ODI Global

Working paper

AI and the potential for universal affordable access to front-line justice

Marcus Manuel, Clare Manuel and Yoko Yoshihara

October 2026



Readers are encouraged to reproduce material for their own publications, as long as they are not being sold commercially. ODI Global requests due acknowledgement and a copy of the publication.

For online use, we ask readers to link to the original resource on the ODI Global website.

The views presented in this paper are those of the author(s) and do not necessarily represent the views of ODI Global or our partners.

© ODI Global 2026

ODI Global
4 Millbank
London SW1P 3JA
United Kingdom
www.odi.org

This work is licensed under CC BY-NC-ND 4.0.

How to cite: Manuel, M., Manuel, C. and Yoshihara, Y. (2026) *AI and the potential for universal affordable access to front-line justice.* ODI Global Working Paper. London: ODI Global (https://odi.org/en/publications/ai-and-the-potential-for-universal-affordable-access-to-front-line-justice/)

Acknowledgements

The authors thank the Charles Stewart Mott Foundation for funding this research.

The authors are grateful to peer reviewers Nate Edwards (Center on International Cooperation, New York University) and Ameya Naik (ODI Global) for their comments and insights, which have enriched the paper, alongside quality assurance from Frederique Dahan, copyediting from Matthew Foley and support from Asiyah Dubuisson.

The authors are grateful to the many people who gave their time to be interviewed or to correspond with us, and who generously shared data, evidence and insights.

We particularly thank the team at the British Columbia People's Law School (Canada): Patricia Byrne, Drew Jackson and Urszula Lipska, who gave so generously of their time to establish a costing of their public-facing chatbot *Beagle+*. We thank our ODI Global colleague Divij Joshi for his contribution on risk and accountability.

We are grateful to those who gave us extensive comments and shared their data, including Jason Tashea (National Center for State Courts, US) and Viviana Brun and Zainab Malik (Hague Institute for Innovation of Law (HiiL)).

We also thank those we interviewed who are delivering and designing other front-line justice services and justice technology: Akshetha M. Ashok (Sama, India); Dan Barrett (Citizens Advice, UK); Deepika Kinhal (CORD, India); Patricia Peace Maaike (LawBot Africa, Uganda); Clifford Msiska and Chimwemwe Ndalahoma (Paralegal Advisory Service Institute (PASI), Malawi); Rich Simon (Luminate Solutions, US) and Sarah Stephens (University of Sussex).

We also thank colleagues in international organisations for their insights and suggestions of best examples: Henk-Jan Brinkman (International Development Law Organization); Gloria Manzotti, Victoria Mello and Napoleon Enayaba (UN Development Programme (UNDP)); Philip John McGrath (World Bank) and Mariane Piccinin Barbieri (Organisation for Economic Co-operation and Development (OECD)).

We also thank other academics, independent experts and commentators for their inputs and advice: Varsha Mahadeva Aithala (National Law School of India University); Daniela Barba (World Justice Project); Lisa Burton (Authentic Legal Ai); Kate Fazio (Justice Innovation Consultant, Australia); Maya Markovich (Justice Technology Association, US); Keith Porcaro (Duke University School of Law, US); James Teufel (Institute for the Advancement of the

American Legal System); Juliet Uraz (PhD candidate at LSE) and Jan van Zyl Smit (Bingham Centre for the Rule of Law).

The opinions expressed in this paper, and any outstanding errors and omissions, are the authors' alone.

About this publication

This paper is part of ODI Global's work on justice finance (https://odi.org/en/about/our-work/justice-finance/) which explores how smarter approaches to justice financing could better address people's everyday justice problems. Our work is inspired in part by lessons on public financial management from other sectors.

About the authors

Clare Manuel is a lawyer and an ODI Global Senior Research Associate. She was previously a member of the UK government's Legal Service and then founding director of The Law & Development Partnership, providing policy advice on access to justice to development agencies and governments in sub-Saharan Africa, Asia, the Caribbean and the South Pacific. She is also the former chair of a free legal advice clinic in North London (where she continues to advise) and is a founding director of a digital dispute resolution start-up in the UK.

Marcus Manuel is an economist by training and an ODI Global Senior Research Associate. His previous roles include Development Policy Adviser to the UK Development Minister, a regional director for the Department for International Development (DfID), a deputy director at the UK Treasury and Senior Economic Adviser in the Ugandan Ministry of Finance and Planning.

Yoko Nishiki Yoshihara is a Representative at the Japan International Cooperation Agency (JICA)'s Myanmar Office. She contributed to this report as an intern at ODI Global during her Masters of Public Policy at the University of Oxford's Blavatnik School of Government. She was previously Deputy Director of JICA's Law and Justice Team, Governance Group.

Contents

Acknowledgements 3
Executive summary 8
Key findings 8
Recommendations 10
1 Introduction 12
1.1 Methodology 13
2 Already on the front line: AI tools in use today 15
2.1 Introduction 15
2.2 The context: AI developments in law firms and judiciaries 15
2.3 Overview of ODI's global scan 16
2.4 Legal information, advice and assistance 16
2.5 AI in informal dispute resolution 18
2.6 Global reach, but limited scale 19
2.7 Where most people go: general-purpose LLMs 20
2.8 Conclusion 21
3 Risks and evidence gaps for AI at the front line 25
3.1 Introduction 25
3.2 The risks of general-purpose LLMs 25
3.3 The risk of AI increasing inequality at the front line 30
3.4 General AI risks and mitigations 31
3.5 Evidence has not kept pace with AI developments 33
3.6 Conclusion 34
4 Closing the justice gap: scaling up AI at the front line 36
4.1 Introduction 36
4.2 Affordability and value for money 36
4.3 Addressing the barriers 41
4.4 Conclusion 45
5 Conclusions and recommendations 47
5.1 Conclusions 47
5.2 Recommendations 48
References 51
Appendix 1 List of examples of AI tools 64
Appendix 2 Characteristics of a well-designed public-facing bespoke legal chatbot 69

Display items

Box 1 Beagle+ ……22
Table 1 Stylised cost effectiveness model for bespoke AI chatbot…….40

Abbreviations and acronyms

AI	Artificial intelligence
API	Application programming interface
AVA	Alaska Virtual Assistant
CBA	Cost–benefit analysis
CEA	Cost-effectiveness analysis
CEPEJ	European Commission for the Efficiency of Justice
CRT	Civil Resolution Tribunal (British Columbia)
DCAF	Geneva Centre for Security Sector Governance
EU	European Union
GDPR	General Data Protection Regulation
HAI	Human-Centered Artificial Intelligence (Stanford University)
HiiL	Hague Institute for Innovation of Law
IDLO	International Development Law Organization
JAC	Justice Action Coalition
LLM	Large language model
LSB	Legal Services Board (England and Wales)
MSME	Micro, small and medium-sized enterprise
NITA	National Institute for Trial Advocacy
NYSBA	New York State Bar Association
ODR	Online dispute resolution
OECD	Organisation for Economic Co-operation and Development
RAG	Retrieval-augmented generation
RCT	Randomised controlled trial
SMS	Short message service
SRA	Solicitors Regulation Authority (England and Wales)
UNDP	United Nations Development Programme
UNESCO	United Nations Educational, Scientific and Cultural Organization
WJP	World Justice Project

Executive summary

Most attention on artificial intelligence (AI) in justice has focused on improving the efficiency of law firms and judiciaries. Both are investing heavily in AI. This paper looks instead at the justice front line: legal information, advice, assistance and informal dispute resolution for the 1.5 billion people worldwide who cannot resolve their justice problems. It draws on an original ODI Global scan of existing bespoke front-line AI tools worldwide, a literature review, interviews and the first-ever detailed costing of a leading public-facing legal chatbot.

Key findings

Bespoke AI tools are starting to be used on the front line
ODI Global's scan found 77 bespoke front-line AI tools in 39 countries. Some are public-facing legal chatbots such as Canada's *Beagle+* in British Columbia*,* or South Africa's *Maya.* Other tools provide advice through a voice service on basic mobile phones (Niger) or enable multilingual access to services (India). Some support providers of front-line services, such as the Citizens Advice *Caddy* tool in England and Wales.

Few public-facing AI tools have reached scale in terms of outreach
Even the most widely used legal chatbot, *Maya*, reaches less than 2% of South African adults, and *Beagle+* only 0.2% of adults in British Columbia. In informal dispute resolution, AI is mainly helping repeat players such as bulk debt collectors, not individuals.

General-purpose LLMs are much more widely used
Most US adults (65%) now use general-purpose large language models (LLMs) for legal help. The most senior judge in England and Wales has recognised general-purpose LLMs to be good enough to submit a claim to court. The question is no longer whether AI will be used by the public to access justice, but how to make it safer and fairer.

General-purpose LLMs have inherent risks
General-purpose LLMs are largely accurate when asked well-phrased questions, but they can give wrong advice, especially when detailed knowledge of local laws is critical. Their answers can be inconsistent and not appropriate to the needs of the user. They have

inherent risks of confidentiality, sycophancy and bias. While some aspects of accuracy are improving, better models are unlikely to solve these wider problems.

Well-designed bespoke AI chatbots are now demonstrably a better option, though not a perfect one

Well-designed, retrieval-augmented generation (RAG) chatbots draw on the capability of general-purpose LLMs but are then limited to working with curated local legal material. Run by independent accountable organisations, their results are tested and overseen by lawyers. *Beagle+* provides legal information, with approximately 99% of its reviewed responses rated 'safe' in relation to legal accuracy and appropriateness. *Maya*, operating in a more permissive regulatory regime, shows it is possible to go further and provide advice and assistance, all the way through to preparing court documents and statements to be read out in court. The use of general-purpose LLMs means these tools share some of their weaknesses. But these two examples provide proof of concept: it is possible to implement a well-designed bespoke AI chatbot that is superior to a general-purpose LLM.

Bespoke AI chatbots are remarkably affordable

The key new finding in this paper is how cheap such bespoke tools are. *Beagle+* cost $110,000 to develop and $58,000 a year to run. Its unit costs of $7 per conversation are projected to fall to $1 at scale, many times cheaper than a telephone information service. These costs include keeping a 'human on the loop' to monitor answers.

Bespoke AI chatbots now look like a low-cost, low-risk, high-return investment

A national scale-up of a bespoke public-facing legal chatbot could transform access to justice for millions. For those needing face-to-face support, it would enable human front-line advisers to work more efficiently, reaching more people for the same budget. A scaled-up bespoke legal chatbot in British Columbia would cost less than 0.4% of the province's legal aid budget. Chatbots for all 41 lower-income countries in sub-Saharan Africa could be developed for $5 million in total.

There are few examples of AI transforming informal dispute resolution – but without this, the justice gap will remain

Online dispute resolution has worked at scale in the private sector (e.g. eBay) for over 15 years, but public-sector experience is mixed, and our global scan found few AI examples. How far AI should go in deciding disputes is a democratic question around what the public would accept. But the public use of general-purpose AI tools is likely to increase the demands on courts, and the resulting potential surge in claims suggests that testing these approaches is urgent.

Without deliberate action, AI will widen the justice gap

Currently, AI in justice mainly benefits those who can pay for lawyers and repeat players such as bulk debt collectors. Public-facing

bespoke tools operating at any significant scale are rare, especially in lower-income countries, and standards for front-line AI are almost entirely missing. This risks reinforcing a two-tier system in which only the resourced and capable can enforce their rights.

The barriers to scaling up AI at the front line could be overcome
There is scope for AI to have a significant impact on closing the justice gap. The technology exists, the curated legal datasets are present in many countries, and the costs are small relative to national budgets. Overcoming the constraints on AI-assisted front-line justice will require new thinking and making choices about finance, regulation and professional culture.

Money: Front-line justice remains at the margins of justice budgets, and small non-profit providers lack capital.

Monopolies: Rules on the practice of law can restrict the operations of non-profit legal advice centres, their non-lawyer advisers and their chatbots, while general-purpose LLMs operate largely unconstrained.

Mindsets: Any hesitancy on the part of the front-line legal advice sector about AI, often for principled reasons, needs to be engaged with.

While the evidence base is good enough to justify investing in front-line AI, there are significant evidence gaps
Investments in front-line AI need to be designed as structured experimentation, with results and costs published. Only 22 of the 77 tools in the global scan publish usage data, and only 12 publish accuracy or evaluation data. Independent audits, randomised controlled trials and cost-effectiveness analyses are rare.

Recommendations

Realising AI's potential at the justice front line will require a collective effort to tackle the barriers of money, monopolies and mindsets, and to build evidence on what works.

Governments and regulators should:

- Facilitate sharing of local knowledge, often held by non-profits, to enable the creation of a national service that can draw on the necessary detailed understanding of local laws and regulations.
- Fund the development of one AI chatbot in a non-profit advice organisation (a trusted body independent of government).
- Allow non-profit organisations to give the full range of advice and assistance for everyday justice problems – delivered by their chatbots and non-lawyer advisers.

- Track the quality of free public legal advice from existing providers, general-purpose and bespoke AI tools through a continuous benchmarking process.
- Invest in research and up-to-date legal needs surveys that capture AI use.
- Accelerate exploration of innovative use of AI in pre-court dispute resolution mechanisms (and initiate consultations on public acceptability).

Philanthropic funders and donors should:
In lower-income countries:

- Fund the development of curated legal knowledge bases.
- Fund the up-front costs of bespoke chatbots including funding $5 million for their development in sub-Saharan Africa.

All funders should:

- Fund research to identify the key features of an effective chatbot and the most cost-effective way to deliver them.
- Make it a condition of their funding that organisations undertake and publish evaluations, including unit costs, accuracy and user outcomes.

1 Introduction

Much of the focus on AI in the justice sector has been on the formal system – on AI's ability to transform how lawyers, law firms and judiciaries work. The World Bank's published examples of AI use in the sector all relate to formal justice (World Bank, 2026a; World Bank, 2026c). Of the 19 justice-related cases cited in OECD's review of AI use cases in government, 17 concern the judiciary (OECD, 2025b: 265ff).

The focus of this paper is on AI at the justice frontline: its potential to transform the provision of legal information, advice, assistance and informal dispute resolution services for people currently without access to justice.[1] The paper looks at how AI could contribute to closing the justice gap for the 1.5 billion people worldwide who cannot resolve their justice problems (WJP, 2019; 2023).

Commentators and researchers on AI and access to justice have mixed views on its potential. Some are optimistic. The World Bank's 2026 World Development Report on *The Promise of Artificial Intelligence* presents justice as one of four citizen-welfare service areas[2] offering a speedy path to greater prosperity (World Bank, 2026a and b). Richard Susskind, a former Special Envoy for Justice and AI to the Commonwealth Secretary-General, regards AI as a gamechanger in service delivery and access to justice (Susskind, 2023; 2026). The public also seems optimistic. Research in the UK suggests there is strong consumer support for legal AI: around three-quarters expect it to make legal services easier to use, more affordable and more accessible (MEL, 2026).

Others are more cautious. The United Nations has noted that digital tools, including AI, have the potential to improve access to justice, but can also deepen inequality and disadvantage marginalised groups (UN Secretary-General, 2024). A 2026 review of AI and justice for the Nuffield Foundation's Public Right to Justice Programme concludes: 'whether [AI] will advance or undermine [the] public's right to justice [that is accessible, fair and effective] is, at present, an open question' (Sargeant, 2026: 5). Others caution about a lack of evidence that AI will improve access to justice (Tashea,

[1] The Justice Action Coalition's (JAC) Justice Financing Framework describes front-line justice services as comprising legal information, advice, assistance and informal dispute resolution services for everyday justice problems (JAC, 2025).
[2] Alongside education, health care and social protection.

2026). Recent work on public perceptions of AI highlights growing anxiety and lack of trust (Stanford HAI, 2026).

This working paper contributes to the debate on AI's potential to make a significant contribution to front-line justice and to closing the justice gap. It analyses AI tools already in use on the front line (Section 2); considers challenges, risks and evidence gaps (Section 3); and reflects on what AI could achieve at scale, and the barriers that would need to be overcome to get there (Section 4). The paper concludes with recommendations on the way forward (Section 5). While the paper takes a global perspective, the authors have also drawn on their experience and location in the UK to illustrate latest developments.

1.1 Methodology

The foundation of the paper is an original global scan of bespoke AI-supported tools in front-line justice. This is supplemented with a literature review.

Front-line justice
Front-line justice refers both to the nature of the provider (accessible and low-cost/affordable) and the type of service provided:

- Legal information, advice and assistance – providing tailored individual information and advice on everyday justice problems, and assistance to resolve them.
- Informal dispute resolution – dispute resolution mechanisms outside the formal system, usually easily accessible for less serious/low-value problems.

The Justice Action Coalition's (JAC) Justice Financing Framework provides a more detailed description of front-line justice services (JAC, 2025).

Global AI scan
The global scan covers both bespoke AI tools providing front-line justice services and simple uses of AI, such as translation and transcription, to enable access to these services. The aim is to identify the extent to which these tools and uses exist, what they do and whether any have gone to scale.

The scan takes a broad view of AI. It includes bespoke generative AI providing information, advice and assistance, drawing on a curated body of trusted information (legal databases, legislation and case law). It also includes predictive AI: machine learning systems that forecast, classify and predict patterns (such as the likelihood of a claim succeeding). Also considered is agentic AI capable of executing complex, multi-step legal tasks autonomously. Rajendran et al. (2025) provide a helpful survey of AI's roles in legal services: automating routine legal tasks, improving the accuracy and speed of

legal research, document review and providing real-time assistance through virtual assistants and chatbots.

ODI Global's scan, while new and extensive, is not statistically sampled, and can only be illustrative rather than comprehensive, particularly given the pace of change. Sources for the scan (and the literature review) were:

- ODI Global's networks (including the Hague Institute for Innovation of Law (HiiL), the International Development Law Organization (IDLO), the OECD, the United Nations Development Programme and the World Bank.
- Interviews with key informants (listed in the acknowledgements).
- Online open sources with information searched for, verified and in some cases translated with the help of AI tools including Claude, ChatGPT and NotebookLM.

While the boundary is not precise (and use of AI keeps changing), the scan does not cover non-AI-based, online/technology-based mechanisms. Examples cited generally do not include pilots, tools that have not been deployed, or tools that have been withdrawn.

2 Already on the front line: AI tools in use today

2.1 Introduction

Formal justice providers – law firms and the judiciary – are adopting AI at pace. This section reviews the extent to which AI is also available at the front line to assist with people's everyday justice problems. ODI Global's scan of bespoke AI tools in front-line justice produced a list of 77 examples (see Appendix 1, with a fuller description of each based on public sources in a separate volume).

After summarising the scan's findings, this section examines the most common front-line use of AI in justice: general-purpose LLMs which millions of people are turning to for legal information, advice and assistance.

2.2 The context: AI developments in law firms and judiciaries

Formal justice providers are making extensive use of AI. Despite challenges (see Garg and Ma, 2025), lawyers are investing heavily, with large firms forming strategic partnerships with AI providers and building their own AI platforms (Scarcella, 2026). Competition is driving AI adoption in law firms, seen as an 'imperative' for productivity gains (Deloitte, 2026). AI tools for legal professionals, such as Thomson and Harvey, draw on proprietary legal research content (Scarcella, 2026; Thomson Reuters, 2026). Google claims that its Gemini Enterprise for Legal's stable of AI agents are able to handle legal and administrative functions without significant human oversight (Scarcella, 2026). A September 2026 LexisNexis survey reported that 94% of UK lawyers use generative AI for legal work, up from 46% in January 2025 (LexisNexis, 2026). England's first AI-driven law firm has now been authorised (Garfield AI, 2025).

AI is also being adopted at pace in judiciaries. Reviews including from the European Commission for the Efficiency of Justice (CEPEJ) and the OECD, and for the United Nations, document how AI tools are assisting judges, although with uneven impact and with challenges (see CEPEJ, 2024; CEPEJ, 2025; OECD, 2025b; OECD, 2026; Satterthwaite, 2025). China's judiciary makes particularly intensive use of AI (DCAF, 2026: 38–39). The 2026 World

Development Report showcases the Kenyan judiciary's use of AI to assign more than 10,000 court cases each year to mediators, helping to reduce case backlog (World Bank, 2026a).

2.3 Overview of ODI's global scan

Of the 77 bespoke front-line AI justice tools and uses, 64 provide or support legal information, advice and assistance, and 13 support informal dispute resolution. Most of the examples of legal information, advice and assistance (57) are public facing, with the rest (7) supporting human advisers. However, public-facing examples are easier to find through internet-based research, and there may be examples of adviser-facing tools that the scan failed to pick up. While the scan's focus was on front-line justice, it picked up more examples (95) of AI being used to support the internal operations of courts/judiciaries.

2.4 Legal information, advice and assistance

Public-facing bespoke AI tools

Bespoke AI tools are being used to provide legal information, advice and assistance directly to the public. There is wide variation in:

- ease of use (from clunky, with lengthy sign-up and registration processes, to immediate anonymous conversational access over voice or internet);
- range of support (from information to assistance); and
- scale (10,000 conversations a year to 255,000 a year).

Examples of the most complex tools are *Beagle+* in Canada (see Box 1) and South Africa's *Maya*. These are public-facing RAG chatbots that respond to conversational questions from the public on a broad range of everyday justice problems. They use general-purpose LLMs drawing on a defined curated body of legal information and experience of lawyers responding to justice problems. *Beagle+* uses this material to offer legal information, while *Maya* goes further and can provide legal advice and assistance, including drafting court documents. The scan includes narrower applications of AI, such as translation and transcription to improve website accessibility.

Sixteen of the 57 public-facing tools providing legal information, advice and assistance focus on a single legal issue, such as housing, family or employment. For example, developed by HiiL with Nigerien justice institutions and customary leaders, Niger's *Hadin kai-bani fondo* ('the Path for Cohesion') provides guidance in French on land and resource disputes, drawing on locally validated content. It combines a web platform with *Illimi*, a voice assistant accessible through a basic mobile phone. As in many countries, customary dispute resolution is widely accepted in Niger, and traditional chiefs

have a statutory power to conciliate customary, civil and land disputes. As such, the tool was designed to support customary leaders in their conciliation work, and is also available to community members (Brun et al., 2025).

In Tanzania *Dada Wakili* supports women dealing with land, property and inheritance rights issues, where statutory, customary and religious legal systems interact. It is built on a structured corpus of Tanzanian inheritance statutes, case law and paralegal training materials, and provides guidance in Kiswahili and English via WhatsApp and SMS.

In jurisdictions where customary justice is widely used, such tools can give all parties, including chiefs, community members and community paralegals, the ability to access the same reliable source of advice in the local language, supporting fair decision-making that appropriately reflects national laws. AI can also offer enhanced transparency: AI-enabled recording, publication and translation of hearings could help spread best practices and reduce the risk of ill-informed, biased or corruptly influenced decisions.

Some public-facing AI acts as the entry point for users to access a human-delivered service. An example is India's *Nyaya Setu*, which offers guidance and can draft simple notices and applications, but also refers users to *Tele-Law* panel lawyers. Another example is Pakistan's *NAZ Assist*, which has a call-back facility to enable users to check AI-generated advice through a free consultation with a lawyer. Nigeria's *Podis* uses AI to triage clients and connect them with *pro bono* lawyers.

Regulation of legal practice can confine chatbots to offering legal information and general guidance. Where the regulatory environment is less restrictive, chatbots can provide the public with a more comprehensive service. For example, in South Africa *Maya* provides advice, drafts court documents, organises evidence and provides court submissions.

Adviser-facing bespoke AI tools

The scan revealed four examples of legal chatbots for advisers and volunteers at legal advice organisations. As with public-facing chatbots, these are grounded in curated local legal sources – legal databases, legislation and case law. One of the best evaluated is *Caddy*, for the Citizens Advice service in England and Wales.

Cost-effective front-line justice services tend to be provided by non-lawyers – such as community justice workers, community-based paralegals and case workers (Manuel et al., 2025). AI has the potential both to enhance their capability and increase productivity.

Human oversight

As well as being used to support services delivered by humans, AI tools have human oversight embedded within them – usually through

ex post checks. Niger's *Hadin kai-bani fondo* logs every interaction for monthly human review and corpus updating (Brun et al., 2025). *Beagle+* uses AI (Claude) to review every conversation, escalating those it flags as high-severity or problematic for review by a lawyer (see Box 1). *Caddy* has an approval button recording that a named supervisor authorised the advice provided.

2.5 AI in informal dispute resolution

Informal (broadly non-court) rules-based online dispute resolution (ODR) operated before the advent of AI. In the private sector, eBay was by 2011 handling 60 million disputes a year without human input – more than the entire US civil court system (Katsh and Rule, 2016: 333).

The public sector record on ODR is mixed: British Columbia's online Civil Resolution Tribunal (CRT) is widely seen as best practice, but usage is low (it closed only 6,928 cases in 2024/2025), and unit costs are relatively high ($1,510 per dispute) (CRT, 2025: 54–55). The European Union's consumer ODR platform and the Netherlands' civil and family *Rechtwijzer* and its scaled-down successor, *Uitelkaar.nl,* had low take-up and were withdrawn (Legal Futures, 2017; HiiL, 2017; see also Cashman and Ginnivan (2019) and Cortes (2018a and b) for reviews of pre-AI ODR).

There is now potential for everyday justice problems/small-scale disputes to be resolved speedily using AI. The UK Ministry of Justice's 2025 AI Action Plan for England and Wales promotes the use of 'upstream' AI to triage, advise and encourage civil and family parties towards out-of-court resolution (Ministry of Justice, 2025; 2026a). The newly designated President of the UK Supreme Court, Lord Briggs, has noted that AI could go further, and provide determinations of small-scale, low-value disputes, depending on public acceptability:

> *I have described this as a 'democratic question' because it ultimately boils down to what the public want, or are prepared to accept … The recent arrival on the scene of Agentic AI makes it by no means implausible that an AI platform could soon be tasked with the whole process of coming to a decision on a case* (Briggs, 2026: 14–15; see also p. 20).

The scan revealed no current examples of AI providing determinations with reasons. It did, however, provide examples of dispute resolution organisations making some limited use of AI. In India, for example, it is now easier to lodge a complaint with the National Consumer Helpline, which uses AI for speech recognition and to classify complaints. There were also examples of AI being used to assist parties with negotiating settlements and predicting outcomes. In the United States, *Mighty* is a free-to-use outcome

predictor for personal injury. But it is new and privately owned, with no publicly available information. Australia's publicly owned *amica* focuses on a single issue – asset splits to assist divorcing couples to reach out-of-court settlements – but there seems to be only modest take-up to date. In the UK a free public online predictor tool is under development (Nuffield Foundation, n.d.). One reason for the limited investment in informal dispute resolution could be negative experiences with earlier online dispute resolution technology.

The use case where AI is making a real impact in non-court dispute determination is in assisting large repeat players with bulk claims. Commercially developed AI tools assist institutional bulk debt collectors to resolve large volumes of low-value payment disputes such as dishonoured cheques and parking fines. India's *CORD*, *Sama* and *Saya* are examples. In Spain, one response to the 2025 introduction of compulsory pre-litigation mediation has been the development of commercial AI negotiation assistants (such as *MASC*) for repeat-player law firms. These tools are inaccessible to single-dispute unrepresented parties.

2.6 Global reach, but limited scale

The scan revealed bespoke front-line justice AI examples in 39 countries across Africa, the Americas, Asia, Europe and Oceania. These range from some of the highest income (US) to the lowest (Niger). However, there were few examples from lower-income countries. A third come from India, the US and the UK, which in part reflects the authors' location and past research connections, resulting in an inherent bias towards English-speaking countries.

Bespoke tools are being adapted to context

The 2026 World Development Report on AI stresses the importance of AI being appropriate to local contexts, with justice identified as an area where AI poorly adapted to local conditions can do real harm due to inaccuracy and bias (World Bank, 2026a).

The need for AI tools to be designed for context is particularly acute in low-income/low-resource environments. Challenges can include limited digital literacy, cost, access and device usage, connectivity challenges, users' over-trust in technology, legal pluralism and the inability of AI to interpret local norms and community discretion, cultural sensitivity and language challenges (Ogoniba, 2023; Okeke, 2026; Stephens, 2025a; b; Stephens and Nchembi, 2025; UNDP, 2026a; WJP/World Bank, 2025; World Bank, 2026a).

AI tools in the scan address these challenges. They use context-specific curated bodies of legal material. Connectivity challenges are addressed by using WhatsApp and SMS, for example in Kenya, Nigeria and Sierra Leone. Niger's *Hadin kai-bani fondo* is a phone service with AI interactive voice recognition. Translation is a key

feature, with examples in Kenya and Nigeria. India's *Nyaya Setu* chatbot accepts spoken queries in more than 20 languages.

India has one of the most advanced AI ecosystems
India has the most examples of front-line AI justice in the scan (11). It has a wide range of public and private providers, paid-for and free services, and all front-line justice functions (information, advice, assistance and informal dispute resolution) are covered.

Few bespoke tools operate at scale
Assessing the extent to which bespoke AI tools have enabled scaled-up access to front-line justice is difficult. Most do not report coverage figures. Where figures are provided, they are self-reported, some from commercial developers, and it is beyond the scope of this paper to assess their validity.

Some of the largest numbers are reported by organisations making limited use of AI. India's National Consumer Helpline, using AI for speech recognition and to classify complaints, deals with 1.9 million annual consumer complaints.[3] More sophisticated and potentially transformative AI tools providing legal information, advice and assistance directly to the public have had limited uptake for reasons that are not clear, but possibly because they are still at an early stage. *Maya* has one of the largest numbers of annual users, 255,000,[4] but this is less than 1% of the adult population.

2.7 Where most people go: general-purpose LLMs

General-purpose LLMs are rapidly changing the face of front-line justice. A December 2025 survey reports that 65% of US adults have used an AI chatbot for legal help (Hollenbeck, 2026). A UK survey also at the end of 2025 showed a much lower level of use: 16% of adults who had reported a legal dispute in the last two years (JUSTICE, 2026). However, this is likely to increase given the rate at which overall use of general-purpose LLMs in the UK is increasing, and since 69% of respondents said (in December 2025) that they trust general-purpose LLMs for legal advice at least somewhat (Enock and Margetts, 2026).

The explosion in public interest in general-purpose LLM use for justice issues is unsurprising given the global justice gap. To give some country examples: in England and Wales 32% of adults who experience a legal problem receive no professional support (LSB and The Law Society, 2024); in the US, less than 10% of people on a low income get any or enough legal help with substantial legal problems (Legal Services Corporation, 2022). General-purpose LLMs seem more popular than bespoke AI tools, even where these tools have the greatest reach: in British Columbia, 10% of adults have used

[3] December 2024 figure annualised.
[4] Annualised figures based on first eight months of 2026. Source: interview with authors.

People's Law School websites, against the newer *Beagle+* chatbot, where usage is only 0.2%. As noted above, less than 1% of adults have used the *Maya* chatbot.

General-purpose LLMs are also impacting organisations providing free legal support. HiiL's survey of legal aid providers suggests widespread use of general-purpose LLMs in their operations, particularly ChatGPT (HiiL, 2025). A recent AI policy template for UK *pro bono* legal advice clinics addresses the risks of using AI, and the importance of using only approved AI tools and for permitted purposes (LawWorks, 2026). The extent to which general-purpose LLM use by legal aid providers is overseen, supported by training and guided by safeguards in practice is unclear.

As discussed in the next section, using general-purpose LLMs for legal advice can be problematic due to issues around accuracy, confidentiality, sycophancy and bias – which can be particularly acute when used by people with little or no knowledge of law.

2.8 Conclusion

General-purpose LLMs are changing the face of front-line justice for both the public and front-line legal advice organisations. But using general-purpose LLMs in this way also presents serious risks. Private sector law firms, rather than relying on general-purpose LLMs, are adopting bespoke AI tools to enhance their service to clients able to pay.

Bespoke AI tools are also becoming available at the front line. There is an emerging global stable of both public- and adviser-facing bespoke legal chatbots based on trusted, contextually appropriate bodies of legal material. Given the size of the justice gap and how little affordable legal advice is available for everyday problems, uptake of public-facing legal chatbots is strikingly low. Part of the reason may be that they are still at an early stage and are not yet widely known. Other barriers to uptake are discussed in Section 4.

There are few examples of AI transforming informal dispute resolution. The private sector (eBay, for example) has demonstrated the efficacy of non-court ODR in resolving small-scale disputes without human intervention. But this experience has not carried across to the public sector, where use is limited and unit costs are high. AI's main impact in dispute resolution has been to assist large players with bulk claims.

There is growing policy interest in 'upstream' AI to triage, advise and nudge parties towards out-of-court resolution and, depending on public acceptability, in AI determinations of small-scale, low-value disputes. How far AI should go in deciding disputes is a democratic question to do with what the public would accept. At the same time, the growing use of general-purpose AI tools is likely to increase

demand on the courts, and the resulting potential surge in claims suggests that testing these approaches is urgent.

The global scan reveals a world of initiatives that demonstrate AI's potential to support and transform front-line justice services. However, based on the interviews for this research, and the publications reviewed, this knowledge base is not widely known, even within the sector. This presents an opportunity for richer cross-country learning – especially lessons from non-OECD countries, and countries such as India, which has one of the most advanced front-line justice AI ecosystems – and for developing shared guidance on best practices.

Taken together, the key policy question that emerges is not whether AI will be used by the public to access justice, but how to make it safer and fairer.

Box 1 *Beagle+*

Beagle+ is a RAG chatbot developed by People's Law School, a not-for-profit organisation based in British Columbia, Canada. Funded mainly by the Law Foundation of British Columbia and the federal government, the School provides free legal education. It hosts two websites: People's Law School[5] and Dial-A-Law,[6] which in 2025 were collectively used 604,400 times (by 452,400 unique users). *Beagle+* is trained on these websites (with automatic daily updates) and uses ChatGPT 5.4 to provide responses. It offers information about the law, explains legal rights, discusses options and lets people know about possible next steps they can take. The chatbot handled 8,604 conversations in the 12 months ending March 2026. It won the 2026 US Michael Mills InnovAction Award.

A key feature of the School's work on *Beagle+* has been its commitment to make public its evolving quality assurance process. To prepare for launch, it assessed Beagle+'s responses against ideal lawyer-crafted answers to 42 nuanced questions asked by the public of the previous version of the chatbot. A panel of lawyers scored each Beagle+ response for safety (was the information legally accurate and unlikely to mislead?) and value (did the tone, clarity and empowerment level support the user in taking next steps?). The School also reviewed differences in system prompts, temperature and retrieval chunking.[7] It published the full dataset to provide a compact benchmark for evaluating the system's ability to handle real user queries with appropriate caution, jurisdictional relevance and

[5] https://www.peopleslawschool.ca/

[6] https://dialalaw.peopleslawschool.ca/

[7] System prompts are the standing instructions that tell the tool how to respond; its temperature is a setting that controls whether its answers are consistent or more varied; and its retrieval chunking is the way reference documents are divided into smaller pieces so the tool can find relevant passages.

helpfulness. Over the first two years of operation of Beagle+, lawyers at People's Law School reviewed all conversations for safety. After two years, the School moved to an interim system that triggered review based on fixed numeric thresholds from automated scoring. Across three review phases, approximately 15,000 items were reviewed for safety: the 42 pre-launch test questions, about 5,400 live conversations reviewed by lawyers, and about 9,600 later conversations screened by an AI tool, with lawyers reviewing those it flagged. Of the roughly 15,000 Beagle+ responses reviewed, approximately 99% were rated safe.

A response was rated safe on the basis of whether it is legally accurate, whether its tone suits the seriousness of the situation, and whether it directs the person to legal representation or other services where the matter calls for it**.** In the latest system, *Claude* reviews every conversation, and anything it flags as high-severity or otherwise problematic goes to mandatory manual review. In earlier internal testing, *Claude'*s assessments matched the team's own review of the same conversations substantially more closely than the automated scoring it replaced. As part of the School's ongoing commitment to review Beagle+, an independent academic evaluation is underway.

The School is now analysing and making available the costs of developing and maintaining *Beagle+*. The full development cost was $110,000, consisting of $60,000 for the original build cost plus $50,000 for staff time. These figures exclude the cost of the underlying website content, since that work is treated as already done. Two-thirds (64%) of the ongoing annual running costs of $58,000 are for the lawyers reviewing conversations. Ongoing trademark, software maintenance and LLM/API costs account for less than a fifth (17%), with marketing costs making up the balance.

The current ongoing unit cost of *Beagle+* is $6.80 per conversation. As the use of the chatbot grows, the School expects the unit cost to fall, as the need for lawyer oversight would not need to increase 1:1 as usage increases. As website use is over 600,000, and the number of people with legal problems in British Columbia is over 500,000[8] (and some will have more than one problem), one potential scenario would be for *Beagle+* to scale up to 360,000 conversations, corresponding to 65% of the total number of people with legal problems. While the costs of LLM/API would rise in line with usage, the need for lawyer oversight would not rise at the same rate. Assuming only a fourfold increase in lawyer oversight for every tenfold increase in usage, the annual running cost for Beagle+ is projected to be $425,000, and the unit cost would fall to $1.20 per conversation. So, if *Beagle+* could help reduce demand on other

[8] Canadian Legal Problems Survey (Statistics Canada, 2022).

parts of Canada's justice system, there would be a major improvement in overall cost-efficiency.

Scaling up *Beagle+* would require a small rebalancing of funding. An annual running cost of $425,000 would require a significant (60%) increase for the School (as its annual budget is only $730,000), but this would be only 0.4% of the total funding for legal aid in British Columbia (Legal Aid BC, 2026).

Sources: People's Law School (2024 and 2025b) supplemented by information provided by the Executive Director and staff at People's Law School (June–September 2026) on costings. Full details of the current and projected costings are available on request from the School (or via the authors).

3 Risks and evidence gaps for AI at the front line

3.1 Introduction

The previous section highlights how general-purpose LLMs are changing the face of front-line justice. This section starts by examining the problems they pose (accuracy, confidentiality, sycophancy and bias), and considers how bespoke AI tools, if more widely available, could help to address these issues. This leads into a discussion of the wider risk that AI, as currently deployed, may be increasing inequality in access to front-line justice – between people and countries with access to appropriate AI tools, and those without. General AI risks and mitigations in the context of front-line justice are also covered.

The section ends with a review of the evidence gaps on AI-assisted front-line justice. It concludes that, while the gaps are significant, there are clear advantages of well-designed bespoke legal chatbots over general-purpose LLMs for delivering front-line justice. Appendix 2 sets out the characteristics of such bespoke public-facing chatbots.

3.2 The risks of general-purpose LLMs

The widespread use of general-purpose LLMs for legal advice and assistance could suggest that the justice gap is being addressed at zero cost to the public purse. However, there are profound concerns about the quality and appropriateness of the advice provided. With the caveat that the company has an interest here as a provider of an AI legal assistant for lawyers, a 2026 LexisNexis survey of UK lawyers reported that, although 94% of them use AI for their legal work, 81% demanded bespoke AI tools grounded in legal sources over general-purpose LLMs (LexisNexis, 2026).

Accuracy

While there is no standardised measure of accuracy, general-purpose LLMs are capable of giving reasonably accurate answers to well-defined everyday legal questions. However, there are concerns about their consistency and reliability, including evidence that they can alternate between correct and incorrect answers to legal

questions depending on phrasing (Faghih et al., 2026).[9] An independent assessment in August 2025 (before ChatGPT 5 was launched) recorded 80% accuracy on well-phrased questions on US law (Vals AI, 2025). In the UK, a slightly earlier assessment (December 2024) by global law firm Linklaters found that accuracy on English law could be as high as 90% for a simple question when the area of law is well known and discussed, although overall accuracy was only 26% against their full set of questions (Linklaters, 2025). Linklaters also noted how quickly general-purpose LLMs were improving: overall performance had doubled in just over a year since their first assessment (October 2023). The only recent assessment ODI Global could find of ChatGPT 5 was by a Canadian panel of two lawyers who reviewed how it handled two everyday legal problems (Dylag, 2026). The conclusion was that the responses mostly identified the relevant issues and provided an accurate summary of the law, but that sometimes the advice was incomplete, included minor errors and did not evaluate options (e.g. the benefits of taking a particular approach and the chances of success). While general-purpose LLMs are improving, OpenAI's own reporting on hallucination rates in general notes that ChatGPT 5.2 returned errors 6% of the time (OpenAI, 2025).

General-purpose LLMs may struggle more in non-Western jurisdictions, where they are less familiar with local languages and legal texts. An assessment in India in 2024 (which included ChatGPT 4) reported that general-purpose LLMs were scored highly by law students on legal advice, but only satisfactorily on research, where they were prone to hallucination (Hemrajani, 2025). A more recent Indian assessment (which included ChatGPT 5) judged that general-purpose LLMs outperformed the top-scoring law student in law exams, although there were problems with 'manufacturing authorities' (Juvekar et al., 2025).

It is possible that general-purpose LLMs are 'accurate enough' to address simple, everyday legal problems for the public. The President of the UK Supreme Court, Lord Briggs, considers that general-purpose LLMs enable the public to prepare submissions that would be 'sufficient to get the complainant through the court door' (Briggs, 2026: 10).

There is evidence that well-designed bespoke legal chatbots (see Appendix 2) can provide better/more accurate answers, once they are fully developed. A key issue is the 'quality' of the questions asked; for example, people could ask questions based on a false premise or that are unclearly or wrongly framed. The counterfactual also needs to be considered: advice from human lawyers may not always be entirely accurate or helpful. The issue of professional indemnity insurance is discussed later.

[9] The paper suggests that consistency testing (self-paraphrasing/ensemble prompting) offers a cheap mitigation.

An independent audit of *Legal Aid* in North Carolina in the first few months of its chatbot's operations found multiple problems; for example, on housing questions half the answers were too basic to be helpful, and 20% were wrong. There were improvements in later iterations (Porcaro et al., 2025), and the organisation reports that the audit catalysed rapid improvements (Porcaro et al., 2025). The *Alaska Virtual Assistant (AVA)* legal chatbot has also struggled with accuracy. The administrative director of the Alaska Court System, and one of the project's leaders, commented in early 2026 that they were not confident that the *AVA* would be able to replicate what human facilitators at the self-help centre were able to do, due to inaccuracies and some incompleteness (Perlo, 2026).

There are more positive results. A (rare) randomised controlled trial (RCT) of the *Caddy* tool found that 80% of responses were good enough to go to advisers without revision (Varotsis, 2025). An audit of Nevada's court chatbot, giving legal advice to self-represented litigants, found that 80% of responses were helpful (Porcaro et al., 2026). *Beagle+* achieved a pre-launch 'safety' rating (in relation to legal accuracy and appropriateness) of 70% measured against a set of defined questions. The rating improved in further reviews. Overall, in three phases, approximately 15,000 items were reviewed for 'safety': 42 pre-launch test questions, about 5,400 live conversations reviewed by lawyers, and about 9,600 later conversations screened by an AI tool, with lawyers reviewing those it flagged. Of the roughly 15,000 *Beagle+* responses reviewed, approximately 99% were rated safe (see Box 1) (People's Law School, 2025a; b).

Confidentiality

People seeking answers to legal questions often need to share highly personal and potentially sensitive information. Exchanges with general-purpose LLMs are retained for a period and recent press reports have noted that even private prompts have been read by their staff (Anton, 2026).

Well-designed bespoke chatbots (see Appendix 2) run by accountable, trusted organisations cannot eliminate confidentiality risks, but they can mitigate some: the underlying information base holds no personal data, the operator is an accountable organisation rather than a distant platform, the chatbot can alert users to confidentiality risks before they share personal information, and operators can choose to remove identifying details before a query leaves the system, or ensure that any such information is automatically deleted in real time.

Alternatively, bespoke chatbots can be built on open-source models that the organisation runs on its own servers. Client information then never passes to an external AI provider, which reduces confidentiality risks, although the organisation becomes responsible for keeping the system secure. This option depends on the availability of secure local servers.

Self-hosting reduces the risk that confidentiality is lost to a third-party provider, but the position on legal privilege (whether exchanges with AI are shielded from disclosure, in the same way as conversations with lawyers) is unclear. A US district court in 2026 held that they do not attract legal privilege and must be disclosed.[10]

Sycophancy

A common concern with general-purpose LLMs is that they are designed to engage the user and so mirror the user's desires – personalised sycophancy. Experiments on people's willingness to act on LLM-generated advice, as compared with lawyer-generated legal advice, suggest that people may actually prefer general-purpose LLMs to lawyers because the advice provided is confident – possibly more encouraging and welcoming, but also less likely to enumerate risks (Schneiders et al., 2025). Competent legal advice goes beyond accuracy to ensure the advice is safe, identifies the objectives and best interests of the client, elicits all the relevant facts, and works with the client to develop a strategy to achieve an objective (Dylag, 2026).

Nori (2026b) argues that a well-designed purpose-built legal chatbot can be meaningfully safer than a general-purpose LLM. Although it typically runs on the same underlying models and inherits some of their weaknesses, it can be restricted to vetted content and can include deterministic hard stops (for example, bypassing generated text entirely and providing immediate contact information when a query indicates a domestic violence emergency). It can also be systematically tested and monitored in ways that are not possible when people use a general-purpose chatbot directly.

Bias and limited relevance to diverse communities

There are also concerns that general-purpose LLMs are not trained on legal language or documents that reflect the legal problems of diverse socioeconomic, educational and geographic groups (Laskowski et al., 2023, cited in Okeke, 2026; Amnesty International, 2026). Well-designed bespoke legal chatbots based on a trusted body of locally relevant and context-specific legal material can address this (Nori, 2026b). This is the approach adopted with South Africa's *Maya* chatbot, which is trained on manuals developed by local paralegals.

Human rights considerations

The UN has noted the human rights risks associated with the use of AI in the administration of justice (UN Secretary-General, 2024). Amnesty International has raised specific concerns that general-purpose LLMs breach human rights standards, including because of ingrained biases, over-censoring of global majority languages and automation bias. Amnesty specifically makes a distinction between

[10] *United States v. Heppner*, No. 25-cr-00503 (United States District Court for the Southern District of New York, Feb. 17, 2026).

general-purpose LLMs and domain-specific, locally trained AI tools (Amnesty International, 2026).

General-purpose LLMs are causing problems downstream

There are growing concerns that general-purpose LLM-generated legal advice is causing problems downstream in formal courts and tribunals. Sargeant notes that, in England and Wales, there is no systematic evidence that AI is increasing the number of claims (Sargeant, 2026: 72). However, Lord Briggs refers to a likely 'tsunami' of AI-generated claims (Briggs, 2026: 11) and the Economist magazine has highlighted the risk of a 'deluge' of claims with unnecessarily long and complex submissions generated by general-purpose LLMs (The Economist, 2026a; 2026b).

As for flood management, knowing the height of any impending wave is key. While the Economist referred to the risk of a 100-fold increase, ODI Global's analysis of past survey data across the civil justice system in England and Wales (WJP data) suggests the increase could be fivefold at most in the civil justice caseload.[11] However, even a doubling would significantly increase backlogs and delays, which would worsen, rather than improve, people's ability to get justice. Planning for any potential increase requires better and more current data on unmet legal needs. The last legal needs survey in England and Wales was conducted in late 2023, when usage of general-purpose LLMs was much lower. As the problem is potentially global, it is concerning that most legal needs surveys date from 2018–2020 (WJP, n.d.).

Helping courts and judges work more efficiently with AI may be part of the answer, but the problem can also be tackled upstream. Bespoke AI chatbots could be more effective than general-purpose LLMs at discouraging people from pursuing claims that are unlikely to succeed or are not in their best interests, reducing the number of weak claims taking up court time. There is also an equality dimension. Large institutions already use commercial AI tools to issue debt recovery claims in bulk. Defendants typically have less legal capability than these institutional claimants. Giving them access to bespoke AI tools to support their defence would help restore equality of arms.

An additional way of dealing with an increased judicial caseload could be to divert low-level and straightforward cases out of the formal justice system altogether. Steering suitable cases into speedier, more informal dispute resolution is known to be cost-effective (JAC, 2025, drawing on Manuel et al., 2025). AI tools that predict outcomes and find zones of agreement are being explored, and there may be scope for AI determinations in suitable cases.

[11] Drawing on WJP data, the authors estimate the number of civil legal problems that are sufficiently severe to need legal assistance to be 16 million in England and Wales, compared to the current number of civil cases filed with courts, tribunals and ombudsman offices of 3.6 million (precise ratio 4.4:1).

Future general-purpose LLM developments and the case for bespoke chatbots

While there are clear potential advantages of bespoke chatbots over general-purpose LLMs, it is less clear whether these advantages will persist in the future. The capacity of general-purpose LLMs has increased rapidly in the last few years and is expected to continue to improve. But while these improvements are likely to reduce the risk of hallucinations and improve accuracy (narrowly measured), they are unlikely to overcome issues around confidentiality and sycophancy. Well-designed bespoke chatbots already mitigate these two concerns. If hosted by trusted front-line justice organisations, issues of legal privilege could be dealt with. In addition, the case for bespoke chatbots could become much more pressing if general-purpose LLMs were ever to decide to withdraw from responding to legal queries. Such a decision could be prompted by either the risk of litigation (as has started to happen for medical queries) or the risk of future regulation.

3.3 The risk of AI increasing inequality at the front line

The UN has noted the risk that AI could increase inequality in the administration of justice (UN Secretary-General, 2024). A recent UNDP regional survey in Asia and the Pacific warns that, without deliberate action, AI risks widening justice gaps, based on unequal digital access, language coverage and institutional capacity (UNDP, 2026a). The evidence in this paper suggests that bespoke AI tools designed for local contexts can address these problems and offer the potential to enhance access to justice.

However, current AI funding and deployment patterns could increase inequality in several ways:

- Bespoke AI legal tools are being developed commercially by and for private sector law firms. Only those who can afford a lawyer will benefit from these AI developments. Public-facing bespoke legal chatbots are rare, mainly in early stage and not scaled up. For most, the only available public-facing legal advice tool is a general-purpose LLM. The efficacy of adviser-facing AI tools to make a significant impact on resolving people's everyday justice problems depends on a nationwide network of front-line justice advisers – unavailable in most countries. In the UK, what has been seen as an emergent 'two-tiered justice system', in which only the resourced and capable can enforce their rights (Curran et al., 2025), could be accentuated by AI (JUSTICE, 2026).
- Because AI involves up-front costs, it is repeat players (such as law firms and bulk debt collectors) who are best placed to pay for AI tools to enhance their efficiency.

- Front-line justice tends not to be a political priority (JAC, 2025), though a significant increase in general-purpose LLM-generated claims before formal courts could increase the attraction of intervening and investing 'upstream' in front-line justice.
- The global scan revealed only a few examples of bespoke AI front-line justice tools from lower-income countries, where challenges to AI adoption include limited digital literacy, cost and connectivity problems.

3.4 General AI risks and mitigations

The inherent risks of relying on general-purpose LLMs to provide legal advice and assistance reflect well-known AI risks in public administration. As applied to the justice system, they can be summarised as:[12]

- **Reliability and accuracy:** Inaccurate or fabricated outputs coupled with the risk of user over-trust in AI due to its perceived authority and apparent plausibility (the 'fluency trap'). Creation of fake information and the threat of AI agents going rogue.
- **Transparency and accountability:** The need to preserve due process to prevent algorithmic bias in decision-making leading to inequity, discrimination or exclusion. Bias in public-facing advisory tools like chatbots shaping the advice given. The risk of algorithmic bias even in routine tasks like transcription and translation.
- **Privacy, confidentiality and security:** Breaches of privacy and confidentiality where case/client data is processed by third-party systems.

Policymakers are responding to these risks. The EU's 2024 AI Act[13] imposes direct limitations on AI use, and the General Data Protection Regulation article 22[14] restricts automatic decision-making with no human assessment.[15,16] Consumer protection regulation is also seen as a way forward, with legislation[17] being considered by the New York Senate that would impose liability on AI developers for professional

[12] See for example Legal Services Research Centre, 2026; Barocas et al., 2023; Cobbe et al., 2023; Coglianese, 2024; HiiL, 2025; IDLO, 2026; Levy et al., 2021; McGrath and Baidya, 2026; New York State Bar Association, 2026; Raworth, 2026; Satterthwaite, 2025; Sinha et al., 2026; Stephens, 2025a; UNDP, 2026a; WJP/World Bank, 2025; World Bank, 2026a.

[13] EU AI Act, Regulation (EU) 2024/1689. Adopted: 13 June 2024. Published in the Official Journal: 12 July 2024. Entered into force: 1 August 2024. CELEX number: 32024R1689

[14] General Data Protection Regulation (EU) 2016/679, art 22. Adopted: 27 April 2016. Published in the Official Journal: 4 May 2016 (OJ L 119) Entered into force: 24 May 2016. Became applicable: 25 May 2018. CELEX number: 32016R0679

[15] The GDPR gives individuals a right not to be subject to decisions based solely on automated processing that have legal or similarly significant effects on them, unless an exception applies (contractual necessity, legal authorisation or explicit consent), in which case safeguards such as a right to human intervention must be provided.

[16] OECD's survey of digital government reported slower AI progress in EU countries, attributed largely to regulation without accompanying support for controlled innovation such as sandboxes (OECD, 2026b).

[17] Bill S7263.

advice provided by their chatbots. The issue of harmful professional advice is also being litigated in the US. In *Winters v. OpenAI*,[18] the complaint alleges that ChatGPT's personalised medical diagnoses and recovery instructions were dangerous and caused harm (Nori, 2026b). A related issue is the potential to outlaw/negate general-purpose LLM liability disclaimers. If liability were to rest with AI developers (including with front-line justice providers), consideration would need to be given to professional indemnity insurance requirements (Necz, 2024).

A range of mitigation measures to address AI risks have been incorporated into a growing body of policies and guidance in the justice sector. These include transparency about AI use; human verification obligations; training on the use of AI; supervision, monitoring and assessment; routes for escalation to a human; redress mechanisms; citizen involvement; and independent oversight (Open Government Partnership, 2026; Shukla, 2025; UNDP, 2026b; UNESCO, 2025).

Many countries have adopted (or are in the process of adopting) policies and guidelines on the use of AI in their judiciaries, including Brazil,[19] Colombia,[20] England and Wales,[21] India,[22] Kenya[23] and Spain.[24] CEPEJ has developed a European Ethical Charter on the use of AI in judicial systems (CEPEJ, 2025). Guidance on AI use is also being developed for law firms and lawyers by their professional bodies. In England and Wales, the Solicitors Regulation Authority has issued a warning notice on the use of AI (Solicitors Regulation Authority, 2026).

What is missing is guidance for front-line organisations, and AI developers, on AI use in front-line justice. In England and Wales, the Legal Services Board (LSB) has identified an 'expectation-reality gap' between consumers' expectations of how AI is used (including lawyer oversight of AI outputs, privacy, accuracy and access to a human) and the reality. The LSB notes that standards specifically for public-facing AI legal tools are almost entirely absent and most standards that do exist are non-binding (LSB, 2026).

[18] Winters v. OpenAI Group PBC et al., filed in San Francisco County Superior Court.
[19] https://atos.cnj.jus.br/atos/detalhar/6001
[20] https://rm.coe.int/Columbia-guidelines-for-the-use-of-artificial-intelligence-in-the-judi/1680b53483
[21] https://www.judiciary.uk/guidance-and-resources/artificial-intelligence-ai-judicial-guidance-october-2025/
[22] https://cdnbbsr.s3waas.gov.in/s3ec0490f1f4972d133619a60c30f3529e/uploads/2026/06/2026060342.pdf
[23] https://judiciary.go.ke/judiciary-to-leverage-ai-to-enhance-justice/
[24] General Council of the Judiciary's AI instruction.

3.5 Evidence has not kept pace with AI developments

Commentators have noted the extremely limited evidence base on the effectiveness and impact of AI use in justice systems (Sargeant, 2026; World Bank, 2026a), with no agreement on even basic metrics (Tashea, 2026). Outcome and impact concerns are rarely accounted for in policy frameworks or in evaluations of AI deployments (Deshmukh and Davies, 2026; OECD, 2025b). A review of AI tools in the UK justice system found that fairness, comprehension, lived experience, differential impact and procedural legitimacy are rarely considered in evaluating the performance or impacts of an AI system (Sargeant, 2026). The OECD has identified monitoring and evaluation as a critical weakness of generative AI deployment in the public sector in general, with few governments systematically measuring performance, impact or compliance (OECD, 2026a).

Evidence of performance, effectiveness and impact of front-line AI-assisted justice

Given the evidence gaps on AI deployment across governments, it is unsurprising that data on performance and effectiveness (let alone impact) of AI deployment for front-line justice is sparse. Of the 77 live tools and uses considered in the global scan carried out for this research, only 22 publish data on the number of cases/queries handled, and only 12 provide measures of accuracy and/or evaluation data. As Sargeant (2026) points out, where there has been performance evaluation, findings tend to be internal or vendor-led. Two rare examples of independently audited tools are Duke University's audits of Legal Aid in North Carolina and the Nevada Court's self-help chatbot. There are also plans for Duke to audit the Alaska Court System's AI virtual assistant (Porcaro et al., 2025; 2026).

The lack of data and evidence is consistent with ODI Global's previous findings of general justice sector underspend on research and development, compared with other sectors such as education (Manuel and Manea, 2025; Sorabji, 2026).

Evidence on AI-assisted front-line justice is needed at various levels ranging from basic legal accuracy to the usefulness of AI tools, including the extent to which they actually contribute to the resolution of justice problems, in line with the OECD's developing Justice Services Measure (OECD, 2026c; Gramatikov, 2026). Legal needs surveys will have to be adapted to better capture emerging forms of AI-enabled assistance (Pleasence and Balmer, 2025). Porcaro (2026) proposes a five-point scale for audits of legal chatbots, ranging from harmful (an incorrect or misleading answer that would likely cause the user some harm) to helpful and actionable (which gives the user a specific, actionable next step). The People's Law School assesses its *Beagle+* public-facing chatbot with reference to safety (was the information legally accurate and unlikely to mislead?)

and value (did the tone, clarity and empowerment level support the user in taking next steps?) (see Box 1).

Front-line legal advice organisations need information about how to incorporate AI into their operations. The evidence from both the private sector and government is that organisational change is needed for AI returns to be realised (Challapally et al., 2025; Michalski et al., 2025). A UK study on the use of Copilot in the civil service shows that time savings were dependent on onboarding, training and change management (Arzilli et al., 2026).

A key issue for front-line justice is the extent to which users (particularly those with vulnerabilities) need interaction with a human to enable them to resolve their justice problems. *Legal Aid* in North Carolina reports that its AI legal advice chatbot logged 50,000 conversations in its first year of operation. It will be interesting to see whether the number of users rises to match the 400,000 users a year of its telephone helpline, or whether most users continue to need (or want) human input.

Evaluations of AI use for front-line justice services need to consider the counterfactual – for example, the alternative to an AI-provided service could be no service at all, rather than a perfect lawyer. Ideally the counterfactual would also consider the quality of existing services, such as legal advice clinics. For an AI informal dispute resolution mechanism, possible alternatives could include non-AI technology such as a rules-based decision tree. Very few front-line justice tools have been subjected to RCTs enabling these kinds of comparisons. Rare examples include *Caddy* in the UK (Varotsis, 2025) and Mexico City's Labour Court outcome predictor.

Given the limited evidence base, it is important that AI tools for front-line justice are built and developed to enable experimentation, failure, learning and change.

Value for money

A final aspect of evaluation should be value for money. Given the scale of justice needs and the funding constraints, there is little value in a highly effective AI pilot if high unit costs make it hard to scale. Value for money assessments of front-line justice tools are rare – even the cost of developing systems is rarely published. Evidence on value for money from ODI Global's scan is discussed in Section 4.

3.6 Conclusion

The rapid expansion in public use of unregulated general-purpose LLMs to provide legal information, advice and assistance has opened a route to justice for many. But there are serious concerns about accuracy, confidentiality, sycophancy and bias.

Well-designed RAG chatbots draw on the capability of general-purpose LLMs but then work only with curated local legal material. Run by independent, accountable organisations, their results are tested and overseen by lawyers. *Beagle+* provides legal information, with approximately 99% of its reviewed responses rated safe in relation to legal accuracy and appropriateness. In the case of *Maya*, South Africa's more permissive regulatory regime makes it possible to go further and provide advice and assistance, all the way through to preparing court documents and statements to be read out in court.

Bespoke AI chatbots are not a panacea: they inherit some of the weaknesses of the underlying models, and their accuracy depends on realistic budgets for ongoing oversight. They cannot eliminate risks around confidentiality, sycophancy and bias, but well-designed models can reduce and mitigate them. So they could be a game changer. And *Beagle+* and *Maya* provide proof of concept: it is possible to implement a well-designed bespoke AI chatbot that is superior to a general-purpose LLM.

At the same time, there is a clear risk that AI, as currently funded and deployed, could increase inequality in access to justice. Bespoke tools are being built commercially for law firms and for repeat players such as bulk debt collectors, while public-facing bespoke chatbots remain rare, mostly at early stage and with few examples in lower-income countries. For most people the only available AI legal tool is a general-purpose LLM. Policymakers are responding to general AI risks, and guidance on AI use is proliferating for judiciaries and lawyers, but guidance and standards for AI in front-line justice are largely missing. The risks of general-purpose LLMs, together with the inequalities caused by current AI deployment across the justice system, suggest that urgent action is needed.

AI use for front-line justice remains experimental. The evidence base has not kept pace with AI developments. Of the 77 tools in ODI Global's scan, only 22 publish data on the cases they handle and only 12 provide measures of accuracy or evaluation; independent audits and RCTs are rare; and value for money is almost never assessed. This reflects the general under-investment in research and development in the justice sector. Evidence is needed at every level, from basic legal accuracy to whether tools actually help people resolve their justice problems, assessed against a realistic counterfactual that may be no service at all, rather than a perfect lawyer.

4 Closing the justice gap: scaling up AI at the front line

4.1 Introduction

Section 3 concludes that there are clear potential advantages of bespoke chatbots over general-purpose LLMs. This section addresses the affordability of developing and running bespoke legal AI chatbots, based on in-depth analysis of British Columbia's *Beagle+* chatbot. It has not been possible to do a similar analysis for the other aspect of front-line justice – informal dispute resolution – because of the absence of data on costs.

The section ends with suggestions for how to overcome three key barriers to AI-assisted front-line justice: money (financing barriers); monopolies (regulatory restrictions on the practice of law); and mindsets (professional hesitancy).

4.2 Affordability and value for money

Development and running costs

Developing a bespoke front-line justice AI tool of the kind seen in the examples provided by the global scan involves up-front costs and investment. These can be perceived as being high, as it is the costs of systems developed for law firms that capture headlines. For example, Thomson 1.0, Thomson Reuters' proprietary LLM for legal, tax and regulatory professionals, cost $40 million to develop (Thomson Reuters, 2026).

However, a key insight of this paper, based on British Columbia's People's Law School's willingness to open their books, is that the costs can be much lower: *Beagle+* cost $110,000 to develop (see Box 1). In Niger, the initial technical development and testing of *Illimi*, the voice assistant behind *Hadin kai-bani fondo*, cost approximately $100,000. Building on an existing provider's voice platform kept this cost lower than it would otherwise have been. The figure excludes the research, co-design, development and validation of legal content, outreach and ongoing operation. Developing a similar voice assistant elsewhere could cost approximately $50,000, but this is an indicative estimate rather than a fixed replication price: the cost would depend on the existing infrastructure, languages, local content, telecom

arrangements and scope of the service. The experience with *Hadin kai-bani fondo* also highlights the need to fund ongoing operation and plan for provider changes (interview, September 2026).

In terms of running costs, recent experience with the *Alaska Virtual Assistant (AVA)* has highlighted that it can be challenging to deliver accurate low-cost bespoke AI chatbots in terms of both time and money. The developer noted that 'under one technical set up 20 *AVA* queries would cost only about 11 cents' (Perlo, 2026). However, the project took 15 months to launch (five times longer than originally planned) and struggled with accuracy.

The experience with *Beagle+* has been much more positive and will hopefully make it easier to create realistic budgets and timetables to develop and run chatbots elsewhere. It highlights that developing an accurate model takes time – 18 months (excluding experience with an earlier version of the chatbot). The full costs of maintaining an accurate and appropriate system need to be factored in up front. The unit costs of *Beagle+*, at $6.80 a conversation, are much higher than the AVA's cited cost of 11 cents (but much cheaper than a human adviser alternative). See Table 1 for details.

What is clear is that the running costs of maintaining an accurate system (which involves significant human oversight) can be lower and much more sustainable than previously thought (see Sargeant, 2026).

Reaching more people for less: AI's transformative potential

Previous ODI Global research shows that the key to closing the justice gap is adopting cost-effective approaches to front-line services and shifting the focus from lawyers and courts to community justice workers/community paralegals/case workers and to informal dispute resolution (JAC, 2025; Manuel et al., 2025). A range of countries have adopted this approach to legal advice and assistance, enabling non-lawyer front-line justice workers to assist clients to resolve their everyday justice problems, including in some cases representing clients in court (see Teufel et al., 2026). ODI Global has provided benchmark unit costs ranging from $20 a case in low-income countries to $790 in OECD countries, as well as extensive evidence that these benchmark unit costs are being achieved by non-lawyer front-line justice providers, with 52 examples from every region in the world (JAC, 2025).

Appropriate use of well-designed bespoke AI tools grounded in curated knowledge bases could have a potentially transformative effect by increasing cost-effective coverage in two ways.

First, AI can deliver **gains in the efficiency of front-line justice providers**. One of the best-documented examples of savings comes from UK Citizens Advice. A randomised controlled trial of their adviser-facing AI tool *Caddy* found it 'halved the response time for advisors seeking supervisory support, unlocking capacity to help

thousands more people' (Hagan, 2025). It also found that 80% of *Caddy'*s responses were good enough to go to advisers without revisions. Advisers with access to *Caddy* were more than twice as likely to say they were confident giving advice as those in the control group. They were more than one-and-a-half times as likely to say they could resolve the client's issue (Varotsis, 2025*).* Similar savings are reported by others. A study of customer-support agents at a software firm found that AI led to 34% more resolutions per hour by the lowest-skilled agents (Brynjolfsson et al., 2023).

As AI improves, it is possible that even greater savings could be achieved. But regardless of future AI developments, if the cost of providing accurate front-line legal advice and assistance could be halved, that would have a dramatic impact – making this kind of service even more affordable than the JAC and ODI Global have previously suggested. The JAC's Justice Financing Framework notes that, if well spent, allocating just 2.5% of a country's total justice sector budget[25] to front-line services could result in a fourfold increase in such services (in a typical OECD country) (JAC, 2025). If costs were halved, only a 1.25% share would be required. This could be found through a gradual, minuscule budget re-prioritisation of just 0.25% a year over a five-year period.

An alternative response to such an efficiency gain could be to double the reach of legal advice clinics. AI chatbots could also increase the number of people supported by empowering advisers in non-legal agencies, such as health care or mental health, to help clients with their justice problems without needing to refer them to a legal advice clinic.

For lower-income countries, where front-line justice services are often provided through customary and informal systems, the *Hadin kai-bani fondo* pilot in Niger provides early evidence that customary justice actors can use and find value in voice-based legal guidance. While the costs of customary and informal systems are already low, AI could reduce the major cost – training.

Second, the experience with *Beagle+* and *Maya* suggests that public-facing bespoke AI legal chatbots can provide **direct assistance to people with sufficient capability and digital access, without the need for adviser support**. If regulatory regimes allow, this could be the end-to-end legal support of the kind provided by *Maya* in South Africa, from the initial download of users' concerns through to preparing court documents and drafting statements for users to read out in court. A spillover benefit is that access to a trusted legal chatbot would reduce the need for users with capability and digital access to call on the services of legal advice clinics, enabling them to

[25] As defined in International Monetary Fund (IMF)'s and OECD's expenditure classification: Public Order and Safety.

spend more time with vulnerable, digitally excluded clients who do need their support.

Assessing the value for money case of public-facing bespoke chatbots

The gold standard for assessing the value for money case for public-facing bespoke chatbots would be a full cost–benefit analysis (CBA). This would go beyond measures of outputs (number of conversations) to look at outcomes (the impact that the conversation had on the client). As measuring outcomes is in general expensive, CBAs are undertaken infrequently in the public sector, even where outcome measuring is well-developed, as in the health sector. CBA is particularly rare in the justice sector, as noted in World Bank and OECD global reviews (Manuel and Manuel, 2021). A much more common assessment technique is cost-effectiveness analysis (CEA). CEAs are undertaken when the case for providing a particular type of service (or treating a specific disease) has already been established through CBA. The focus then is on finding the most cost-effective method.

Several justice CBAs show a high rate of return of providing legal information, advice and assistance. Citizens Advice in England and Wales delivers benefits 33 times the cost of the service, and community legal centres in Australia 18 times (Manuel and Manuel, 2021). There are also broader benefits beyond the resolution of individual legal problems. Enhancing access to justice strengthens democratic resilience and trust in institutions, helping people to understand and assert their rights (OECD, 2023; OECD, 2025a; Commission on Legal Empowerment of the Poor, 2008).

These CBAs suggest it would be possible to establish a strong CBA case for public-facing bespoke chatbots that provide a similar range of services to Citizens Advice (which is limited to information and advice). As *Maya* provides an even wider range of services the case should be even stronger. Once the CBA case is established in principle for public-facing bespoke chatbots, the next question is how cost-effective they are compared to other ways of delivering these services. Directly assessing cost-effectiveness is not straightforward, as public-facing chatbots are not directly comparable with other existing services. One challenge is that published service data for front-line legal advice providers is based on cases, whereas for bespoke chatbots it is based on the number of conversations (a case could involve several conversations). A greater challenge is the extent of the support offered. In-person advisers may provide personal support to clients through court or tribunal processes. The unit costs of such cases are much higher than the costs of a conversation with a chatbot. A more relevant comparison would be the costs of a telephone service, such as British Columbia's *Legal Aid FamilyLine*. Here unit costs are $118 per client.[26] Although

[26] Authors' calculation based on British Columbia Legal Aid report 2025–26 figures of spend of CAD 440,844 and 2,792 unique clients supported.

Beagle+ only provides information whereas the telephone service offers advice, it is striking that FamilyLine costs are 17 times higher than *Beagle+*'s current costs of $6.80 a conversation, and almost 100 times higher than its projected costs of $1.20 a conversation once scaled up.

Another approach to estimating cost-effectiveness is to compare a bespoke AI chatbot's running costs with two stylised models. The first assumes a conversation is the equivalent of a 20-minute consultation with a lawyer providing information. The second assumes a conversation would also cover legal advice and assistance (e.g. with court documentation). This is assumed to require two hours of the lawyer's time.[27] Table 1 compares costs, using salary rates for a Canadian lawyer with one year's experience working at a small law firm.[28] The costs for a bespoke AI chatbot are based on *Beagle+* estimates for current conversations ('small scale') and its projected costs when scaled up ('full scale') (see Box 1).

Table 1 Stylised cost effectiveness model for bespoke AI chatbot

Cost of lawyer per conversation	Chatbot cost per conversation	Cost of lawyer relative to chatbot
20 mins = $17.60	Small scale $6.80	3 times higher
20 mins = $17.60	Full scale $1.20	16 times higher
2 hours = $105.70	Small scale $6.80	16 times higher
2 hours = $105.70	Full scale $1.20	88 times higher

Source: Authors' calculations based on *Beagle+* estimates for current conversations ('small scale') and its projected costs when scaled up ('full scale') (see Box 1). Costings for *Maya* are not available.

Table 1 suggests that a bespoke AI chatbot would be cost-effective. Even operating at a small scale, the costs could be a third of those of a lawyer. At full scale, it could be 16 times cheaper. The greatest cost-effectiveness gains would come from a bespoke chatbot that was scaled up and provided the full scope of assistance, including drafting court documents. In that case the bespoke chatbot could be 88 times cheaper than the lawyer.

Table 1 is a simple initial exploration of the cost-effectiveness of bespoke AI chatbots. It assumes that the chatbot is as accurate as the lawyer. It also assumes that the client does not need human interaction with a lawyer. Some clients clearly will. Finally, Table 1

[27] Stylised time estimates based on one of the authors' experience of advising at a free legal aid clinic in London.

[28] Authors' calculation based on CAD 90,000 annual salary for a one-year-qualified lawyer in a small firm in Ontario (as figures not available for British Columbia) (https://www.zsa.ca/salary-guide) plus a 30% uplift to allow for non-salary costs. A fuller comparison would look at a wider range of firms and provinces.

understates the wider benefits of a chatbot – accessible at clients' convenience, including any time of day or week.

The $1.20 cost per conversation figure is of course only indicative. It is a projection based on one specific model in one country. The cost of alternative models, for example ones that for confidentiality reasons do not involve the use of an LLM, will be different. It is also possible that the cost of accessing general-purpose LLMs may rise. However, as this access accounts for only a small proportion of ongoing costs,[29] such changes are unlikely to significantly reduce the scale of the cost benefits noted above. Finally, as two-thirds of ongoing costs are for lawyer oversight (see Box 1), it is likely that unit costs in lower-income countries would be lower than $1.20.[30]

While traditional legal advice providers will always be needed for some clients, especially the most vulnerable and digitally excluded, the analysis suggests that bespoke chatbots would enable some groups to be effectively supported much more cheaply than current approaches – and hence much more affordably.

4.3 Addressing the barriers

4.3.1 Money: financing barriers

High-income countries can afford to invest in front-line justice AI, but face practical challenges

The cost of bespoke front-line justice AI tools can be low. The *Beagle+* costing (see Box 1) demonstrates how affordable developing and running a bespoke legal RAG chatbot can be (in a context where the work on curating legal texts and information relating to everyday problems has already been done – as is the case in many countries). Using England and Wales as an example, a similar chatbot, scaled up for population size, could be built for $1.4 million and run for around $6 million a year – 0.2% of the costs of providing legal aid (of $3.2 billion).[31] While such costs are vanishingly small, there are four reasons why funding remains a challenge in practice.

First, there is a **longstanding and widespread failure to prioritise front-line justice services** (JAC, 2025; Manuel and Manea, 2025). Instead, the focus of justice ministries and judiciaries is on courts and formal justice. This is reflected in spending patterns. Few countries meet even the small recommended share of total justice spending allocation for front-line services (2.5% of total spend). This failure

[29] Two-thirds of *Beagle+*'s ongoing costs are projected to be for lawyers' time to check responses.

[30] As most of the $1.20 unit cost for *Beagle+* is oversight by Canadian lawyers, the equivalent in a lower-income country, where salaries are much less, might only be $0.40. This is based on comparative pay for a lawyer with one year's experience of $1,200 in Lagos (https://www.profolio.ng/lawyer/salary-guide), against $60,000 equivalent in Ontario.

[31] Authors' calculations based on a net operating cost of £2.40 billion https://www.gov.uk/government/publications/legal-aid-agency-annual-report-and-accounts-2024-to-2025/annual-report-and-accounts-2024-25-financial-statements-html

persists despite growing evidence of the impact of such investment – for example, that free specialist advice could save the UK up to $15 billion over 10 years, and that there is a link between unmet legal need, poverty and poor health (Curran et al., 2025). Uraz (2026a) provides granular evidence showing how reductions in legal aid are associated with rising mortality rates.

Second, and this is a common problem in many sectors, **'budget silos' result in funding for justice being split across multiple agencies**. It can be hard to coordinate across silos, so cross-cutting activities often end up underfunded. The UK is an example of this: the Ministry of Justice funds legal aid, local governments fund law centres and Citizens Advice offices are funded by a combination of the Department for Business, Innovation, Science and Trade, the Department for Work and Pensions and local governments. Only a limited degree of coordination would be needed to fund a low-cost *Beagle+*-style chatbot. Simple budget mapping[32] might be sufficient to prompt the necessary coordination, but if needed there are many tried and tested mechanisms for more formal coordination and management of funds for shared objectives (e.g. in the UK, Sure Start funding, the Shared Outcomes Fund and the Better Care Fund).[33]

Third, in many countries there is an **incorrect perception that national civil legal aid systems are unaffordable**. For example, plans in the UK for national legal aid provision in the 1940s (in parallel with the establishment of the national health service) were initially well funded, but civil legal aid has subsequently been gradually eroded and is now widely seen as unaffordable (Uraz, 2026b). The efficacy of adviser-facing AI tools in making a significant impact on resolving people's everyday justice problems depends on a nationwide network of front-line justice advisers. This is lacking in many countries. ODI Global's evidence is that cost-effective and affordable national systems providing legal advice, assistance and informal dispute resolution are possible, provided they include limited service provision by lawyers (Manuel et al., 2025).

Fourth, **front-line justice providers tend to be small non-profit organisations with little or no funding available for capital investment**. As a result, they find it difficult to meet the cost of purchasing AI tools and solutions (HiiL, 2025).

None of these is an insuperable barrier to funding AI chatbots. ODI Global has shown how other sectors, particularly health, have taken on board evidence on the effectiveness of rebalancing their budgets and have adjusted their spending towards front-line services (Hart et al., 2025).

[32] Showing current funding flow to front-line justice services.

[33] Respectively, for childcare, health and family support for early years; cross-departmental working; and integrated health and social care.

Funding needs
Chatbots for everyday legal needs are unlikely to be developed by the private sector (as typical claim values are too low to allow for a commercial return). As discussed above, development and running costs, while low, are still too high for a single not-for-profit front-line legal advice organisation. Costs and risks would be lower if a single chatbot could be used by multiple organisations. Financing constraints could be overcome through a minuscule reprioritisation of government spending, or funding from philanthropic organisations, many of which are looking for opportunities to support AI uses.

Lower-income countries would need additional external funding
Bespoke legal chatbots have the potential to transform access to justice for millions of people in lower-income countries, including by informing and making more transparent the operation of customary and informal justice systems. But the relatively small development costs of a public-facing chatbot represent a large part of the limited budget for capital investment in the justice sector in lower-income countries.

The foundations for bespoke legal chatbots in lower-income countries are already in place. There are examples of multilingual chatbots operating over existing telecommunication systems (Interactive Voice Response, SMS and WhatsApp) as well as online. Many countries have datasets covering everyday legal needs, often developed by legal aid boards, universities or civil society organisations.

As the chatbot infrastructure is essentially the same, there may be scope for a common, open-source approach, updated as general-purpose LLM capability improves. Shared components would bring the cost down, but consideration would need to be given to what parts of the architecture can be common and what must be customised at a national or sub-national level, including to address concerns about sensitive data flowing across borders. This could be funded by a collective effort of governments and private philanthropic donors, in consultation with appropriate regional organisations (including civil society) to ensure the design captures what is needed and is relevant in multiple country contexts. While aid from OECD countries for justice projects is falling, philanthropic grants are rising (although only 3% is provided to low-income countries) (Manuel and Sisodia, forthcoming). The amounts required are small. Based on *Beagle+* and *Hadin kai-bani fondo* costings, developing national chatbots in all 41 lower-income countries in sub-Saharan Africa would cost $5 million. This assumes they are developed independently. If part of the development were shared, the costs would be even lower.

4.3.2 Monopolies: regulatory barriers and unauthorised practice of law

General-purpose LLMs provide a wide range of legal services including advice and drafting legal and court documents – all directly to the public. To avoid claims for losses incurred because of incorrect advice, and unauthorised practice of law challenges, developers rely on disclaimers that a general-purpose LLM is not a lawyer, and that the advice it provides could be incorrect. These approaches are currently being contested in the US courts and potentially addressed by US legislators.

Developers of bespoke legal chatbots (usually front-line legal advice organisations) are more cautious about unauthorised practice of law concerns. In some jurisdictions (Canada, India and some US states) lawyers have a strict exclusive monopoly on providing legal advice (see Nori (2026c) on India, and Edwards and Cram (2026: 33) on the US). As noted, unauthorised legal practice considerations limit the scope of services offered by *Beagle+* to legal information, while South Africa's more permissive regulatory regime enables *Maya* to provide a full range of services.

Some commentators see regulatory regimes that inhibit the adoption of public-facing AI as a 'protectionist trap' (Nori, 2026a). This echoes Susskind's longstanding argument that professional resistance to technology can put lawyers' interests ahead of meeting people's legal needs (Susskind, 2023). Regulatory sandboxes – supervised spaces for testing new products and services – provide an opportunity to engage with the regulatory barriers inhibiting front-line justice services. In Utah, this approach has enabled community justice workers to provide front-line legal advice and assistance, waiving unauthorised practice of law regulations.[34] The UK government's Advisory AI Growth Lab (an AI sandbox) has selected legal services as its first sector focus, with potential applications for legal aid charities and advice bodies concerned with public access to justice.

Bespoke legal chatbots, both public- and adviser-facing, offer the potential to transform the scope, reach, cost-effectiveness and affordability of front-line advice and assistance for everyday justice problems. But this potential is only likely to be fully realised in a more permissive regulatory environment. Such an environment would enable non-profit free legal advice organisations to offer the full scope of advice and assistance for everyday justice problems and, with appropriate safeguards (including professional indemnity insurance), to do so cost-effectively by deploying non-lawyers, such as community justice workers and community paralegals (Manuel et al., 2023; Manuel and Manuel, 2024). Greater regulatory freedom could also give AI developers the confidence to develop and deploy

[34] The judiciaries of six other states – Alaska, Arizona, Delaware, Hawaii, Texas and South Carolina – and Washington DC have also authorised community justice workers, in some cases giving them rights of audience (see Teufel et al., 2026). However, as press reports in Arizona illustrate, it is much more challenging when deregulating to allow for-profit organisations to operate (Gersony, 2026).

bespoke public-facing AI legal chatbots with the range of legal advice and assistance currently available from general-purpose LLMs.

4.3.3 Mindsets: professional reluctance

Private law firms are embracing AI for productivity gains, but legal aid and public interest lawyers can be hesitant. Some parts of the legal advice sector are enthused about AI's potential to address the justice gap and help under- and unserved communities (NYSBA, 2026; Ogoniba, 2023). Others urge caution, particularly for vulnerable clients, and stress the need for human advisers to exercise judgment, show empathy and manage uncertainty. An examination of AI's potential to help Community Legal Centres in Australia concluded that, while well-designed AI tools can augment physical 'shop fronts', they cannot replace them (Hastings et al., 2026). In a July 2026 survey of public interest lawyers in the US, three-quarters agreed at least somewhat with the statement 'I am pretty unfriendly to AI' (NITA, 2026, unpublished data). However, another US study suggests that legal aid organisations are embracing AI at twice the rate of other lawyers (Ambrogi, 2025).

Practical, sector-specific training, delivered through trusted peer networks, could build confidence: two-thirds of legal aid practitioners surveyed by the New York State Bar Association wanted AI training (NYSBA, 2026), and evidence from other public services shows that guidance on using AI tools increases uptake (World Bank, 2026a). Clear guidance from regulators and funders could help reduce concerns about confidentiality and client safety. Above all, AI should be framed as a way to extend human advice services rather than replace them (Hastings et al., 2026), with advisers involved in designing the tools and independent evaluations building the evidence on what works.

4.4 Conclusion

The costs of developing and running accurate, public-facing bespoke chatbots dealing with everyday justice problems are remarkably low, and substantially lower than human-delivered information. There is also evidence that AI chatbots improve the efficiency of front-line justice advisers. The barriers to realising their potential are real, but none is insuperable:

Money: ODI Global estimates that a public-facing legal chatbot for England and Wales could be built for $1.4 million and run for around $6 million a year – 0.2% of the costs of providing legal aid. But the cost of purchasing AI tools and running them can be prohibitive for small, non-profit front-line justice providers.

The challenge is a longstanding and widespread failure to prioritise front-line justice services.

Monopolies (regulatory barriers): Regulatory risks and the cost of litigation for any errors are inhibitors. Relaxing the monopoly on legal services for everyday legal problems handled by non-profit free legal advice organisations and non-lawyer advisers, coupled with professional indemnity insurance, would go a long way to addressing these issues. The initial step could be to secure ‘sandbox’ treatment for chatbot development.

Mindsets: Overcoming reluctance on the part of legal aid lawyers about AI adoption would involve taking the principled concerns behind that reluctance seriously, including involving advisers in designing tools, and framing AI as a way to extend human advice services rather than replace them.

Our analysis rests on one costed public-facing chatbot and one adviser-facing randomised trial; more organisations will need to follow the People’s Law School’s example and open their books before realistic benchmark costs can be established. There is no cost data at all on AI-assisted informal dispute resolution, and a CBA measuring outcomes rather than outputs has yet to be attempted for any bespoke chatbot. Independent evaluation, with advisers involved in design, will be essential to build the evidence and the confidence of the sector.

Bespoke AI chatbots now look like a low-cost, low-risk, high-return investment. Proof of concept has been established, curated legal datasets exist in many countries, and the costs are small. A national scale-up of a bespoke public-facing legal chatbot could transform access to justice for millions. For those needing face-to-face support, such a chatbot would enable human front-line advisers to work more efficiently, reaching more people for the same budget.

The constraints on AI-assisted front-line justice are choices about finance, regulation and professional culture.

5 Conclusions and recommendations

5.1 Conclusions

Since the landmark 2019 Justice for All Report (Task Force on Justice, 2019), there has been increased focus on closing the justice gap and investing in front-line justice services, as well as the costs of failing to do so. Despite this, the focus of ministries of justice and judiciaries remains on courts and formal justice.

Formal justice providers – law firms and the judiciary – are making extensive use of AI. Current AI deployment across the justice system is likely to increase inequality: benefiting large repeat players with bulk claims and those able to pay for legal services. Lower-income countries face significant challenges: reliance on AI that is ill-suited to the local context could do real harm due to inaccuracy and bias. There has been little focus to date on the potential for AI to enhance access to front-line justice and help close the justice gap.

However, there is now a new dynamic. For the first time, general-purpose LLMs are offering millions the prospect of personalised, tailored legal information, advice and assistance. Concerns that a flood of AI-enabled claims could start hitting the courts could potentially jolt policymakers into reallocating funding towards the front line. With marginal reallocation of existing resources to front-line justice, AI offers two opportunities to reduce the justice gap.

The first (and the focus of this paper) is to develop bespoke legal chatbots providing public-facing legal information, advice and assistance, backed by human advisers to assist clients with vulnerabilities or who are digitally excluded. Well-designed RAG chatbots draw on the capability of general-purpose LLMs but are then limited to working with curated local legal material. Run by independent accountable organisations, results are tested and overseen by lawyers. Risks around confidentiality, sycophancy and bias cannot be eliminated. But *Beagle+* and *Maya* provide proof of concept: it is possible to implement a well-designed bespoke AI chatbot that is superior to a general-purpose LLM.

Not only can accurate bespoke AI chatbots be developed: the key new piece of evidence is that their costs are low, including the costs of continued human oversight to ensure accurate and helpful answers. Further adoption and scaling up of chatbots is held back by

lack of public funding for the small non-profit organisations that are developing them. Their potential to assist the public is also constrained by regulations. Both funding and regulatory barriers could be overcome. Running costs are affordable (possibly less than 0.4% of the costs of legal aid), and deregulation to allow non-profits to provide legal assistance for everyday legal problems – cases that are too small and uneconomical for a lawyer – would have minimal impact on the legal profession. While governments could host such chatbots, this is better done by independent organisations, as some legal needs involve challenging government decisions e.g. on benefit payments.

The second opportunity is to use AI to scale up informal, out-of-court dispute resolution. While the private sector has demonstrated the efficacy of non-court ODR to resolve small-scale disputes without human intervention, ODR uptake in the public sector is low. The pattern is repeated with AI. We found very few examples of AI-assisted informal dispute resolution, and no public sector examples of AI making determinations in disputes. How far AI should go in deciding disputes is a democratic question of public acceptability. But the public use of general-purpose AI tools is likely to increase demands on the courts, and the potential surge in claims suggests that testing these approaches is urgent.

A key overall conclusion is that the use of AI in front-line justice is still at an experimental stage. Investment should therefore leave room for failure, learning and change, with lessons openly shared.

5.2 Recommendations

Realising AI's potential at the justice front line will require a collective effort. Governments and regulators, philanthropic funders, and front-line advice organisations each have a part to play in tackling the barriers of money, monopolies and mindsets, and in building the evidence on what works.

For governments, especially ministries of justice and finance, and regulators

- **Facilitate sharing of local knowledge**, often held by non-profits, to enable the creation of national services that can draw on the necessary detailed understanding of local laws and regulations.
- **Fund the development of one AI chatbot** built on legal material pooled from advice organisations and hosted and maintained by a trusted non-profit organisation seen as independent of government. Use budget mapping, and pooled funds where needed, to overcome departmental silos.
- **Create a clear route for accountable non-profit chatbots.** Allow non-profit advice organisations, their chatbots and non-

lawyer advisers to give the full range of advice and assistance for everyday legal problems. Use regulatory sandboxes to test front-line AI justice services, prioritising applications from advice charities and legal aid providers.

- **Track the quality of public free legal advice** from existing providers and general-purpose and bespoke AI tools through a continuous benchmarking process.
- **Develop standards for public-facing legal bespoke chatbots**, covering accuracy testing, human oversight, confidentiality, transparency and redress. The characteristics set out in Appendix 2 could provide a starting point.
- **Invest in research, monitoring and evaluation** on AI's effectiveness and impact on front-line justice. Fund up-to-date legal needs surveys that capture AI use to help plan for the potential surge in claims.
- **Accelerate exploration of innovative use of AI in pre-court dispute resolution mechanisms** (and initiate consultations on public acceptability).

For philanthropic funders and donors

In lower-income countries:

- **Pay the up-front development costs** of bespoke public-facing chatbots. These costs are small but are still beyond the means of most advice organisations.
- **Fund the foundations for AI tools.** Support governments and civil society to build curated legal knowledge bases, including on customary law, in local languages.
- **Back shared, open-source approaches.** Fund the development of a common chatbot model that can be adapted across countries. The cost for all lower-income countries in sub-Saharan Africa is an estimated $5 million or less.

For front-line advice organisations

- **Collaborate rather than build alone.** Pool curated legal materials so that one well-funded tool can serve many organisations.

For all funders

- **Fund research to identify the key features of an effective chatbot** and the most cost-effective way to deliver it.

- **Make evidence a condition of funding.** Require published evaluations, including unit costs, accuracy and user outcomes, and pay for independent audits.
- **Fund evaluation as structured experimentation**, with clear metrics, realistic counterfactuals (often no service at all, rather than a perfect lawyer) and public reporting of failures as well as successes.

References

Ambrogi, R. (2025) *Legal Aid Organizations Embrace AI at Twice the Rate of Other Lawyers, New Study Reveals* LawSites, 25 September (www.lawnext.com/2025/09/legal-aid-organizations-embrace-ai-at-twice-the-rate-of-other-lawyers-new-study-reveals.html).

Amnesty International (2026) *Unlawful by design: exposing the human rights costs of generative AI.* Index POL 40/0996/2026. London: Amnesty International (www.amnesty.org/en/documents/pol40/0996/2026/en/).

Anton, G. (2026) 'Inside Project Lily: OpenAI contractors are reading ChatGPT users' private prompts, new report finds' *International Business Times Australia* 16 September (www.ibtimes.com.au/openai-contractors-review-chatgpt-conversations-1875634).

Arzilli, F., Lynch, C.S. and Page, L. (2026) *An Evaluation of DWP's Microsoft 365 Copilot Trial.* Department for Work and Pensions. Published 29 January 2026. Available at gov.uk.

Barocas, S., Hardt, M. and Narayanan, A. (2023) *Fairness and machine learning: limitations and opportunities.* Cambridge, MA: MIT Press (https://fairmlbook.org/).

Briggs, M. (2026) 'AI and civil justice: preparing for the tsunami'. Speech to the Civil Justice Systems in the 21st Century Conference. Oxford, 20 May (https://supremecourt.uk/speeches/speech-lord-briggs-200526).

Brun, V., Lenz, R. and Díaz Pérez, S. (2025) 'A voice-first AI service for people-centred justice in Niger'. Paper presented by HiiL team at the Workshop on AI for Access to Justice, Dispute Resolution, and Data Access (AIDA2J), 38th International Conference on Legal Knowledge and Information Systems (JURIX 2025), Turin, 9–11 December (aida2j.com).

Brynjolfsson, E., Li, D. and Raymond, L.R. (2023) *Generative AI at work.* NBER Working Paper 31161. Cambridge, MA: National Bureau of Economic Research (https://doi.org/10.3386/w31161)

Cashman, P. and Ginnivan, E. (2019) 'Digital justice: online resolution of minor civil disputes and the use of digital technology in complex litigation and class actions' Macquarie

Law Journal 19:39–79 (www.mq.edu.au/__data/assets/pdf_file/0012/866289/Digital-Justice.pdf).

CEPEJ – European Commission for the Efficiency of Justice (2024) *European judicial systems – CEPEJ evaluation report – 2024 evaluation cycle (2022 data).* Strasbourg: Council of Europe (www.coe.int/en/web/cepej/-/the-cepej-publishes-its-new-evaluation-report-on-european-judicial-systems).

CEPEJ (2025) *1st AIAB report on the use of artificial intelligence in the judiciary.* Strasbourg: Council of Europe (https://rm.coe.int/1680b49def).

Challapally, A., Pease, C., Raskar, R. and Chari, P. (2025) 'The GenAI Divide: State of AI in Business 2025' MIT NANDA (https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf).

Civil Resolution Tribunal (2025) 2024/2025 annual report: April 1, 2024 to March 31, 2025. Victoria, BC: Civil Resolution Tribunal (https://civilresolutionbc.ca/wp-content/uploads/CRT-Annual-Report-2024-2025.pdf).

Cobbe, J., Veale, M. and Singh, J. (2023) 'Understanding accountability in algorithmic supply chains' in *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency (FAccT '23).* New York: ACM (https://dl.acm.org/doi/10.1145/3593013.3594073).

Coglianese, C. (2024) 'Procurement and artificial intelligence' in *Handbook on public policy and artificial intelligence.* Cheltenham: Edward Elgar Publishing (www.elgaronline.com/edcollchap/book/9781803922171/book-part-9781803922171-26.xml).

Commission on Legal Empowerment of the Poor (2008) *Making the law work for everyone, volume I: report of the Commission on Legal Empowerment of the Poor*. New York: Commission on Legal Empowerment of the Poor and UNDP.

Cortes, P. (ed.) (2018a) *The New Regulatory Framework for Consumer Dispute Resolution*. Oxford: Oxford University Press.

Cortes, P. (2018b) 'Using technology and ADR methods to enhance access to justice' *International Journal of Online Dispute Resolution* 5(1–2): 103–121 (https://doi.org/10.5553/IJODR/235250022018005102011).

Curran, L., Ching, J. and Jarman, J. (2025) *Regulatory leadership on access to justice: report to the Legal Services Consumer Panel on unmet legal need – barriers and creative solutions to access to justice.* Nottingham: Nottingham Law School, Nottingham Trent University, and London: Legal Services Consumer Panel (www.legalservicesconsumerpanel.org.uk/wp-

content/uploads/2025/04/25.04.01-Regulatory-Leadership-on-Access-to-Justice-Report.pdf).

DCAF – Geneva Centre for Security Sector Governance (2026) *Artificial intelligence and access to justice in fragile and conflict-affected situations*. Geneva: DCAF (www.dcaf.ch/artificial-intelligence-and-access-justice-fragile-and-conflict-affected-situations).

Deloitte (2026) *The AI imperative: reshaping of the legal industry.* AI Survey Publication. Deloitte in partnership with RSGI. (https://www.deloitte.com/uk/en/services/legal/research/ai-imperative-reshaping-of-the-legal-industry.html)

Deshmukh, S. and Davies, M. (2026) Measuring up: evaluating claims about AI and productivity in the UK public sector. Policy briefing. London: Ada Lovelace Institute (www.adalovelaceinstitute.org/policy-briefing/measuring-up/).

Dylag, M. (2026) 'Is it competent AI-generated legal advice?'. Presentation to the panel discussion '*Beyond the algorithm: people-centred design in an AI age*'. Unpublished (https://justicedevelopmentgoals.squarespace.com/spring-2026-conference-1).

Edwards, N. and Cram, S. (2026) *A people-centered justice approach to implementing AI governance: how the justice sector can support inclusive, equitable, and enforceable AI governance*. New York: NYU Center on International Cooperation (https://cic.nyu.edu/wp-content/uploads/2026/02/A-People-Centered-Justice-Approach-to-Implementing-AI-Governance-March-2026.pdf).

The Economist (2026a) 'How AI is breaking the British state'. *The Economist*, 6 August (www.economist.com/leaders/2026/08/06/how-ai-is-breaking-the-british-state).

The Economist (2026b) 'The tragedy of the commons, AI edition'. *The Economist*, 6 August (www.economist.com/britain/2026/08/06/the-tragedy-of-the-commons-ai-edition).

Enock, F.E. and Margetts, H.Z. (2026) *Large language models in the UK: public use, trust, and attitudes.* Oxford: Oxford Internet Institute, University of Oxford (www.oii.ox.ac.uk/wp-content/uploads/2026/07/Large-language-models-in-the-UK-public-use-trust-and-attitudes_FINAL.pdf).

Faghih, K., Cheng, Y., Saha, S. et al. (2026) 'Same question, different answers: evaluating LLM reliability beyond accuracy'. College Park, MD: University of Maryland (https://arxiv.org/abs/2607.22554).

Garfield AI (2025) ‘Ministry of Justice features Garfield AI as world’s first AI-driven law firm in AI Action Plan’ (https://www.garfield.law/press/garfield-ai-featured-in-moj-legal-technology-report).

Garg, A. and Ma, M. (2025) *Opportunities and challenges in legal AI*. White paper. Stanford, CA: Stanford Law School (https://law.stanford.edu/wp-content/uploads/2025/01/Legal-AI-Opportunities-and-Challenges-Whitepaper.pdf).

Gersony, L. (2026) ‘Arizona lets investors own law firms. Consumers pay the price’. The Arizona Republic, 8 February (https://www.azcentral.com/story/news/local/arizona-investigations/2026/02/08/arizona-alternative-business-structures-program-draws-consumer-complaints/87214158007/).

Gramatikov, M., Díaz Pérez, S. and Taylor, R. (2026) *Measuring justice from the users’ perspective: a first look at the Justice Services Measure.* The Hague: HiiL.

Hagan, M. (2025) ‘AI+A2J 2025 Summit takeaways’. *Justice Innovation* blog, 17 December. Stanford, CA: Stanford Legal Design Lab (https://justiceinnovation.law.stanford.edu/aia2j-2025-summit-takeaways/).

Hart, T., Manuel, M. and Manuel, C. (2025) *More money for justice and more justice for the money: lessons from the health sector for financing front-line justice services.* ODI Global Working Paper. London: ODI Global (https://odi.org/en/publications/more-money-for-justice-and-more-justice-for-the-money-lessons-from-the-health-sector/).

Hastings, C., Cotterell, A. and Bruce, F. (2026) ‘Artificial intelligence and access to justice at the “shop front”: the potential and limitations of meeting legal need through technology’ *Australian Journal of Social Issues* 1–14 (https://doi.org/10.1002/ajs4.70115).

Hemrajani, R. (2025) ‘Evaluating the role of large language models in legal practice in India’ (https://arxiv.org/abs/2508.09713).

HiiL (2017) *Rechtwijzer: why online supported dispute resolution is hard to implement*. The Hague: HiiL (https://www.hiil.org/news/rechtwijzer-why-online-supported-dispute-resolution-is-hard-to-implement/)

HiiL (2025) *AI and access to justice: a snapshot of current trends and future potential.* The Hague: Hague Institute for Innovation of Law (www.hiil.org/?p=25138).

Hollenbeck, S. (2026) ‘Survey: 65% use AI legal advice, but accuracy concerns remain’. Blog (www.rev.com/blog/ai-legal-advice-index)

IDLO (2026) *Statement to the United Nations Global Dialogue on Artificial Intelligence Governance delivered by M. Cassayre*. Geneva, 7 July (https://idlo.int/news/policy-statements/united-nations-global-dialogue-artificial-intelligence-governance)

Juvekar, K., Bhattacharya, A., Khadloya, S. and Saxena, U. (2025) 'Are LLMs court-ready? Evaluating frontier models on Indian legal reasoning' *Proceedings of the Natural Legal Language Processing Workshop 2025*, Suzhou, China. Association for Computational Linguistics (https://aclanthology.org/2025.nllp-1.26.pdf).

JUSTICE (2026) 'What AI chatbots can teach us about unmet legal needs'. London: JUSTICE (https://justice.org.uk/reports/what-ai-chatbots-can-teach-us-about-unmet-legal-needs).

Justice Action Coalition (2025) Justice Financing Framework. November (www.sdg16.plus/justice-financing-framework).

Katsh, E. and Rule, C. (2016) 'What we know and need to know about online dispute resolution' *South Carolina Law Review* 67(2): 329–344 (https://scholarcommons.sc.edu/sclr/vol67/iss2/10).

Laskowski, C., Griffin, C.L. Jr and Thumma, S.A. (2023) *Generative artificial intelligence and access to justice: possibilities, concerns, best practices, and how to measure success.* Arizona Summit on Artificial Intelligence Law and the Courts, 1 December. Williamsburg, VA: National Association for Court Management (https://nacmnet.org/wp-content/uploads/AI-and-Access-to-Justice-Final-White-Paper.pdf).

LawWorks (2026) 'Template clinic AI policy'. London: LawWorks (www.lawworks.org.uk/solicitors-and-volunteers/resources/template-clinic-ai-policy).

Legal Aid BC (2026) *2025/26 Annual Service Plan Report*. Vancouver, BC: Legal Aid BC (https://www.bcbudget.gov.bc.ca/Annual_Reports/2025_2026/pdf/agency/labc.pdf).

Legal Futures (2017) 'Pioneering ODR platform to rein in ambitions after commercial setback', *Legal Futures* 3 April. (https://www.legalfutures.co.uk/latest-news/pioneering-odr-platform-to-rein-in-ambitions-after-commercial-setback).

Legal Services Board and The Law Society (2024) *2023 Individual Legal Needs Survey: Exploring Unmet Legal Needs*. London: Legal Services Board (https://legalservicesboard.org.uk/wp-content/uploads/2025/01/ILNS-unmet-needs-topic-report-FINAL.pdf).

Legal Services Board (2026) 'AI tools show real promise to increase access to legal services, but consumers lack basic protections'. News release, 11 June. London: LSB

(legalservicesboard.org.uk/news/ai-tools-show-real-promise-to-increase-access-to-legal-service).

Legal Services Corporation (2022) *The justice gap: the unmet civil legal needs of low-income Americans*. Washington DC: LSC (justicegap.lsc.gov/the-report/).

Legal Services Research Centre (2026) *Generative AI use in the legal profession: findings from the 2025 Victorian Lawyer Census.* Melbourne: Victorian Legal Services Board and Commissioner (https://lsbc.vic.gov.au/research-centre/what-we-do/publications/generative-ai-use-legal-profession-findings-2025-victorian).

Levy, K., Chasalow, K.E. and Riley, S. (2021) 'Algorithms and decision-making in the public sector' *Annual Review of Law and Social Science* 17: 309–334 (https://doi.org/10.1146/annurev-lawsocsci-041221-023808).

LexisNexis (2026) 'Lawyer preference for AI grounded in legal sources rises to 81%'. Press release (www.lexisnexis.com/community/pressroom/b/news/posts/lawyer-preference-for-ai-grounded-in-legal-sources-rises-to-81).

Linklaters (2025) 'UK – The LinksAI English Law Benchmark (Version 2). DigiLinks blog. 10 February (https://www.linklaters.com/en/insights/blogs/digilinks/2025/february/uk-the-linksai-english-law-benchmark-version-2 accessed 19 September 2026).

Manuel, C. and Manuel, M. (2024) 'Cost-effective front-line justice services in Sierra Leone: a case study in frugal innovation and domestic resourcing'. London: ODI (https://odi.org/en/publications/cost-effective-front-line-justice-services-in-sierra-leone-a-case-study-in-frugal-innovation-and-domestic-resourcing/).

Manuel, C. and Manuel, M. (2026) 'Financing scaled-up front-line justice services: data and evidence gaps' in M. Burnett and R.L. Sandefur (eds) *Global perspectives on people-centered justice: exploring the evidence.* Chicago, IL: American Bar Foundation (https://url.uk.m.mimecastprotect.com/s/_ivDCE92zcW8RRvcNfxU7wgSF?domain=americanbarfoundation.org/).

Manuel, C., Manuel, M. and Sharp, S. (2025) 'Front-line justice services with the potential to scale up: evidence from LMICs'. London: ODI (https://odi.org/en/publications/front-line-justice-services-with-the-potential-to-scale-up-evidence-from-lmics/).

Manuel, M. and Manea, S. (2025) 'State of justice financing 2025 annual review: domestic financing and aid'. ODI Policy Brief. London: ODI (https://media.odi.org/documents/Justice_Financing_Annual_Review_2025_Sep25_rev081225.pdf).

Manuel, M. and Manuel, C. (2021) ‘People-centred justice for all – a route to scaling up access to justice advice and assistance in low-income countries’. ODI Policy Brief. London: ODI (https://odi.org/en/publications/people-centred-justice-for-all/)

Manuel, M. and Manuel, C. (2023) ‘“Small is beautiful, but scale is necessary”: front-line justice services in lower-income countries with the potential to scale up’. ODI Policy Brief. London: ODI (https://odi.org/en/publications/small-is-beautiful-but-scale-is-necessary-front-line-justice-services-in-lower-income-countries-with-the-potential-to-scale-up/).

Manuel, M. and Msiska, C. with Ndalahoma, C. and Nanthambwe, H. (2023) ‘Cost-effective front-line justice services in Malawi: a case study in frugal innovation’. London: ODI (https://odi.org/en/publications/malawi-case-study/).

Manuel, M. and Sisodia, A. (forthcoming) ‘State of justice aid 2026 annual review’. ODI Policy Brief. London: ODI.

Manuel, M. and Uraz, J. (forthcoming) ‘State of justice finance 2026 annual review’. ODI Policy Brief. London: ODI.

MEL Research (2026) *AI in legal services*. Report for the Legal Services Board. London: LSB (available via legalservicesboard.org.uk/news/ai-tools-show-real-promise-to-increase-access-to-legal-service).

McGrath, P. and Baidya, A. (2026) *Public institutions in the age of AI: emerging practices.* Prosperity Insight Series. Washington DC: World Bank (https://documents1.worldbank.org/curated/en/099051226113013516/pdf/P502259-818d6459-642d-4c32-829b-7c76a085b561.pdf).

Michalski, J., Winters, S., Gunn, D. and Holland, J. (2025) ‘AI ROI: The paradox of rising investment and elusive returns’. Deloitte Global. 22 October (https://www.deloitte.com/global/en/issues/ai/ai-roi-the-paradox-of-rising-investment-and-elusive-returns.html).

Ministry of Justice (2025) *AI action plan for justice.* London: Ministry of Justice (www.gov.uk/government/publications/ai-action-plan-for-justice/ai-action-plan-for-justice).

Ministry of Justice (2026a) *AI action plan for justice: one year on.* London: Ministry of Justice (https://www.gov.uk/government/publications/ai-action-plan-for-justice-one-year-on).

Necz, D. (2024) ‘Rules over words: regulation of chatbots in the legal market and ethical considerations’ *Hungarian Journal of Legal Studies* 64(3): 472–485 (https://doi.org/10.1556/2052.2023.00472).

NITA – National Institute for Trial Advocacy (2026) AI survey results. Unpublished survey data and analysis workbook.

Nori, S. (2026a) ‘Legal futurisms: the Toronto epiphany – why legal aid must embrace consumer-facing tech’. Legal Futurisms/LawDroid, LinkedIn, 14 July (www.linkedin.com/pulse/legal-futurisms-toronto-epiphanywhy-aid-must-embrace-tech-nori-wzabf/)

Nori, S. (2026b) ‘Legal futurisms: OpenAI sued for unauthorized practice of medicine: key takeaways for access to justice’. Legal Futurisms/LawDroid, LinkedIn, 28 July (www.linkedin.com/pulse/legal-futurisms-openai-sued-unauthorized-practice-medicine-nori-ixwye/).

Nori, S. (2026c) ‘When the obstacle becomes the path: UPL and Indian legal tech’. Legal Futurisms/LawDroid, LinkedIn, 30 June (www.linkedin.com/pulse/when-obstacle-becomes-path-upl-indian-legal-tech-sateesh-nori-hohxc/).

Nuffield Foundation (n.d.) ‘Access to justice through artificial intelligence’. Project page (www.nuffieldfoundation.org/project/access-to-justice-through-artificial-intelligence).

NYSBA – New York State Bar Association (2026) *Report and recommendations on artificial intelligence and access to justice in 2025.* Approved by the House of Delegates, 16 January. Albany, NY: New York State Bar Association, President’s Committee on Access to Justice (https://nysba.org/wp-content/uploads/2023/02/Approved-Report-and-Recommendations-on-AI-and-Access-to-Justice-in-2025.pdf).

OECD (2023) *Recommendation of the Council on Access to Justice and People-Centred Justice Systems*. OECD/LEGAL/0498. Paris: OECD (https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0498)

OECD (2025a) *Toolkit for access to justice and people-centred justice systems*. Paris: OECD Publishing (https://www.oecd.org/en/publications/toolkit-for-access-to-justice-and-people-centred-justice-systems_aecf7f78-en.html).

OECD (2025b) *Governing with artificial intelligence: the state of play and way forward in core government functions*’. Paris: OECD Publishing (www.oecd.org/en/publications/governing-with-artificial-intelligence_795de142-en/full-report.html).

OECD (2026a) ‘Generative AI experimentation in government: learning from emerging guidelines’. OECD Working Paper on Public Governance No. 93. Paris: OECD Publishing (https://doi.org/10.1787/42815683-en).

OECD (2026b) *Digital government outlook 2026: from foundations to transformational impact.* Paris: OECD Publishing (https://doi.org/10.1787/0496b2bc-en).

OECD (2026c) 'Towards a measurement framework for people-centred justice systems'. OECD working paper. Paris: OECD Publishing (www.oecd.org/content/dam/oecd/en/publications/reports/2026/06/towards-a-measurement-framework-for-people-centred-justice-systems_b1446e3b/f4c9dccb-en.pdf).

Ogoniba, B.U. (2023) 'AI and access to justice: improving legal services and equality' *Nigerian Bar Journal* 13(1): 54 (www.ajol.info/index.php/nba/article/view/254148).

Okeke, U.C. (2026) 'Enhancing access to justice through technology in developing countries: technology-based initiatives, challenges and emerging technologies' *International Journal for Court Administration* 16(3): 5 (https://doi.org/10.36745/ijca.635).

OpenAI (2025) 'Introducing GPT-5.2'. Product release, 11 December. San Francisco, CA: OpenAI (https://openai.com/index/introducing-gpt-5-2/)

Open Government Partnership (2026) 'Factsheet: justice and artificial intelligence' (www.opengovpartnership.org/stories/justice-and-artificial-intelligence/).

People's Law School (2024) 'Beagle+'. Vancouver, BC: People's Law School (www.peopleslawschool.ca/about/beagle/).

People's Law School (2025a) 'Beagle+ fetches access to justice'. (https://www.peopleslawschool.ca/news/beagle-fetches-access-to-justice/).

People's Law School (2025b) 'How to test a GenAI chatbot', 17 March. Vancouver, BC: People's Law School (www.peopleslawschool.ca/news/testing-genai-chatbot/; underlying data at www.justicebench.org/dataset/beagle-data).

Perlo, J. (2026) 'Alaska's court system built an AI chatbot. It didn't go smoothly.' NBC News, 3 January (https://www.nbcnews.com/tech/tech-news/alaskas-court-system-built-ai-chatbot-didnt-go-smoothly-rcna235985).

Pleasence, P. and Balmer, N.J. (2025) *Report of the pilot Victorian Legal Understanding and Lawyer Use (V-LULU) survey*. Melbourne: Victorian Legal Services Board and Commissioner (https://lsbc.vic.gov.au/node/1139).

Porcaro, K. et al. (2025) Audit report for Legal Information Assistant (LIA). Spring 2025. Durham, NC: Duke Center on Law & Technology (https://scholarship.law.duke.edu/cgi/viewcontent.cgi?article=1001&context=dclt).

Porcaro, K., D'Antuono, C. and Jeter, W. (2026) *Audit report for self-help chatbot.* Spring 2026. Durham, NC: Duke Center on Law & Technology (https://scholarship.law.duke.edu/cgi/viewcontent.cgi?article=1003&context=dclt).

Rajendran, R.K. et al. (2025) 'The role of AI in enhancing access to justice and legal services' in H. Bajraktari (ed.) *Exploration of AI in contemporary legal systems.* Hershey, PA: IGI Global (https://doi.org/10.4018/979-8-3693-7205-0.ch008).

Raworth, H. (2026) 'Legal AI vendors should share the verification burden' *Lawyers Weekly* 12 August (https://www.lawyersweekly.com.au/corporate-counsel/44827-legal-ai-vendors-should-share-the-verification-burden).

Sargeant, H. (2026) *Artificial intelligence & justice: an evidence scoping review.* Public Right to Justice programme. Nuffield Foundation (www.nuffieldfoundation.org/wp-content/uploads/2026/07/AI-and-justice.pdf).

Satterthwaite, M. (2025) *AI in judicial systems: promises and pitfalls. Report of the Special Rapporteur on the independence of judges and lawyers*. UN Doc. A/80/169. New York: United Nations General Assembly (https://www.ohchr.org/en/documents/thematic-reports/a80169-ai-judicial-systems-promises-and-pitfalls-report-special)

Schneiders, E. et al. (2025) 'Objection Overruled! Lay People can Distinguish Large Language Models from Lawyers, but still Favour Advice from an LLM' *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM (https://doi.org/10.1145/3706598.3713470).

Scarcella, M. (2026) 'Google expands Gemini Enterprise AI platform for law firms, lawyers'. *Reuters*, 25 August (https://finance.yahoo.com/technology/ai/articles/google-expands-gemini-ai-platform-120333158.html).

Shukla, A.K. (2025) 'AI for citizens: empowering public institutions worldwide through inclusive, transparent and ethical AI'. SSRN working paper (https://doi.org/10.2139/ssrn.5342156).

Sinha, A. et al. (2026) 'Human-on-the-loop orchestration for AI-assisted legal discovery' Mountain View, CA: Google LLC (https://arxiv.org/abs/2606.19812).

Solicitors Regulation Authority (2026) *Misuse of AI: warning notice*. (https://www.sra.org.uk/solicitors/guidance/misuse-ai/)

Sorabji, J. (2026) Improving civil justice reform: an analysis of major reviews in England and Wales. Full report. Public Right to Justice programme. Nuffield Foundation (www.nuffieldfoundation.org/wp-

content/uploads/2026/06/Improving-civil-justice-reform-major-reviews.pdf).

Stanford HAI (2026) *The AI Index 2026 Annual Report*. Stanford, CA: AI Index Steering Committee, Institute for Human-Centered AI, Stanford University (https://arxiv.org/abs/2606.15708).

Statistics Canada (2022) *Canadian Legal Problems Survey (CLPS), 2021*. Ottawa: Statistics Canada.

Stephens, S. (2025a) 'The development and evaluation of a legal chatbot for women in Tanzania'. 1 April. SSRN working paper (https://ssrn.com/abstract=5251610).

Stephens, S. (2025b) 'Justice in my pocket? How can the use of a justice chatbot which is accessible on a mobile phone enable access to justice for women in Dar es Salaam?' Brighton: University of Sussex (https://hdl.handle.net/10779/uos.30363844.v1).

Stephens, S. and Nchembi, D. (2025) 'Exploring the opportunities and limitations for artificial intelligence to reduce legal aid backlogs in Malawi'. Discovery Report. Frontier Technologies Hub (www.frontiertechhub.org/insights/ai-legal-assist-discovery).

Susskind, R. (2023) Tomorrow's Lawyers: An Introduction to Your Future. 3rd edn. Oxford: Oxford University Press.

Susskind, R. (2026) *The future of law: reflections and predictions.* 2nd edn (30th anniversary edition). Oxford: Oxford University Press.

Tashea, J. (2026) 'A research agenda for justice technology' in M. Burnett and R.L. Sandefur (eds) *Global perspectives on people-centered justice: exploring the evidence.* Chicago, IL: American Bar Foundation (https://www.americanbarfoundation.org/wp-content/uploads/2026/04/Global-Perspectives-on-People-Centered-Justice.pdf).

Task Force on Justice (2019) *Justice for all.* New York: Pathfinders for Peaceful, Just and Inclusive Societies (www.justice.sdg16.plus/report-old2022).

Teufel, J., Burnett, M., and Sandefur, R.L. (2026) 'Research Brief Analysis of the Social and Economic Impact of the Delaware Qualified Tenant Advocates Program (2022–25)'. American Bar Foundation Access to Justice Research Initiative. Chicago, IL: American Bar Foundation.

Thomson Reuters (2026) 'Thomson Reuters leverages its world-class data assets to launch its own frontier model'. Press release, 24 August (https://www.thomsonreuters.com/en/press-releases/2026/august/thomson-reuters-leverages-its-world-class-data-assets-to-launch-its-own-frontier-model)

UNDP – United Nations Development Programme (2026a) *The algorithm in the courtroom: how artificial intelligence is reshaping justice and the rule of law across Asia and the Pacific.* Bangkok: UNDP Regional Bureau for Asia and the Pacific (www.undp.org/asia-pacific/publications/algorithm-courtroom).

UNDP (2026b) *Responsible AI in justice: regional policy guidance for UNDP programming in Asia-Pacific.* Bangkok: UNDP Regional Bureau for Asia and the Pacific (www.undp.org/asia-pacific/publications/algorithm-courtroom).

UNESCO (2025) *Guidelines for the use of AI systems in courts and tribunals*. Paris: UNESCO (https://unesdoc.unesco.org/ark:/48223/pf0000396582).

UN Secretary-General (2024) *Human rights in the administration of justice: report of the Secretary-General*. UN Doc. A/79/296, 7 August. New York: United Nations General Assembly (https://digitallibrary.un.org/record/4060365).

Uraz, J.N. (2026a) 'Justice under austerity: the impacts of reduced access to legal assistance in England and Wales'. SSRN working paper (https://ssrn.com/abstract=6516818).

Uraz, J.N. (2026b) 'Legal aid in England and Wales: a historical and budgetary perspective'. SSRN working paper (https://ssrn.com/abstract=6516838).

Vals AI (2025) 'Vals Legal AI Report (VLAIR): legal research' (www.vals.ai/industry-reports/vlair-10-14-25).

Varotsis, A. (2025) 'Transforming civic engagement with Caddy'. Blog. London: Incubator for AI, Department for Science, Innovation and Technology (https://ai.gov.uk/blogs/transforming-civic-engagement-with-caddy/).

WJP – World Justice Project (2019) *Measuring the justice gap: a people-centered assessment of unmet justice needs around the world*. Washington DC: World Justice Project (worldjusticeproject.org/our-work/research-and-data/measuring-justice-gap).

WJP (2023) *Dissecting the justice gap in 104 countries: WJP Justice Data Graphical Report I*. Washington DC: World Justice Project (worldjusticeproject.org/our-work/research-and-data/wjp-justice-data-graphical-report-i).

WJP (n.d). Atlas of legal needs (https://worldjusticeproject.org/legal-needs-atlas/)

WJP and World Bank (2025) *Advancing access to justice via information and communications technology: a literature review.* Washington DC: World Justice Project and World Bank (https://documents1.worldbank.org/curated/en/0990613251110

58349/pdf/P174073-228a7e33-deb9-4e8b-957a-adbef402e74d.pdf).

World Bank (2026a) *World development report 2026: the promise of artificial intelligence.* Washington DC: World Bank (https://hdl.handle.net/10986/45296).

World Bank (2026b) 'AI offers lifeline to developing economies in an era of weak growth'. Press release, 4 August (www.worldbank.org/en/news/press-release/2026/08/04/ai-offers-lifeline-to-developing-economies-in-an-era-of-weak-growth).

World Bank (2026c) *The AI repository* (https://airepository.worldbank.org)

Appendix 1 List of examples of AI tools

Definitions

Legal information, advice and assistance – providing tailored individual information and advice on everyday justice problems, and assistance to resolve them.

- *Information* refers to the provision of generic or tailored information/guidance.
- *Advice* refers to tailored advice on actions to take.
- *Assistance* refers to support with preparing court or other documents.

Informal dispute resolution – dispute resolution mechanisms outside the formal system, usually for less serious/low-value problems, accessed using informal processes.

Country	Name	Function
Argentina	'Vicky' – Ministerio Publico de Buenos Aires citizen virtual assistant (first judicial chatbot in Argentina)	Information (rights, procedures and referral, including free legal advice for family violence)
Argentina	Asistente Virtual del Poder Judicial de la Nacion (GPT-4.1-based portal assistant)	Information (procedural navigation of the federal case-management portal; escalation to human operators)
Australia	amica	Informal dispute resolution (family)
Australia	Justice Connect SmartAssist diagnostic AI (legal-problem classifier, with University of Melbourne)	Information (AI legal-problem diagnosis feeding intake, triage and referral)
Bangladesh	'Mina' LLM legal assistant (research prototype)	Information/Advice (*prototype*)
Bangladesh	JusticeNetBD (RAG chatbot on women's legal rights)	Information/Advice (women's legal rights: violence, dowry)
Brazil	Defensoria Publica AI: 'Ivo' federal public-defender chatbot, built by Serpro)	Assistance (AI-augmented free legal aid via public defenders)
Canada	Beagle+ (People's Law School)	Information
Canada	LLMediator (GPT-4 assisted ODR, Cyberjustice Laboratory) – research prototype	Informal dispute resolution (AI-assisted negotiation and mediation) – research stage

Country	Name	Function
China	12348 China Legal Services Network smart consultation; Baidu Faxingbao ('free AI lawyer'); district-government deployments	Information/Advice
Colombia	LegalApp (Ministry of Justice citizen 'Google juridico')	Information (AI-powered natural-language search over guided justice routes)
Czechia	Frank Bold GPT-4 legal-aid chatbot (pilot)	Information/Advice
Dominican Republic	SARA – Sistema de Atencion y Respuesta Automatizada (Ministry of Women)	Information/Advice (rights, services and referral) – gender-based violence only
Ecuador	LegalBot-EC (academic RAG chatbot on Ecuadorian law)	Information (*research prototype)*
Egypt	Shawwir (free AI legal consultation + lawyer booking)	Information/Advice with referral to Assistance
France	Juribot (free public legal AI) and the French consumer legal-chatbot market	Information/Advice (free and freemium public chatbots)
Greece	i-ACCESS My Rights (conversational AI for child victims); Ministry of Justice document AI	Information (co-designed legal chatbot for children)
India	CADRE (Centre for Advanced Dispute Resolution Excellence) – ODR platform	Informal dispute resolution – financial and commercial matters (debt recovery, credit, MSME and consumer money claims)
India	CORD (Centre for Online Resolution of Disputes) – ODR platform	Informal dispute resolution – financial and commercial matters (debt recovery, credit, MSME and consumer)
India	Jugalbandi – original schemes chatbot deprecated; continues as Jugalbandi Manager (OpenNyAI)	Information
India	Jupitice Justice Technologies – private digital court and ODR platform, with the *Saya* AI agent	Informal dispute resolution – financial and commercial matters (debt recovery, credit, MSME and consumer money claims)
India	LegalKart (AI legal information + lawyer matching)	Information/Advice with referral to Assistance
India	Multilingual consumer legal-AI market (Legal Setu, NyayGuru, NyayaCircle, Nyaypeeth, LawChat AI)	Information/Advice (with document generation and low-cost lawyer referral)
India	National Consumer Helpline (NCH)/INGRAM – AI-enabled pre-litigation consumer grievance redressal	Informal dispute resolution (pre-litigation redressal; company responds directly)
India	Nyaya Setu (voice-first multilingual AI legal assistant, built by the Digital India BHASHINI Division)	Information (with step-by-step procedural guidance)
India	Presolv360 – ODR platform	Informal dispute resolution – financial and commercial matters (debt recovery, credit, MSME and consumer money claims)

Country	Name	Function
India	Sama (ODRWays Solutions Private Limited) – ODR platform	Informal dispute resolution – financial and commercial matters (debt recovery, credit, MSME and consumer money claims)
India	Webnyay – AI-driven ODR and arbitration platform	Informal dispute resolution – financial and commercial matters (debt recovery, credit, MSME and consumer money claims)
Indonesia	Hukumonline AI suite – LIA (2018), AI Regulatory Compliance System (2020), Ask Hukumonline AI (2024), Allex	Information (consumer legal education; local-government legal literacy)
Japan	Bengo4.com AI legal consultation chat (with Bennavi AI online and other commercial Japanese legal-AI services)	Information/Advice
Kazakhstan	Adilet.zan.kz state legislation-portal AI chatbot (Ministry of Justice)	Information/Advice
Kazakhstan	Qazaq Law AI (virtual legal consultant on the AlemLLM national model, NITEC/MAIDD)	Information/Advice
Kenya	SheriaBot	Information/Advice
Kenya	Sheriaplex	Information/Advice
Kenya	Wakili.org/M-Wakili and 'Wakili AI'	Information/Advice
South Korea	Bubbi – Ministry of Justice everyday-law AI chatbot	Information/Advice
South Korea	Free AI legal counselling chatbots (DR & AJU AI and successors)	Information/Advice
South Korea	Korea Legal Aid Corporation (KLAC) – 'Beopryul Ttokttoki' AI legal information Q&A service (web chatbot, callbot and app) and the AI guidance layer of the Legal Aid Platform	Information/Advice
Kuwait	LegalBot.pro (Gulf legal AI assistant, law firm-backed)	Information/Advice
Malaysia	'Rachel' (GLT Law IBM Watson-based legal chatbot)	Information (common legal queries across several practice areas)
Mexico	Mexico City Labour Court outcome-prediction tool ('the calculator') – Sadka/Seira/Woodruff RCT	Informal dispute resolution (settlement support through personalised outcome prediction)
New Zealand	Rentbot, Workbot and Lagbot (Citizen AI with Community Law Centres)	Information/Advice (incl. document assistance) – service since wound down
Niger	Hadin kai-bani fondo ('Path for Cohesion')	Information/Advice
Nigeria	Podus AI	Information/Advice
Pakistan	NAZ Assist (Legal Aid Society/Sindh Legal Advisory Call Centre) – multilingual legal advisory chatbot with free human-lawyer verification	Information/Advice, with referral onto free human legal advice and litigation support
Pakistan	Pakistan Law Bot	Information/Advice

Country	Name	Function
Pakistan	QanoonAI	Information/Advice
Portugal	Guia Pratico da Justica (GPJ) – Practical Guide to Justice	Information (AI-supported legal information with guided procedural pathways to online services)
Senegal	LLM legal assistant research (knowledge graph of Senegalese law) – research stage	Information (research stage – not deployed)
Sierra Leone	Law Reform Commission AI chatbot (in development)	Information (planned – not yet operational)
South Africa	Maya (Legal Fundi)	Information/Advice (with referral to Assistance)
South Africa	My AI Lawyer (Legal&Tax/Legal Interact)	Advice/Assistance
Spain	Ministry of Justice accessible decision explanation tools (co-designed for different access needs)	Information (personalised plain-language explanation of documents users have actually received)
Spain	Private MASC pre-litigation negotiation platforms with AI assistants (post Ley Organica 1/2025)	Informal dispute resolution (AI-supported pre-litigation negotiation)
Sri Lanka	LankaLAW (Chat2Find) – AI legal research assistant	Information (free public legal knowledge platform with an AI assistant; no paywall)
Tanzania	Dada Wakili ('Sister Lawyer')	Information/Advice
Tanzania	Wakili Msomi (Sheria Kiganjani)	Information/Advice
Thailand	LAW-U sexual-crime victim guidance chatbot (Thammasat University)	Information/Advice
Uganda	BarefootLaw with 'Winnie' AI	Information/Advice/Assistance
Uganda	LawBot Africa – 'David' AI legal assistant (lawbot.africa)	Information/Advice (Uganda-specific legal Q&A on a free public tier; document generation and analysis on paid professional tiers – frontline fit partial)
Ukraine	LawLink Bot (Legal Development Network) – AI legal assistant for people in Ukraine	Information/Advice/Assistance (incl. automated court-document generation)
Ukraine	Pravoman – legal assistance platform for Ukrainians abroad (Legal Development Network with Paladin and Legal Innovations/LinGo)	Information/Advice
United Kingdom	Caddy, Citizens Advice	Advice (advisor-facing augmentation)
United Kingdom	Garfield AI (Garfield.Law) – first SRA-authorised AI-driven law firm	Assistance (AI-conducted litigation, end-to-end, under solicitor supervision) – financial and commercial matters
United Kingdom	Simply Resolved ('Molly' AI negotiation bot)	Informal dispute resolution (AI-facilitated negotiation, escalating to human mediation)
United Kingdom	Contend Legal – AI legal assistant for everyday civil problems (England and Wales)	Information/Advice, extending into Assistance (action plans and document drafting)

Country	Name	Function
United Kingdom	Yerty – AI employment-rights assistant and case-management platform (beta)	Information/Advice, with claim assessment and document generation
United States	AVA – Alaska Virtual Assistant (Alaska Court System, with LawDroid and NCSC)	Information (court-run procedural guidance; legal information, not legal advice)
United States	ExpungeMate (Legal Aid Society of Middle Tennessee & The Cumberlands)	Assistance (automated document preparation within legal aid clinics)
United States	Frontline Justice AI knowledge hub (Alaska community justice workers)	Advice/Assistance (advisor-facing augmentation of community justice workers)
United States	Knowledge Bot (Disability Rights Texas, with ASU AI Cloud Innovation Center and AWS)	Information (bilingual RAG assistant over the organisation's own legal content)
United States	LANC-LIA / LIA 2.0 – Legal Information Assistant (Legal Aid of North Carolina, with LawDroid)	Information (RAG over the organisation's own curated content; five problem areas)
United States	LawDroid and legal-aid AI copilots	Advice/Assistance (advisor-facing)
United States	Mighty (free AI case-value and settlement decision tool for accident victims)	Informal dispute resolution (client-side outcome prediction and settle-or-refer decision support)
United States	Rentervention	Information/Advice/Assistance
United States	Roxanne the Repair Bot (Housing Court Answers, built on Josef)	Information (housing repairs and conditions)
Vietnam	National Legal Portal AI legal assistant (with VNeID integration)	Information (statute navigation and basic legal guidance-type support)

Appendix 2 Characteristics of a well-designed public-facing bespoke legal chatbot

Section 3 identifies four main risks of using general-purpose LLMs for legal advice – accuracy, confidentiality, sycophancy and bias. A well-designed bespoke chatbot is one that is at least as accessible as a general-purpose LLM, is designed to mitigate each of these four risks, and is operated by an accountable organisation that ensures it is audited, publishes the results and budgets realistically for the human oversight needed to keep it accurate.

Accessibility maximised (so that it is at least as easy to use as a general-purpose LLM) by:

- Being free at point of use.
- Available 24/7, at the user's convenience.
- Providing quick replies to questions (matching general-purpose LLM response time).
- Accepting conversational questions in plain language, including spoken questions.
- Requiring no sign-up or registration – immediate anonymous access, rather than lengthy registration processes.
- Ideally offering multiple access channels, including interactive voice response (usable from a basic mobile phone), SMS, WhatsApp and online, so that people with limited connectivity or digital literacy are not excluded.
- Ideally offering multiple language options, including local languages.

Accuracy maximised by being:

- Grounded in (through RAG) a curated local knowledge base of trusted legal information, based on local law and on lawyers' experience of responding to everyday justice problems.

- Constrained to respond only based on this knowledge base, which is updated daily.
- Tested, before launch, against real-world local questions, with responses scored by a panel of lawyers for safety (legally accurate and unlikely to mislead) and value (tone, clarity and support for the user in taking next steps), and launched only once a high accuracy threshold is met.
- Tested for consistency – giving the same answer when the same question is phrased differently.
- Designed to elicit the relevant facts and to clarify unclear or wrongly framed questions, rather than answering a question based on a false premise.
- Reviewed continuously by AI so that complex, high-severity or risky conversations are flagged for mandatory review by a lawyer ('human on the loop') – moving to this only after an initial period in which lawyers review every conversation ('human in the loop') and accuracy has been demonstrated.
- Audited by lawyers on a continuous sample basis to check conversations are being accurately and helpfully answered. This is the final guard against model drift and the risks that the underlying general-purpose LLM may draw on the open internet for its answers or make things up.

Confidentiality risks minimised through:

- Simple automatic warning at start about confidentiality risks, before users share personal information.
- No requirement to enter personal details (and a knowledge base that itself holds no personal data).
- Removal of personal and identifying details before a query leaves the system, or automatic deletion of such details in real time.
- Contractual arrangements with any third-party general-purpose LLM providers, ideally with the zero-data-retention terms that large enterprise customers enjoy.
- Alternatively, where secure local servers are available, self-hosting an open-source model so that client information never passes to an external AI provider (accepting that the organisation then becomes responsible for keeping the system secure).
- A designated person in the organisation responsible for ensuring confidentiality.

- Making clear to users that, unlike a conversation with a lawyer, exchanges with the chatbot may not be protected by legal privilege.

Sycophancy avoided, and clients' best interests served, by:

- Providing objective responses that explain risks and the client's best interests, rather than reinforcing the user's existing beliefs or mirroring what they want to hear.
- Identifying the client's objectives and options, and the chances of success, rather than only summarising the law.
- Being candid about weak claims – discouraging people from pursuing claims that are unlikely to succeed or are not in their best interests, rather than encouraging them.
- Building in deterministic 'hard stops': where a query indicates an emergency (for example domestic violence where the client needs to leave the property), bypassing generative text entirely and giving immediate, direct contact information and clear directions, with escalation to a lawyer or human adviser.

Bias and limited relevance to diverse communities addressed by:

- A knowledge base built from locally relevant, context-specific legal material that reflects the everyday problems of diverse socioeconomic, educational and geographic groups.
- Design for the local context – including legal pluralism and customary law, local languages, connectivity, device usage and digital literacy – rather than importing a model from somewhere else.

Accountability, transparency and evaluation ensured by:

- Being operated by a trusted, accountable organisation (seen as independent of national and local government) rather than a distant platform.
- Being transparent that users are dealing with AI, and clear about the scope of the service (legal information and guidance, or – where regulation permits – advice and assistance such as drafting court documents), with professional indemnity insurance in place.
- Providing routes to escalate to a human adviser (for example a call-back from a lawyer to check AI-generated guidance), and a redress mechanism when things go wrong.
- Publishing the quality assurance process, test questions, accuracy results and costs (including unit costs), and ideally submitting to independent audit.

- Measuring value as well as safety – whether the conversation helped the user take the next step towards resolving their problem – and, over time, outcomes.

Sustainability and reach secured by:

- A realistic budget and timetable.
- Investment in marketing and endorsement by trusted organisations and influencers, so that people know the chatbot exists and prefer it to a general-purpose LLM.
- Links to face-to-face and telephone advice services, so that vulnerable and digitally excluded clients who need human support are referred on, and advice clinics can concentrate on them.