flavia Kenyon, 33 Chancery Lane
- Stephen Dowling SC, TrialView • Merlin Lindsay, Velocity
- Sophia Adams Bhatti, JUSTICE UK • Dr Peter Zheng, Robosloth
- Dr Alastair Moore, UCL Technologies • Dominic Courage
- Dr Geoffrey Goodell, UCL • Richard Bayfield (Part), ICE
- Rasto Koval, Cequence • Tom Oakden, Bluelight
- Ryland Ash, Watson Farley Williams • Dr Alex Mankoo, The British Academy
- Ian Grigg, Peer For Peer Foundation • Will Roberts, Weber Shandwick
- Anonymous

The notes below do not represent the views of any individual, nor those of the employers or companies of the attendees and have been purposely anonymised.

With special thanks to Dr Stacy Sinclair for meticulously summarising and anonymising the speaker's notes, and to Karl Reimer for his help in preparing the final document, as well as the drafting committee (noted above*).

Foreword by Sir Geoffrey Vos, Master of the Rolls and Head of Civil Justice in England & Wales

Artificial intelligence is no longer a distant prospect but an everyday reality within the justice system. Judges and lawyers, like other professionals, are already using AI tools to streamline routine work and enhance their legal reasoning. The central question is how they can do so in a way that preserves public trust and the integrity of the judicial process. The rapid evolution of machine intelligence requires us to reflect on the proper ethical boundaries of its use in the legal and judicial sphere. There may come a point when certain tasks can be performed by AI with a speed and consistency that surpass human capability, but justice is not merely a technical exercise in correctness. It is rooted in empathy, moral judgment, and the interpretative understanding of human circumstances. As we navigate this new terrain, it is essential to identify which aspects of legal work in general, and of judicial decision-making in particular, must remain inherently human if the rule of law is to retain its legitimacy and moral force. When deployed responsibly, AI can enhance efficiency, improve access to justice, and assist in the accurate and timely handling of information. Yet its integration must always be accompanied by rigorous verification, transparency, and adherence to the ethical standards that underpin legal work. This roundtable sought to inform that continuing debate and to consider how technological progress can be harnessed to strengthen, rather than redefine, the principles of justice.

Executive summary

1. This paper is a summary of a roundtable on AI and empathy.
2. The Roundtable concluded that while AI cannot yet feel empathy, it may increasingly be trained to recognise and respond to human context, thereby supporting fairer, more accessible judicial reasoning. AI should assist, not decide, ensuring that moral understanding, compassion, and discretion remain distinct human responsibilities. Participants agreed that any AI system designed to support judicial empathy must be transparent, auditable, and compliant with recognised governance and data standards, to avoid opaque or “black box” decision-making. Accountability, explainability, and public trust were identified as essential preconditions for legitimate use of AI in the justice system.
3. The Roundtable took empathy as its central theme, recognising it as a defining human quality in judicial decision-making and a benchmark against which the limits of AI must be assessed. The discussion confirmed that while AI may simulate certain empathic behaviours, it cannot replicate the wisdom, discretion, and contextual sensitivity that underpin legitimate legal judgment.
4. Participants also agreed that empathy cannot be considered in isolation.

For AI to play a meaningful and legitimate role in legal services, issues of trust, accountability, and institutional legitimacy must be addressed. Discussion extended to practical mechanisms for transparency and auditability, including blockchain-based record-keeping, Ricardian contracts, and the debated concept of legal personhood for AI agents – framed not as a concession of status, but as a possible mechanism for securing responsibility and accountability.

5. Shortly after the Roundtable, the decision in *VP Evans & Ors v The Commissioners for HMRC* offered a timely example of judicial engagement with AI tools. Particular attention is drawn to paragraph 43 of the judgment, which sets out the precise context in which AI was used: a discrete case-management application, determined solely on papers, written submissions, without a hearing or any assessment of credibility or honesty.
6. This instance reflects one type of judicial application of AI envisioned in the Roundtable: supporting the process of decision-making, as opposed to substantive determinations. It also reinforces the distinction between written (document-based) and oral (credibility-assessing) forms of advocacy.
7. What emerges is a spectrum of AI use cases. At one end are applications that support efficiency – such as document review, extracting data from evidence, and generating summaries of written submissions or case papers for procedural decisions. These tools can improve judicial capacity without compromising fairness or independence. At the other, more contested end, are applications that seek to influence or determine substantive legal judgments. Panellists agreed that this domain must remain the preserve of uniquely human qualities: empathy, judgment, identity, personhood, and respect for human rights. These are not incidental virtues but essential safeguards.
8. Accordingly, this paper provides both a synthesis of participants’ views on empathy and its limits and an account of the wider governance structures necessary to ensure that AI adoption in legal services strengthens, rather than undermines, public confidence in justice.

Roundtable highlights

Opening session. The Roundtable began with reflections on the evolving balance between machine and human roles in the justice system. English law was highlighted as a system of global influence, underpinning a significant portion of international trade and providing trusted mechanisms for dispute resolution. Even in highly codified domains such as tax law, where statutes strictly define liabilities, human judgment remains essential. The opening remarks stressed that oversight, empathy, and discretion must remain central to legal processes as technology is adopted. Recent developments, such as the Electronic Trade

Documents Act 2023, were cited as evidence of the UK’s leadership in adapting legal frameworks to digital change. Smart contracts and Ricardian contracts were noted as future tools for automating transactions, but with the reminder that the integrity of the common law tradition depends on careful and balanced integration of such technologies.

Panel 1: Judicial empathy and AI. The first panel examined whether AI can or should be trained to exhibit judicial empathy. Speakers noted that empathy in justice involves qualities such as neutrality, respectful treatment, and giving parties the sense that they are being heard. While AI cannot replicate wisdom or moral authority, it may be able to simulate some empathic behaviours. Studies were cited where AI chatbots were perceived as more empathetic than physicians, suggesting cognitive empathy – machines detecting and classifying emotions – may be on the horizon. Even so, panellists agreed that genuine empathy is beyond the reach of current AI systems. AI’s strengths lie in its efficiency, particularly in handling legal research, disclosure, and case management. In specific areas, such as construction adjudications that concern specific issues of payment, AI could provide rapid, objective first-tier rulings, with human judges reviewing outcomes. Trust, rather than empathy, was seen as the decisive factor for adoption. AI systems must demonstrate reliability, transparency, and accountability if they are to complement, not replace, human judgment.

Panel 2: Public policy, ethics and trust. The second panel examined the societal and legal implications of AI-driven decisions. Panellists highlighted the importance of procedural fairness, noting that legitimacy depends not only on efficiency but also on the fairness of the process. Trust in systems – whether human or machine – is built through transparency and accountability. Concerns were raised about over-reliance on AI, especially as models like LLMs evolve from advisory tools to potentially authoritative systems. Their “black box” nature poses challenges to trust, even as their performance rivals that of top-tier human expertise. Human rights and the rule of law were identified as foundational principles, not just codified law, but deeply held societal values.

Panel 3: AI, FinTech and accountability. The third panel explored the intersection of AI, law, and financial technologies. Panellists discussed the tension between the efficiency of automation and the need for fairness and human judgment. Semantic technologies and ontologies – such as JSON-LD and the Web Ontology Language – were highlighted as tools for structuring machine- and human-readable contracts. Ricardian contracts were presented as a model for binding agreements bridging law and code.

Cross-cutting themes. Trust emerged as the central requirement for the acceptance of AI. Systems must be transparent, reliable, and demonstrably effective. Blockchain-based audit trails and other accountability mechanisms were viewed as crucial to maintaining transparency and openness in decision-making. AI must be designed and deployed in line with ethical principles and legal norms. Standards-based governance frameworks, drawing on international initiatives and ISO standards, were emphasised as essential to embedding accountability. Speakers stressed that empathy, wisdom, and discretion are irreplaceable in

judicial contexts. While AI may simulate empathy or enhance efficiency, human oversight remains essential for decisions that require contextual and moral judgment.

Practical recommendations. AI should be introduced gradually, starting with low-risk applications in international trade, supply chains, and technical legal tasks. Attention should be paid to the hierarchy of oversight and to ensuring that human rights remain central to AI design and deployment. Trust and transparency are essential; systems must earn trust by demonstrating reliability and transparency. Blockchain and audit tools should provide tamper-evident trails of AI decisions. Practical frameworks, such as ISO 42001, should guide deployment, ensuring that technological innovation is aligned with principles of justice and authority. Collaboration across regulators, banks, trading companies, and professional bodies will be vital in designing governance frameworks.

Conclusion. The Roundtable on Law, Ethics and AI: Can or Should AI Be Trained for Judicial Empathy? provided a profound exploration of AI's role in legal and institutional frameworks. While AI holds transformative potential to improve efficiency and consistency, it cannot replace the human values of empathy, wisdom, and discretion. Trust and accountability, grounded in transparency and robust governance, are essential for public acceptance. The discussions called for a cautious and collaborative approach: beginning in low-risk areas, embedding ethical standards, and preparing institutions for the societal impacts of AI. Post-meeting reflection: the decision in *VP Evans & Ors v The Commissioners for HMRC* affirms the cautious integration pathway supported by the Roundtable. It provides a tangible illustration of AI use under judicial supervision, confirming that transparency and limited, well-defined applications are already being trialled within the legal system. Notably, the distinction between applications suitable for AI assistance (e.g. written-only, procedural matters) and those requiring human judgment (e.g. oral hearings, assessments of credibility) has emerged as a guiding principle for future adoption.