

Existing Standards for AI-Powered Business-to-Consumer Lawtech

A Landscape Review

Commissioned by the Legal Services Board

Prepared by o-x.ai

April 2026

Contents

| | |
|---|----|
| Executive Summary | 4 |
| Scope and Method | 4 |
| Landscape Overview | 5 |
| Thematic Findings | 5 |
| Cross-Cutting Patterns | 6 |
| Conclusions | 7 |
| Potential implications | 8 |
| 1. Introduction | 9 |
| 1.1 Background | 9 |
| 1.2 Purpose of This Research | 9 |
| 1.3 Promises of AI-Powered B2C Lawtech – and a Precondition | 9 |
| 1.4 The B2C Lawtech Landscape | 10 |
| 1.5 The Standards Ecosystem | 10 |
| 1.6 Key Definitions | 12 |
| 1.7 Report Structure | 12 |
| 2. Methodology | 13 |
| 2.1 Research Design | 13 |
| 2.2 Documents | 13 |
| 2.3 Key Informant Interviews | 13 |
| 2.4 Analytical Approach | 14 |
| 2.5 Limitations | 15 |
| 3. The Standards Landscape: An Overview | 17 |
| 3.1 Quantitative Overview | 17 |
| 3.2 Who Sets Standards and How | 19 |
| 3.3 AI-Specificity | 20 |
| 3.4 Overall Thematic Coverage | 21 |
| 4. Thematic Analysis | 24 |
| 4.1 Output Quality | 24 |
| 4.1.1 Consistency | 24 |
| 4.1.2 Reliability | 25 |
| 4.1.3 Accuracy | 26 |
| 4.2 Fairness and Inclusion | 26 |
| 4.2.1 Bias | 27 |
| 4.2.2 Fairness | 28 |
| 4.2.3 Non-Discrimination in AI | 28 |
| 4.2.4 Accessibility and Effective Use | 29 |
| 4.3 User Protection | 30 |
| 4.3.1 Transparency | 31 |
| 4.3.2 Signposting and Referral | 31 |
| 4.3.3 Identifying vulnerability users | 32 |
| 4.3.4 Human Oversight | 33 |

| | |
|--|----|
| 4.4 Governance and Assurance | 33 |
| 4.4.1 Quality Assurance | 34 |
| 4.4.2 Accountability | 35 |
| 4.4.3 Redress and Complaints | 35 |
| 4.4.4 Traceability | 36 |
| 4.5 Data and Privacy | 37 |
| 4.5.1 Privacy | 38 |
| 4.5.2 Data Protection | 38 |
| 5. Taxonomy Analysis: Cross-Cutting Dimensions | 40 |
| 5.1 System Type | 40 |
| 5.2 User Function | 41 |
| 5.3 Legal Scope | 42 |
| 5.4 Oversight Model | 43 |
| 5.5 Standards Approach | 43 |
| 5.6 Reserved and Non-Reserved Activity | 44 |
| 6. Cross-Theme Synthesis | 45 |
| 6.1 Convergences | 46 |
| 6.2 Tensions and Trade-Offs | 47 |
| 6.3 Standards Development Spectrum | 48 |
| 6.4 Cross-Sector Lessons | 50 |
| 7. Summary Gap Analysis | 52 |
| 7.1 Thematic Gaps | 52 |
| 7.2 Structural Gaps | 53 |
| 7.3 Taxonomy-Specific Gaps | 55 |
| 8. Interview Findings | 56 |
| 8.1 Awareness of Existing Standards | 56 |
| 8.2 What Interviewees Prioritise | 57 |
| 8.3 The Case For and Against Standards | 59 |
| 8.4 The Information-Advice Distinction | 60 |
| 8.5 Consumer Perspectives | 61 |
| 8.6 Design Choices and Quality Assurance in Practice | 62 |
| 8.7 Human Oversight and Phased Automation | 64 |
| 8.8 Consumer Education and Expectation Management | 65 |
| 8.9 Beyond Standards | 66 |
| 9. Conclusions and Implications | 68 |
| 9.1 Answering the Research Questions | 68 |
| 9.2 Potential implications | 69 |
| Annexes | 70 |
| Annex A: Glossary | 70 |
| Annex B: References | 70 |

Executive Summary

The Legal Services Board (LSB) commissioned this landscape review to understand the existing standards relevant to AI-powered business-to-consumer lawtech (B2C lawtech).ⁱ The research examines the implications of two related developments: the fast-paced development and roll out of artificial intelligence (AI) tools that deliver legal information, guidance, and advice directly to consumers, and the absence of a coherent standards framework governing such tools. As the oversight regulator for legal services in England and Wales, the LSB has a statutory interest (grounded in the regulatory objectives of the Legal Services Act 2007) in ensuring that innovation in legal services delivery supports improved access to justice, protects consumers, and maintains public confidence in the rule of law.

This review was carried out against the backdrop of the UK Government's growth agenda and its pro-innovation approach to AI regulation, set out in the AI Regulation White Paper (DSIT, 2023) and subsequent policy statements that emphasise proportionate, context or sector-specific governance rather than more general horizontal legislation. The tension between encouraging and enabling innovation and protecting consumers (particularly vulnerable consumers who may be most at risk when using low-cost AI-powered legal tools) forms the central policy challenge informed by this research. If the realistic alternative to an imperfect AI tool is no legal support at all, the risk-benefit calculus is not straightforward.

Scope and Method

We adopted a systematic, AI-supported approach to map and understand the current standards landscape. In total, we reviewed 234 documents from 70 organisations across 8 jurisdictions and found that 160 contain insights relevant to at least one of the themes under review. This was complemented by interviews with stakeholders across the legal ecosystem, including legal professionals, lawtech providers, regulators, consumer groups, and standards bodies.

To analyse the material, we used an AI tool called *Kompass*, which enabled us to extract, organise, and compare information at a scale and speed that would not have been feasible manually.

Documents were identified through a structured search process, including scanning organisational websites, querying standards databases, and following citation trails between documents. Each document was then assessed to ensure it had substantive relevance to AI, rather than only incidental references to technology.

All documents were systematically coded against 17 themes, grouped into five core areas:

- quality of outputs
- fairness and inclusion
- user protection
- governance and assurance
- data and privacy

In addition, we analysed the documents using six further lenses, as follow:

- the type of system
- its intended use
- the area of law it applies to
- oversight mechanisms
- the type of standard (e.g. formal or informal)

- whether it relates to reserved or non-reserved legal activities

Landscape Overview

The standards landscape for AI-powered B2C lawtech is characterised by breadth without depth.ⁱⁱ Guidance documents dominate (135 documents, 58% of the documents identified), reflecting an ecosystem in which most standard-setting activity takes the form of principles, recommendations, and best practice frameworks rather than binding requirements with enforcement mechanisms. Formal standards from bodies such as the International Organization for Standardization (ISO), Institute of Electrical and Electronics Engineers (IEEE), and British Standards Institution (BSI) account for 11% of the documents (26 documents), while binding regulation represents just 8% (18 documents). Codes of practice (12 documents, 5%) and assurance schemes (5 documents, 2%) complete the picture of a governance architecture that is advisory and fragmented.

Within this landscape, the EU AI Act emerges as one of the most significant regulatory instrument, establishing a risk-based classification framework with binding obligations for high-risk AI systems. Its implications for B2C lawtech, however, remain subject to interpretation as the regulatory framework is applied in practice. Alongside this regulatory framework, a number of influential international governance instruments shape the broader standards environment, most notably the US (National Institute of Standards and Technology) NIST AI Risk Management Framework and the ISO/IEC 42000 series of standards addressing AI management systems and lifecycle governance.

A wide range of organisations produce relevant standards. Government bodies are the most prolific contributors (73 documents or 31%), followed by standards bodies (62 or 26%) and industry associations (44 or 19%). Cross-sector regulators contribute 22 documents (9%), academic and research institutions 16 (7%), and legal regulators 12 (5%). Consumer and advice bodies account for 4 documents (2%).

Interview data suggests that in the absence of clear, accessible, and AI-specific standards, lawtech providers have arrived at their own working frameworks through a combination of adjacent regulatory requirements, professional instinct, and trial and error.

Thematic Findings

Thematic analysis across the 17 areas reveals inconsistent and patchy coverage. Data protection and transparency benefit from established and mature regulatory frameworks (culminating in the UK GDPR and the EU General Data Protection Regulation and recent transparency requirements in AI-specific instruments). Consequently, these two areas receive the most consideration across the majority of documents reviewed.

Important areas, especially those providing consumer protection in B2C lawtech, receive much less attention. Even though identifying vulnerability is central to consumer protection (particularly where individuals are, for whatever reasons, at greater risk of vulnerability), vulnerability identification emerges as the weakest area across the entire landscape. Few documents provide operationally meaningful guidance on how AI systems should identify, respond to, and accommodate vulnerable users. Signposting and referral mechanisms, that direct users to appropriate human assistance when AI tools reach the limits of their competence, are similarly underdeveloped. Redress and complaints frameworks, essential for accountability when AI-powered tools produce harmful outcomes, also receive limited treatment, particularly in the context of non-regulated providers.

Figure 11, from the cross-theme synthesis in this report, shows this variation in a heatmap. It shows coverage (number of documents) containing content related to each standard areas (rows) and taxonomy element (columns), which is the type of context of the AI use. So for example, the top left cell has 61 documents about data protection relating to either general-purpose AI platforms, or legal

information platforms or issue-specific legal tools.

Figure 11: Heatmap showing coverage (number of documents) by theme and taxonomy

| | System Type | User-Facing Function | Legal Scope | Oversight Model | Standards Approach | Reserved / Non-Reserved | |
|------------------------------|-------------|----------------------|-------------|-----------------|--------------------|-------------------------|------------------------|
| Data Protection | 61 | 52 | 27 | 73 | 73 | 15 | Data & Privacy |
| Privacy | 55 | 49 | 24 | 69 | 69 | 11 | |
| Transparency | 101 | 92 | 47 | 126 | 128 | 21 | Governance & Assurance |
| Accountability | 94 | 88 | 46 | 115 | 117 | 19 | |
| Human Oversight | 84 | 78 | 44 | 101 | 101 | 20 | |
| Quality Assurance | 72 | 63 | 33 | 95 | 96 | 14 | |
| Traceability | 71 | 62 | 31 | 90 | 89 | 12 | |
| Bias | 71 | 67 | 33 | 90 | 89 | 15 | |
| Fairness | 71 | 65 | 32 | 88 | 89 | 14 | Fairness & Inclusion |
| Non-discrimination | 60 | 56 | 29 | 77 | 76 | 11 | |
| Accessibility | 43 | 43 | 24 | 46 | 46 | 12 | |
| Accuracy | 68 | 63 | 33 | 84 | 84 | 13 | Output Quality |
| Reliability | 66 | 60 | 32 | 86 | 86 | 15 | |
| Consistency | 33 | 32 | 17 | 45 | 46 | 9 | |
| Signposting & Referral | 54 | 51 | 28 | 59 | 60 | 12 | User Protection |
| Vulnerability Identification | 66 | 62 | 29 | 78 | 79 | 12 | |
| Redress & Complaints | 41 | 39 | 24 | 48 | 48 | 12 | |

Taxonomy dimensions: System Type classifies AI tools as general-purpose (e.g. ChatGPT, Claude), legal information platforms (e.g. legal databases), or issue-specific legal tools (e.g. housing, debt, employment). User-Facing Function distinguishes between legal information provision, personalised legal advice, and intermediate functions such as triage and eligibility checking. Legal Scope indicates whether standards address AI in legal services broadly or target narrow, function-specific applications. Oversight Model captures how human oversight is structured: human-in-the-loop review, autonomous operation, or hybrid arrangements. Standards Approach classifies instruments by regulatory modality: rules-based, principles-based, or outcomes-based. Reserved / Non-Reserved indicates whether standards differentiate between AI used for reserved legal activities (e.g. rights of audience, conduct of litigation, etc.) and non-reserved legal services.

Cross-Cutting Patterns

Three strong cross-cutting patterns emerge:

1. **Assumption that a legal professional is involved.** The existing standards landscape assumes that there is a human professional intermediary (a solicitor, barrister, or other regulated professional) between the AI system and the end user. This assumption is

embedded in the regulatory architecture of legal services, where obligations attach to the regulated individual or firm rather than to the technology. For B2C lawtech tools that operate without the involvement of legal professionals, this assumption creates a structural gap: the regulatory obligations that would ordinarily protect the consumer do not apply, or may apply to no identifiable duty-holder.

2. **Standards not designed for AI powered lawtech.** The distinction between AI-specific standards and general obligations applied to AI proves analytically significant. Many standards that are relevant to AI-powered lawtech were not designed with AI in mind; they represent professional conduct rules, data protection requirements, or consumer protection obligations that apply to AI systems by extension rather than by design. While such standards provide a baseline, they do not address AI-particular risks including hallucination (the generation of plausible but fabricated legal information), emergent bias arising from training data, opacity in reasoning processes, and the difficulty of attributing responsibility for automated outputs.
3. **AI platform providers' policies acting as proxy 'standards'.** AI platform providers (principally OpenAI, Anthropic, Google, and Microsoft) exercise a form of de facto governance through their usage policies, content policies, and technical safeguards. These instruments shape the capabilities and limitations of the foundation models upon which many B2C lawtech tools are built, yet they operate entirely outside formal regulatory frameworks, are modifiable at the provider's sole discretion, and offer no meaningful redress mechanism for affected consumers.

Conclusions

RQ1: Are there any standards that apply to B2C lawtechs?

Yes, but very few are designed specifically for consumer-facing legal technology. Most relevant standards are general AI, data protection, or professional regulatory frameworks. As a result, consumers using unregulated AI legal tools are primarily protected by broad regimes such as the GDPR and consumer protection law, which do not directly address the risks of AI-generated legal advice or guidance.

RQ2: Who sets these and how are they developed?

Standards are set across multiple layers of governance, including international technical bodies (e.g. ISO/IEC, NIST, IEEE), supranational legislators such as the EU, UK cross-sector regulators (ICO, CMA, FCA, Ofcom), legal services regulators, and AI platform providers. Each develops standards through different processes and with different enforcement powers, producing a fragmented regulatory landscape.

RQ3: What format do they take and how are they communicated?

Most standards take a soft approach, primarily guidance documents, with a smaller number of formal standards, regulations, and codes of practice. Communication is fragmented: some standards sit behind paywalls, others are dispersed across organisational websites, and platform policies change unilaterally, contributing to low awareness among those developing and deploying B2C lawtech tools. As a result, providers may not actively engage with these standards, meaning their existence does not consistently translate into meaningful consumer protection.

RQ4: What do the standards cover?

Coverage is uneven. Data protection and transparency are relatively well developed, while bias, accountability, privacy, and human oversight receive moderate attention. Areas most relevant to consumer protection (such as vulnerability identification, signposting and referral, redress, accessibility, and reliability) are weakly covered.

RQ5: Are there any obvious gaps?

Significant gaps exist in standards for vulnerability identification, referral pathways, and redress mechanisms. Structural gaps also persist against a backdrop of limited regulation, including the unclear boundary between legal information and advice in AI contexts, and the assumption of professional intermediation that does not hold for direct-to-consumer AI tools.

Potential implications

The following observations emerge from the evidence gathered in this review:

1. The review identifies key gaps (namely vulnerability identification frameworks, signposting and referral protocols, mechanisms for redress of AI-related complaints, and clear consumer-facing transparency requirements). Because these areas currently present the greatest consumer risk while workable models already exist in other sectors.
2. There are existing standards in other sectors that could help inform approaches to developing standards for AI-powered lawtech. For example, operational frameworks from healthcare (such as pre-market validation, escalation pathways for vulnerable users, and controlled AI modification processes). Also financial services (including conduct-of-business rules, senior manager accountability, and outcomes-based consumer protection duties). Both provide tested structures that could be tailored to legal services, where regulation covers individuals and entities for set activities instead of services and products, rather than invented anew. Any such adaptation would need to reflect the current statutory position, under which legal services regulators do not currently have powers to regulate consumer-facing AI products offered by unregulated providers, and are limited to overseeing authorised professionals and the carrying out of specific reserved legal activities.
3. The evidence highlights a “gatekeeper gap”. Where those who interact with a regulated lawyer who uses AI tools are protected by regulation but those who interact with AI tools directly are not.
4. Risks vary significantly across different types of B2C lawtech tools. Tools that provide personalised legal guidance directly to consumers, particularly in higher-risk or sensitive contexts, present different and potentially higher risks. Compared to those offering general legal information or operating under professional supervision. This points to the potential value of approaches that differentiate expectations according to use case and risk.
5. The traditional distinction between legal information and legal advice is difficult to apply in AI contexts. Many tools generate outputs that are tailored to a user’s specific situation, regardless of how the service is labelled. This suggests a need for a more practical way of distinguishing between the two, potentially based on whether users are likely to rely on the output when making decisions.

1. Introduction

1.1 Background

AI-powered B2C lawtech presents both an opportunity and a challenge for the legal services sector. The opportunity is substantial: AI tools have the potential to address the significant unmet legal need documented in successive surveys of the public's experience of legal problems (Legal Services Board, 2023). They offer accessible, affordable, and immediate assistance to individuals, some of whom might otherwise receive no help at all. The challenge is equally significant: AI tools operate in a domain where inaccurate, incomplete, or inappropriate outputs can have serious consequences for individuals' legal rights, financial security, and personal wellbeing. They also operate where the existing regulatory architecture was not designed for (and may not accommodate) direct-to-consumer AI deployment.

1.2 Purpose of This Research

This report presents the findings of a landscape review commissioned by the Legal Services Board (LSB) to map and analyse existing standards, regulations, guidance, and related instruments relevant to artificial intelligence-powered business-to-consumer (B2C) lawtech. The research was designed to provide the LSB with a comprehensive evidence base to understand the existing standards, regulations, guidance, and related instruments relevant to AI-powered tools delivering legal information, guidance, or advice directly to consumers.

Specifically, it addresses the following research questions (RQs):

- RQ1: Are there any standards that apply to B2C lawtechs?
- RQ2: Who sets these and how are they developed?
- RQ3: What format do they take and how are they communicated?
- RQ4: What do the standards cover in terms of service offer, transparency and accessibility and consumer protection?
- RQ5: Are there any obvious gaps?

1.3 Promises of AI-Powered B2C Lawtech – and a Precondition

AI-powered B2C lawtech carries two distinct promises. The first is the promise of better **access to justice**: the prospect that AI tools can extend the reach of legal assistance to the estimated 32% of adults in England and Wales who experience a legal problem but do not obtain professional legal help, or receive inadequate or untimely help, and are left with an unmet legal need (Legal Services Board, 2023). For individuals and small businesses deterred by cost, complexity, or lack of awareness, AI-powered tools offer a cheap (or free) and accessible potential route to understanding their legal position, identifying their options, and taking appropriate action. The economic significance of this promise extends beyond individual welfare; the UK Government's growth agenda positions lawtech innovation as a contributor to economic productivity, with the AI Opportunities Action Plan (Department for Science, Innovation and Technology, 2025) and the AI Action Plan for Justiceⁱⁱⁱ (Ministry of Justice, 2025) suggesting legal services as a sector where AI adoption can drive efficiency and expand market access. The Department for Business and Trade's SME Growth Plan notes that SMEs 'tolerating' legal issues such as contractual disputes costs firms £11.6 billion annually (Federation of Small Businesses, 2024).

The second promise is **economic growth** within the legal services sector itself. AI-powered tools can enable legal service providers (both regulated and unregulated, including advice services) to serve a wider market at lower cost. Potentially expanding the overall size of the legal services economy while

improving national productivity. The UK Government's pro-innovation approach to AI regulation, articulated in the AI Regulation White Paper (Department for Science, Innovation and Technology, 2023^{iv}), explicitly seeks to create conditions in which such innovation can flourish without disproportionate regulatory burden.

Yet both promises depend on a precondition: **trust**. Consumers need to be able to trust that AI-powered legal tools will provide accurate information, treat them fairly, protect their data, recognise their vulnerabilities, and direct them to human assistance when the situation demands it. Trust depends on the trustworthiness of AI-powered legal tools.

1.4 The B2C Lawtech Landscape

The market for AI-powered B2C lawtech tools has expanded rapidly since the public release of large language models in late 2022, building on an earlier generation of rule-based and template-driven legal technology. The current landscape encompasses a spectrum of tools that vary in their sophistication, their relationship to the regulated legal profession, and the nature of the service they provide to consumers. The UK is home to 44% of Europe's lawtech startups (Ministry of Justice, 2025).

At one end of this spectrum are tools that provide **legal information**: factual statements about the law, explanations of legal concepts, or descriptions of legal processes, presented without any consideration of the user's specific circumstances. General-purpose AI assistants such as ChatGPT (OpenAI), Claude (Anthropic) and Microsoft Copilot are increasingly used for this, despite not being designed or validated for legal queries. At the other end are tools that provide legal advice: they apply legal principles to the user's specific facts and circumstances, recommend a course of action and may even assist in carrying out those steps. Between these poles lies a range of tools offering guided navigation of legal processes, document assembly, dispute resolution support, and triage services that help users understand whether they need professional legal assistance.

Dedicated B2C lawtech providers illustrate this spectrum. **Contend Legal**, for example, offers AI-powered legal guidance positioned explicitly as accessible and affordable assistance for individuals who cannot or choose not to engage a solicitor. **AskEllie+** provides AI-assisted guidance on family law matters, targeting an area of law where unmet need is particularly acute. **Garfield AI** and **Legal Utopia** offer AI-driven legal assistance aimed at consumers navigating specific legal problems. **DoNotPay**, which is amongst the earliest and most high-profile B2C lawtech service and attracted significant attention (and litigation) in the United States, offered an AI tool marketed as a "robot lawyer" capable of assisting with consumer disputes, parking tickets, and similar matters. These purpose-built tools operate alongside the incidental use of general-purpose AI systems for legal queries, a practice that is widespread but largely unmeasured.

1.5 The Standards Ecosystem

Standards play a central role in building trust by establishing shared expectations across the ecosystem. At their simplest, standards can be understood as agreed rules, frameworks, or benchmarks that guide behaviour, design, and performance. They encompass a broad spectrum of instruments – from formal technical specifications to informal guidance and industry norms – varying in their degree of authority, enforceability, and scope.

Depending on how they are designed and applied, they can either enable innovation and growth or constrain it. For providers, trust lies in having a clear and proportionate framework that supports responsible innovation without arbitrary restriction. For regulators, it lies in confidence that standards will deliver meaningful consumer protection in practice. Without this mutual trust, adoption will be limited, and the potential gains in access to justice and economic growth will not be fully realised.

Understanding the governance landscape for AI-powered B2C lawtech requires clarity about what standards are, how they differ from regulation, and how various types of instruments interact.

Formal standards are established by consensus through recognised standards bodies and approved by a recognised standardisation organisation. In the AI domain, the principal bodies are the International Organization for Standardization (ISO) and the International Electrotechnical Commission (through their joint technical committee ISO/IEC JTC 1/SC 42 on Artificial Intelligence), the Institute of Electrical and Electronics Engineers (IEEE). As well as national bodies such as the British Standards Institution (BSI) and the US National Institute of Standards and Technology (NIST). Formal standards carry authority through their development process (typically involving multi-stakeholder consultation and iterative drafting) but are generally voluntary unless incorporated into regulation by reference. ISO/IEC 42001:2023 (AI Management Systems) and the IEEE 7000 series on ethical AI represent significant contributions to this category.

Regulation comprises legally binding requirements imposed by bodies with statutory authority. In the AI context, the EU AI Act (Regulation (EU) 2024/1689) is the most significant example, establishing a risk-based classification framework with graduated obligations. In the UK, the regulatory landscape is shaped by the Government's decision to pursue a principles-based, sector-specific approach rather than horizontal AI legislation (i.e. legislation applying a single, cross-economy regulatory framework for AI across all sectors). With existing regulators (the ICO, FCA, CMA, Ofcom, and others) expected to apply cross-cutting principles (safety, transparency, fairness, accountability, contestability) within their respective domains. For legal services specifically, the regulatory framework established by the Legal Services Act 2007 imposes obligations on the regulatory bodies (the SRA, BSB, CLC, and others) and, through them, on regulated legal professionals and entities. Alongside these statutory requirements, sector-specific certification schemes such as the ICO's Legal Services Operational Privacy Certification Scheme (LOCS) provide additional, tailored standards for data protection practices within legal services.

Guidance (the largest category in this review) occupies an intermediate position. Issued by regulators, government bodies, industry associations, and standards organisations, they articulate expectations, interpret existing obligations, and recommend practices without imposing legally binding requirements. Their authority derives from the standing of the issuing body and their alignment with regulatory expectations. The ICO's guidance on AI and data protection, the SRA's guidance on the use of AI in legal services, and NIST's AI Risk Management Framework exemplify this category. The ICO also operates the Legal Services Operational Privacy Certification Scheme (LOCS), which provides sector-specific privacy standards for legal services.^v

Codes of practice represent collective commitments by industry participants to specified standards of conduct. In the legal services context, the Law Society's practice notes and the Council of Bars and Law Societies of Europe's (CCBE's) guidance on AI for European lawyers operate in this mode.

Platform policies (the usage terms, content policies, and acceptable use policies of AI platform providers such as OpenAI, Google, and Microsoft) represent a distinctive and increasingly significant category of de facto standards. They shape the capabilities and limitations of the foundation models upon which many B2C lawtech tools are built.

Assurance schemes provide independent verification that an organisation or product meets specified standards, offering a mechanism for demonstrating compliance and building consumer trust.

The interaction between these different types of instrument is complex and consequential. Voluntary standards may inform regulatory expectations; regulatory requirements may reference or incorporate formal standards; platform policies may impose constraints that supersede or supplement formal regulation. For B2C lawtech specifically, this interplay is further complicated by the boundary between regulated and unregulated legal services. In England and Wales, only certain activities (so-called

reserved legal activities) fall within the scope of legal services regulation. Meaning that many tools, particularly those providing general legal information or support, may sit outside formal regulatory oversight despite posing material risks to consumers.

This complexity is reinforced by the fact that, in legal services, regulation applies primarily to practitioners and entities rather than to services or products themselves. As a result, lawtech tools are typically regulated indirectly (through the obligations imposed on the professionals and firms that develop, deploy, or rely on them) rather than being subject to direct product-level regulation.

1.6 Key Definitions

For the purposes of this report, the following definitions apply:

- **Lawtech** refers to technology applied to legal services, encompassing tools used by legal professionals and those used directly by consumers.^{vi}
- **B2C lawtech** denotes lawtech tools that deliver services (whether characterised as information, guidance, or advice) directly to consumers without requiring the intermediation of a regulated legal professional.
- **AI-powered** refers to tools that employ artificial intelligence techniques, including but not limited to large language models, machine learning classifiers, and natural language processing, as a core component of their functionality.
- **Standards** is used broadly to encompass formal standards (issued by recognised standards bodies), regulations, guidance, codes of practice, platform policies, and assurance schemes that establish expectations for the development, deployment, governance and quality of AI-powered tools.

1.7 Report Structure

The remainder of this report is organised as follows. Section 2 sets out the methodology, covering the research design, body of documents identified, semi-structured interviews, analytical approach, and limitations. Section 3 presents a quantitative overview of the standards landscape, examining the distribution of documents by type, organisation, jurisdiction, and thematic coverage.

Subsequent sections present detailed thematic analysis (Section 4), cross-cutting findings (Section 5), synthesis of patterns and tensions across themes (Section 6), and a gap analysis identifying thematic, structural and taxonomy-specific gaps in the current landscape (Section 7).

Section 8 presents findings from stakeholder interviews, including levels of awareness of existing standards, practitioner priorities, perspectives on the value of standards, and views on key issues such as the information–advice distinction and consumer needs. Finally, Section 9 draws together the overall conclusions and discusses potential policy implications. Annexes provide supporting materials: specifically, a glossary and references.

2. Methodology

2.1 Research Design

This review used a rapid landscape review methodology, designed to map the breadth of a diverse and rapidly evolving field within a short timeframe. The research drew on three complementary evidence streams:

- LSB and research team input, which helped frame the scope, identify priority issues, and interpret findings in the context of legal services regulation, AI governance, and standards development;
- Structured document analysis, which provided a systematic mapping of existing standards, guidance, and regulatory instruments; and
- Semi-structured stakeholder interviews, which tested emerging findings, filled gaps in the documentary record, and provided practical insight into how these instruments operate in practice.

The research was conducted using an AI-augmented research tool, Kompass, that enabled systematic identification, extraction, coding, and cross-referencing of documents at a scale and speed not achievable through purely manual methods. The tool facilitated web scraping and PDF extraction from source organisations, structured coding of documents against the analytical framework, and quantitative analysis of coverage patterns across the body of documents identified. All AI-assisted outputs were subject to human expert review and validation, ensuring that the efficiencies of AI augmentation did not compromise analytical rigour.

2.2 Documents

The review examined **234 documents** from **70 organisations** across **8 jurisdictions**. We found 160 that contained insights relevant to at least one theme.

Documents were included, based on human judgement, where they covered at least one of the themes of interest and that content related to AI. They were sourced through a multi-stage identification process. The primary method used was **organisation-level scanning**: we systematically reviewed publications, guidance, and policy documents of organisations with plausible relevance to AI governance, legal services regulation, consumer protection, or technology standards. This was supplemented by **standards database searches** across the ISO, IEEE, BSI, and NIST catalogues, and by **snowball referencing**, following citations and cross-references within identified documents to locate further relevant materials.

Jurisdictions covered include the UK, European Union, United States, Canada, and Australia, alongside international standards bodies. These were selected to capture leading approaches to AI governance and legal services regulation. This reflects the reality that AI systems, standards, and providers operate across borders, meaning developments in other jurisdictions are often directly relevant to the UK B2C lawtech market.

The review covers the six document types outlined in section 2: regulations, formal standards, codes of practice, guidance, platform policies, and assurance schemes. White papers and research reports with standards-relevant content were also included where they included substantive content around expectations for AI governance.

2.3 Key Informant Interviews

Document collection was supplemented by semi-structured interviews with stakeholders across the B2C lawtech ecosystem. The interview protocol was informed by the sector the interviewee came

from and by the thematic areas and taxonomy dimensions of the analytical framework. With particular emphasis on issues that are not easily identified through documentary analysis alone. Such as how standards are interpreted in practice, operational challenges in applying them, and areas where stakeholders perceive gaps in existing governance arrangements. Interviews explored participants' awareness of existing standards, their assessment of coverage gaps, their experience of applying standards in practice, and their views on the potential design of a voluntary standards framework for B2C lawtech. Interview findings are both presented in a separate section and integrated throughout the body of this report.

Interview participants were purposively selected to represent a range of stakeholder perspectives relevant to AI-enabled B2C lawtech. These included lawtech providers developing consumer-facing tools, regulators with responsibility for legal services or AI governance, consumer and advice organisations representing legal services users, and industry associations representing legal sector stakeholders.

2.4 Analytical Approach

The analytical framework for the review employed a dual structure (see Figure 1). First, **thematic analysis** coded each document against 17 thematic areas. These were organised into five domains, and each domain analysed whether the document provided substantive coverage (addressing the theme with specific, actionable provisions), partial coverage (referencing the theme without detailed treatment), or no coverage.^{vii} This coding enabled quantitative analysis of coverage patterns, identifying which themes are well-served by existing standards and which represent gaps.

The thematic areas address the substantive content of standards and are grouped as follows:

- An **output quality** domain encompasses consistency (whether the tool produces reliable outputs across similar queries), reliability (whether the tool functions dependably over time and under varying conditions), and accuracy (whether outputs are factually and legally correct).
- A **fairness and inclusion** domain covers bias and fairness (whether the tool treats users equitably regardless of protected characteristics), non-discrimination (whether the tool avoids unlawful discrimination in its outputs or operation), and accessibility (whether the tool is usable by individuals with disabilities or limited digital literacy).
- A **user protection** domain addresses transparency (whether users understand they are interacting with AI, how it works, and its limitations), signposting and referral (whether the tool directs users to appropriate human assistance), vulnerability identification (whether the tool recognises and responds appropriately to vulnerable users), and human oversight (whether adequate human supervision of the tool's operation exists).
- A **governance and assurance** domain encompasses quality assurance (systematic processes for ensuring output quality), accountability (clear allocation of responsibility for the tool's outputs and impacts), redress and complaints (mechanisms for users to challenge outcomes and obtain remedies), and traceability (the ability to reconstruct and examine the tool's reasoning).
- Finally, a **data and privacy** domain covers privacy (protection of personal and sensitive information) and data protection (compliance with data protection legislation).

Second, **taxonomy analysis** classified each document across six dimensions:

- **System type** i.e. the technological architecture of the AI tool.
- **User function** i.e. the role the tool plays for the consumer.
- **Legal scope** i.e. the areas of law addressed.

- **Oversight model** i.e. the governance and supervision arrangements.
- **Standards approach** i.e. principles-based, rules-based, or outcomes-based.
- **Distinction between reserved and non-reserved legal activities**, which determines the applicability of legal services regulation.

This dual structure (see Figure 1 below) enabled both detailed analysis of individual themes and cross-cutting synthesis across the landscape. By mapping identified standards across these domains and taxonomy dimensions. To highlight where standards are concentrated, where gaps exist, and how coverage varies across different types of tools and use cases.

The resulting dataset supports descriptive statistical analysis of coverage patterns alongside qualitative interpretation of the standards landscape and its gaps. Cross-theme synthesis examined relationships between thematic areas, identifying patterns such as the co-occurrence of transparency and accountability coverage or the correlation between AI-specificity and operational detail. Gap analysis identified themes, sectors, and contexts where existing standards are absent, insufficient, or structurally unsuited to the B2C lawtech context.

2.5 Limitations

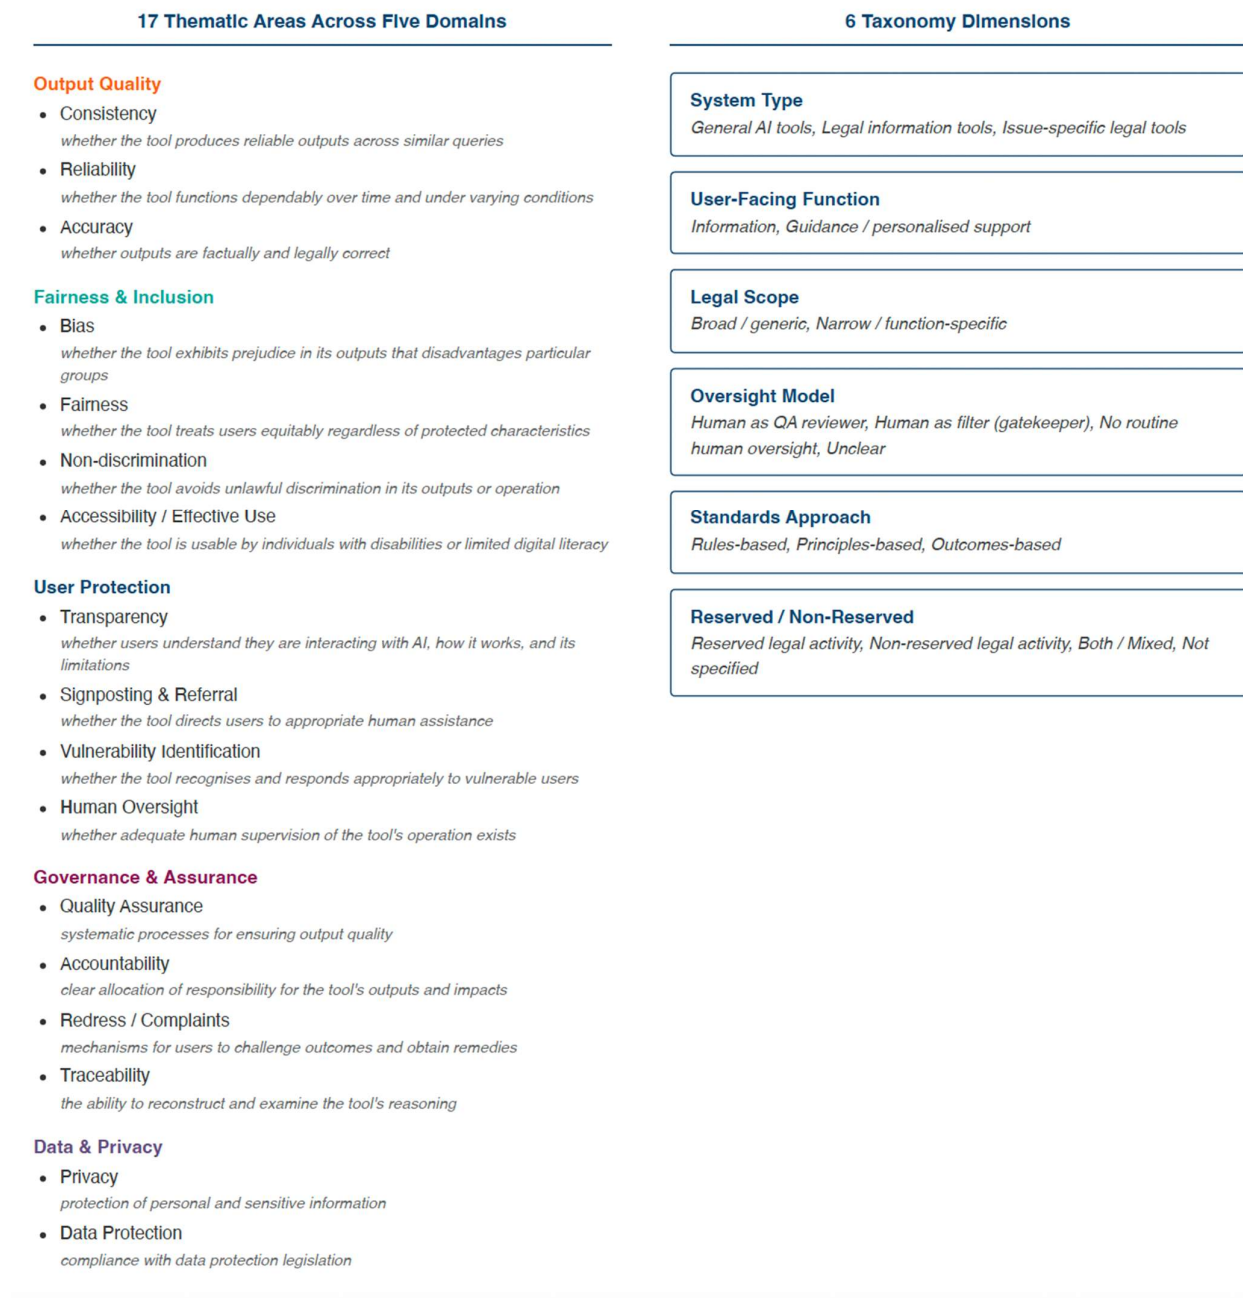
Several limitations and scoping decisions should be acknowledged. AI governance is **evolving rapidly**: new standards, guidance, and regulatory instruments are continually being produced, and documents published after the cut-off date of 6 March 2026 for data collection are not captured. The body of documents therefore reflects the state of the landscape at the point of data collection and should be understood as a snapshot rather than a definitive inventory.

Access constraints affected coverage of formal standards. Several ISO and IEEE standards require purchase and are costly, meaning full-text analysis was not possible in all cases. Where access was restricted, we drew on publicly available abstracts, scope statements, and secondary literature, which may not capture the full detail of the standard's provisions.

The **interview sample** was purposively selected to capture a range of perspectives and should not be treated as representative of any stakeholder population. The insights derived from interviews are therefore indicative rather than generalisable and are intended as a complement to the documentary evidence rather than serve as standalone findings.

The **jurisdictional scope** of the review is primarily England and Wales, reflecting the LSB's regulatory remit. International documents are included as comparators and sources of transferable practice. This approach was intended to ensure the review remained focused and manageable while still capturing relevant global developments, rather than to provide exhaustive coverage of all jurisdictions' standards landscapes. The EU, US, Australian, and Canadian materials included were selected on this basis.

Figure 1: Analytical framework



3. The Standards Landscape: An Overview

3.1 Quantitative Overview

Key finding: Guidance documents dominate (58%), followed by white papers and reports (15%), formal standards (11%), regulation (8%), and codes of practice (5%), with assurance schemes and other document types making up the small remainder. The governance architecture for AI in legal services is predominantly advisory rather than enforceable.

The body of documents provides a sufficiently large and diverse evidence base to support a robust analysis of the standards landscape. The distribution of documents by type reveals a governance architecture that is predominantly advisory in character.

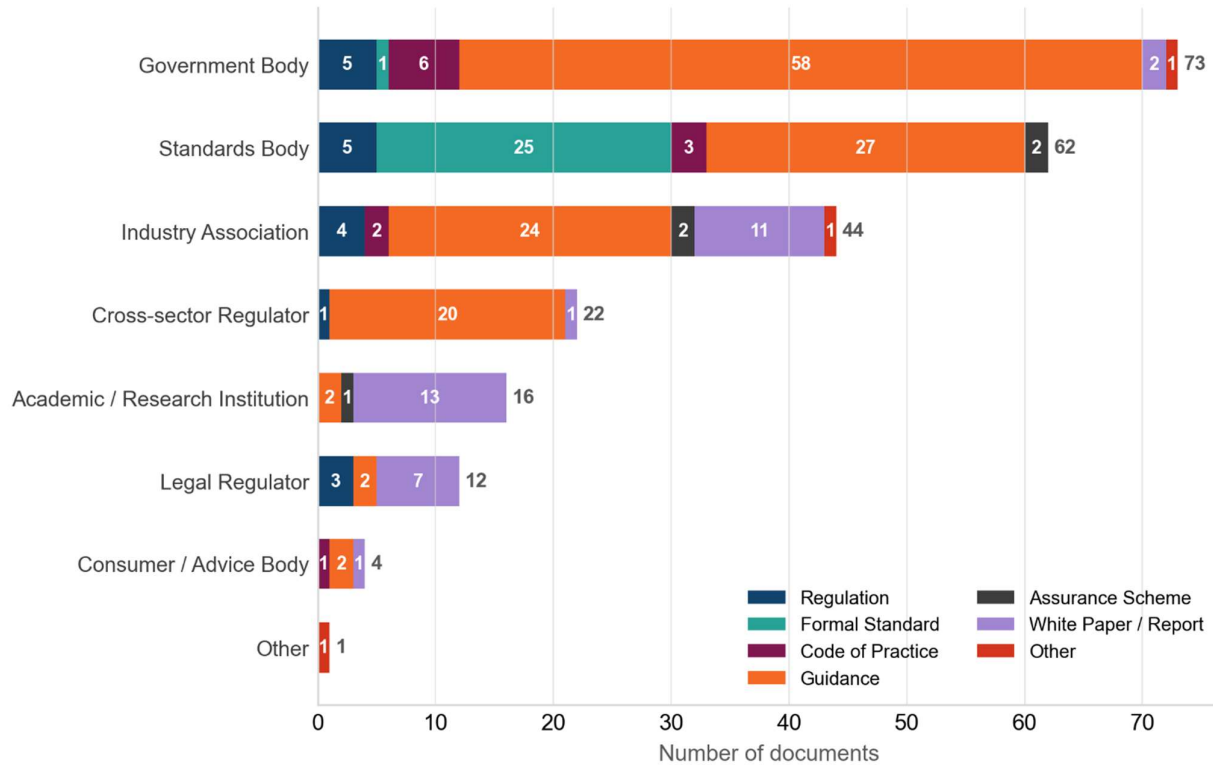
By **document type**, most (135) are guidance documents (58% of the total). White papers and reports account for a substantial share (35 documents, 15%), followed by formal standards (26 documents or 11%), regulations (18 or 8%), codes of practice (12 or 5%), assurance schemes (5 or 2%), and a small number of other document types (3 or 1%). This distribution is significant: it indicates that the primary mode of governance for AI (including AI applied to legal services) is non-binding guidance. While guidance can influence practice and shape regulatory expectations, it lacks the enforceability of regulation and the procedural legitimacy of consensus-based formal standards.

By **organisation type**, government bodies are the most prolific contributors (73 documents or 31%), followed by standards bodies (62 or 26%) and industry associations (44 or 19%). Cross-sector regulators contribute 22 documents (9%), academic and research institutions 16 (7%), and legal regulators 12 (5%). Consumer and advice bodies account for 4 documents (2%).

The relatively low contribution from legal regulators is a finding of material significance but to be expected. This is primarily explained by the fact that they regulate people (legal professionals) and legal firms and not unregulated legal services and products. Much of the applicable governance framework is shaped by cross-sector regulators and industry bodies whose instruments are not tailored to the particular characteristics of legal services delivery.

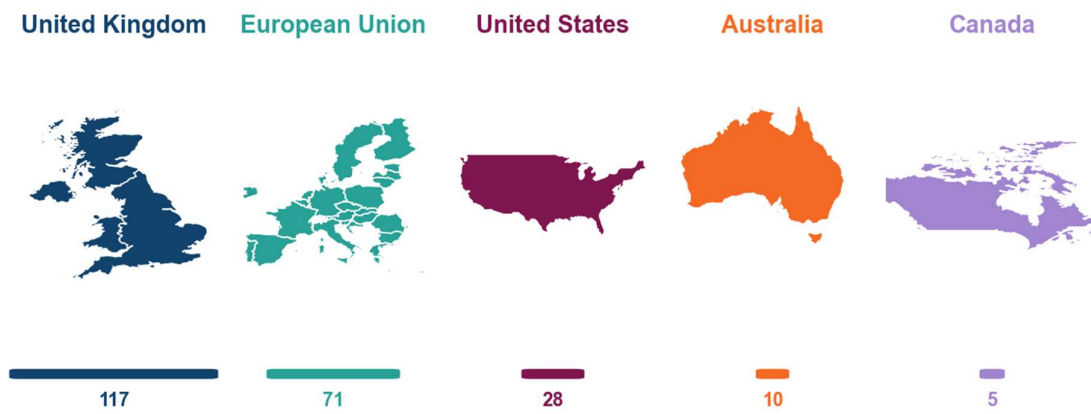
See Figure 2 for an overview of documents reviewed by type and source organisation.

Figure 2: Documents by Type and Source Organisation



By **jurisdiction**, the United Kingdom accounts for 117 documents (50%), the European Union for 71 (30%), the United States for 28 (12%), Australia for 10 (4%), Canada for 5 (2%), and others for 3 (1%) (see Figure 3). The UK's leading position reflects both the density of its regulatory infrastructure and the active engagement of bodies such as the ICO, CMA, Digital Regulation Cooperation Forum (DRCF), and DSIT in AI governance. The EU's substantial contribution is driven largely by the AI Act and its associated implementing measures, technical standards requests, and guidance documents. The US contribution centres on NIST's AI Risk Management Framework and guidance from the Federal Trade Commission, the White House Office of Science and Technology Policy, and state-level initiatives.

Figure 3: Geographic Distribution of Sources



Note: Singapore (1) and International (3) not shown.

3.2 Who Sets Standards and How

Key finding: Standard-setting is fragmented across international bodies, UK cross-sector regulators, legal services regulators, industry associations, and AI platform providers – each operating through different processes with different enforcement powers. Legal services regulators have been slower than their counterparts in financial services and healthcare to address AI governance specifically.

The standard-setting landscape for AI-powered B2C lawtech is very mixed. It is made up of different organisations and bodies operating at different scales, with different mandates, and through different mechanisms. Understanding who sets standards (and the authority and reach of their instruments) is essential to understanding why B2C lawtech is not regulated by the key legal regulators.

International standards bodies operate through consensus-based processes involving national member bodies and produce documents intended for global application. ISO/IEC JTC 1/SC 42 (Artificial Intelligence) has developed a growing portfolio of AI standards, including ISO/IEC 42001:2023 on AI management systems^{viii}, ISO/IEC 23894:2023 on AI risk management^{ix} and the ISO/IEC 5338 series on AI system lifecycle processes^x. IEEE's contributions include the IEEE 7000 series on ethical considerations in system design^{xi}. NIST, while a US national body, exercises global influence through instruments such as the AI Risk Management Framework (AI RMF 1.0, 2023)^{xii} and the companion AI 600-1 profile for generative AI.^{xiii} These formal standards provide architectural frameworks for AI governance but typically require contextualisation and supplementation for specific sectors and use cases.

UK Government and cross-sector regulators have been active in articulating AI governance expectations within the Government's principles-based framework. DSIT's AI Regulation White Paper (2023) established five cross-cutting principles (safety, security, and robustness; appropriate transparency and explainability; fairness; accountability and governance; and contestability and redress) to be implemented by existing regulators within their domains.^{xiv} The Digital Regulation Cooperation Forum (DRCF), comprising the ICO, CMA, Ofcom, and the FCA, has produced joint guidance and a multi-regulator AI toolkit. The ICO's guidance on AI and data protection is among the most operationally detailed instruments in the body of documents, addressing lawful bases for processing, data protection impact assessments, and individual rights in the context of AI decision-making.^{xv} The FCA's framework for AI in financial services provides a comparator that highlights the relative underdevelopment of legal-services-specific guidance.^{xvi}

Legal services regulators have begun to address AI however, there is an important structural limitation in the legal services landscape where it is professionals and firms that are regulated rather than the technology itself. Nonetheless, the LSB's guidance promotes the use of technology and innovation to help address unmet legal need and states that regulators should be proactively engaged in fostering a regulatory environment that encourages technological and innovative solutions to meet consumer need. The LSB also recognises that legal regulators regulate different professions, reserved legal activities and authorised persons (including both individuals and firms/entities), and therefore the different regulators may adopt different approaches when considering outcomes.^{xvii}

Industry bodies contribute a substantial volume of guidance. The Law Society's practice notes, the CCBE's guidance for European lawyers, and LawTechUK's publications articulate professional expectations and best practice recommendations.^{xviii} These instruments are influential within the regulated profession but have limited reach to unregulated lawtech providers, particularly technology companies without a legal services heritage.


Consumer organisations such as Citizens Advice, the Legal Services Consumer Panel and Which? contribute a user-centred perspective that is otherwise underrepresented in the standards landscape

– foregrounding consumer experience, vulnerability, and redress, themes that receive comparatively thin treatment in product-focused standards.

AI platform providers warrant particular attention as de facto standard-setters. OpenAI's usage policies, Google's responsible AI principles, and Microsoft's responsible AI standards shape the behaviour of foundation models upon which many B2C lawtech applications are built. For example, if OpenAI prohibits the use of its models for providing specific legal advice, or if Google implements safety filters that constrain outputs on sensitive topics. These decisions then function as governance mechanisms with material effects on the B2C lawtech landscape. Yet platform policies are set unilaterally, may be changed without notice, and offer no formal accountability or redress mechanisms for affected users or downstream developers.

Figure 4 illustrates the extent to which the instruments identified are binding or mandatory, recommended, voluntary, or of uncertain status. What is significant, is that for over two-thirds (68%) of the documents, the status of the instrument was unknown so it cannot be ascertained if they are binding or unbinding. Together, figures 2 and 4 highlight both the diversity of actors involved in standard-setting and the predominance of non-binding guidance within the landscape.

Figure 4: Binding Status of Standards Sources



3.3 AI-Specificity

Key finding: The apparent breadth of the standards landscape overstates the depth of AI-specific governance. Most standards relevant to B2C lawtech were not designed with AI in mind; they are general professional, data protection, or consumer protection obligations that apply to AI incidentally. When filtered for genuinely AI-specific provisions, coverage thins considerably.

Most standards coverage relevant to AI-powered B2C lawtech derives from general professional and regulatory obligations rather than from AI-specific instruments. Data protection standards, for example, apply to AI systems by virtue of them processing personal data. Not because they were designed to address the particular data protection challenges of large language models (such as the inability to delete training data or the risk of memorisation and regurgitation of personal information). Professional conduct obligations requiring competence, integrity, and client care apply to regulated professionals who use AI tools. However, they do not address the governance of the tool itself when deployed without professional intermediation.

The consequence is that the apparent breadth of the standards landscape overstates the depth of AI-specific governance. When the documents are filtered for genuinely AI-specific provisions, the coverage thins considerably, particularly in the legal services domain. The EU AI Act, NIST AI RMF, and the ISO/IEC 42000 series standards are among the few instruments that were designed from inception to address AI-specific governance challenges. Even these, however, require sector-specific interpretation and supplementation to address the particular context of B2C legal services delivery.


3.4 Overall Thematic Coverage

Key finding: Coverage is markedly uneven. Data protection and transparency are well served by established regulatory frameworks. Vulnerability identification, signposting and referral, and redress mechanisms receive the weakest coverage — precisely the areas most directly relevant to protecting consumers who use AI tools without professional involvement.

Analysis of thematic coverage across the documents reveals a markedly uneven pattern: some themes are addressed extensively while others receive little substantive treatment. The analysis suggests that existing standards are strongest in areas already served by established regulatory frameworks (e.g. transparency) and weakest in areas that require novel governance solutions tailored to the specific characteristics of AI-powered direct-to-consumer services (vulnerability, signposting, redress).

Figure 5 shows the distribution of thematic coverage across the documents and, specifically, the number of documents providing coverage of each theme.

Figure 5: Theme Coverage Across the Standards Landscape



Transparency emerges as the theme with the broadest coverage, receiving treatment in the majority of documents. Transparency has received sustained attention in AI governance discourse, driven by concerns about the opacity of machine learning systems and reinforced by specific transparency obligations in the EU AI Act (notably Articles 13 and 50) and in the UK Government's cross-cutting AI principles.

Accountability and **quality assurance** receive strong coverage, reflecting the emphasis in AI governance frameworks on establishing clear responsibility for AI system outcomes and implementing systematic processes for monitoring and maintaining output quality.

Bias, fairness, and non-discrimination receive attention in a significant number of documents, though the coverage is frequently abstract (i.e. articulating the principle that AI systems should not discriminate) without providing operational detail (e.g. testing protocols, fairness metrics, monitoring requirements) that would enable practical implementation.

Data protection, while benefitting from the mature and well-understood regulatory framework established by the UK GDPR and the EU General Data Protection Regulation, which impose specific obligations on any system processing personal data, receives coverage in around a third of documents.

Vulnerability identification receives notably weak coverage, mentioned in around one-third of all documents. Few documents (around 10% of the total) provide substantive guidance on how AI systems should identify that a user may be in vulnerable circumstances (whether through the nature of their query, their interaction patterns, or explicit disclosure) and fewer still specify how the system should adapt its behaviour in response.

Signposting and referral (the mechanisms by which an AI tool directs users to appropriate human assistance when the tool reaches the limits of its competence or the user's situation requires professional intervention) is similarly underdeveloped. While several documents reference the general principle that AI systems should know their limitations, few specify how this should be operationalised in practice or: what triggers a referral, where the user is directed, how the handover is managed, and who bears responsibility for ensuring the user reaches appropriate help. Even fewer documents explicitly recognise that these limitations should extend to the identification of and response to vulnerable users.

Redress and complaints mechanisms also receive limited treatment, particularly in the context of unregulated providers. Existing redress frameworks in legal services are designed for complaints about regulated professionals and entities. They do not extend to unregulated B2C lawtech providers, leaving consumers of the majority of unregulated AI-powered legal tools without a clear route to challenge outcomes or obtain remedies when things go wrong.^{xix}


4. Thematic Analysis

4.1 Output Quality

Key finding: No legal services standard specifies consistency benchmarks, reliability testing protocols, accuracy validation requirements, or hallucination detection obligations for AI tools. Consumer-facing lawtech operates with no mandated quality assurance requirements. In sharp contrast to healthcare where pre-market validation and post-deployment monitoring are mandatory.

The three themes that fall within the domain of output quality (consistency, reliability, and accuracy) address whether AI systems produce dependable results. Figure 6 shows the number of documents addressing consistency, reliability, and accuracy across different source organisation types.

Figure 6: Theme Coverage by Organisation Type – Output Quality



4.1.1 Consistency

Consistency (the expectation that an AI system produces comparable outputs under comparable conditions) receives markedly uneven attention. As there is no product regulator on legal services, this is not surprising. Healthcare and medical device regulation requires rigorous treatment. The Medicines and Healthcare products Regulatory Agency (MHRA) has published Software and AI as a Medical Device guidance, which specifies that there must be pre-market testing and ongoing surveillance (MHRA, 2023). The US FDA's draft guidance on AI-enabled devices similarly guides developers to demonstrate consistent performance across intended patient populations and use conditions, and recommends post-market surveillance (US FDA, 2025).

International standards bodies provide structured frameworks. The EU AI Act requires that high-risk AI systems achieve appropriate levels of accuracy, robustness and cybersecurity throughout their lifecycle, with conformity assessments to verify these properties prior to market placement (European Commission, 2024, Art. 15 and Art. 43). ISO/IEC 25010 and the associated SQuARE quality model establish measurable quality characteristics (including functional correctness, maturity, and fault tolerance) against which system consistency can be evaluated (ISO/IEC, 2023).

Legal services guidance treats consistency as a matter of professional diligence rather than system-level assurance. The SRA guidance instructs solicitors to verify AI outputs but specifies no consistency benchmarks, testing protocols, or monitoring requirements (SRA, 2023; 2026). The Law Society's guidance similarly warns of variability (i.e. unreliability or hallucinations) in AI outputs without prescribing how firms should measure or manage it (The Law Society, 2025). Apart from the minority of B2C lawtech products that have deliberately put themselves under the regulatory umbrella because they have been developed by regulated legal professionals (for example, Garfield AI), consumer-facing lawtech products operate with few consistency assurance requirements: there is no specific requirement to demonstrate that the same query produces stable results over time. This stands in sharp contrast to the healthcare sector, where the consequence of inconsistent outputs is treated as a system design problem rather than an individual clinician's responsibility.

4.1.2 Reliability

Consistency (subsection 4.1.1) is concerned with whether the system produces comparable outputs under comparable conditions, that is to say, the same question in similar circumstances produces similar answers. Reliability is whether the system performs as intended over time and under varying conditions. This includes its availability, fault tolerance, graceful degradation under stress, and resilience to adversarial inputs. Consistency is one dimension of reliability, but reliability also encompasses system uptime, performance under load, and recovery from errors.

Reliability (the capacity of an AI system to perform as intended under specified conditions over time) is accepted as a governing principle but is rarely operationalised with technical specificity. That is, it is seldom translated into, for example, defined performance metrics, testing protocols, validation procedures, or quantitative thresholds for acceptable error rates, robustness, or system failure.

ISO/IEC standards provide one of the most structured approaches, distinguishing between repeatability (consistent results under identical conditions), reproducibility (consistent results across environments), and resilience (graceful degradation under stress or adversarial conditions). ISO/IEC 25010 further decomposes reliability into maturity, availability, fault tolerance, and recoverability sub-characteristics, each amenable to quantitative measurement (ISO/IEC, 2023). The NIST AI Risk Management Framework identifies reliability as a core trustworthiness function, linking it to validity and robustness, though its guidance remains voluntary and non-prescriptive (NIST, 2023).

The EU AI Act imposes binding reliability requirements for high-risk systems. Providers must demonstrate that systems are resilient against errors, faults or inconsistencies and must implement risk management processes addressing foreseeable misuse and reasonably foreseeable adverse effects (European Commission, 2024, Art. 9, 15). Healthcare regulators give reliability operational teeth: the FDA mandates pre-market validation of analytical and clinical performance with ongoing post-market surveillance (US FDA, 2025).

Financial services regulators address reliability through outcome-focused obligations. The FCA Consumer Duty requires that products and services work as consumers would reasonably expect, which implicitly demands reliable system performance, though the regulatory framework does not prescribe technical reliability testing (FCA, 2024).

Legal services guidance addresses reliability only at the level of general principle. The SRA instructs solicitors to check and verify AI-generated outputs, a requirement that acknowledges the possibility of unreliable outputs but assigns the detection burden entirely to the practitioner (SRA, 2026). No legal services regulator specifies acceptable error rates, testing cadences, resilience requirements, or performance degradation thresholds. The gap is particularly consequential for consumer-facing AI

legal tools, where no professional intermediary stands between the system's output and the user's reliance on it.

4.1.3 Accuracy

Accuracy (the degree to which AI outputs correspond to correct or true values) is widely acknowledged as essential but rarely operationalised for legal services contexts.

The ICO provides one of the most developed UK regulatory treatments, indicating that organisations using AI in automated decision-making to implement statistical accuracy testing and ongoing performance monitoring (ICO, 2023). The FDA mandates pre-market demonstration of both analytical validity (does the system measure what it claims to measure?) and clinical validity (do outputs correspond to real clinical outcomes?). Creating a two-stage accuracy assurance model with no parallel in legal services (US FDA, 2025).

Several critical accuracy dimensions are not addressed across the standards landscape. No framework distinguishes between factual accuracy (correct identification of dates, names, and events) and legal accuracy (correct interpretation of statutes, case law, and regulatory requirements). Despite these being fundamentally different properties requiring different validation approaches. No standard requires confidence scoring that communicates to users the system's assessed certainty in its own outputs. Despite this being technically feasible and widely discussed in the AI research literature. The question of input quality is also neglected: accuracy depends not only on system capabilities but on the quality of information users provide, yet no standard addresses how AI tools should manage ambiguous, incomplete, or contradictory user inputs.

Legal services guidance instructs practitioners to "check the accuracy of AI-produced outputs" (SRA, 2026) but provides no framework for what constitutes acceptable accuracy in legal reasoning tasks. Byrom (2024) identifies case outcome prediction technologies as warranting particular regulatory attention, noting that accuracy claims in this domain are difficult to verify and potentially misleading. For consumer-facing lawtech, the gap is stark: there is no legal-services specific requirement for pre-deployment accuracy testing, no obligation to disclose hallucination rates, and no standard against which accuracy claims could be assessed. The cross-sector regulatory guidance (most notably from the ICO) does impose accuracy-related obligations in the context of automated decision-making. However, these are limited in scope, non-sector-specific, and do not establish operational benchmarks for legal reasoning tasks.


4.2 Fairness and Inclusion

Key finding: Bias and fairness are widely acknowledged but inconsistently operationalised. Financial services has developed sector-specific fairness metrics and testing methodologies; legal services has not. No legal services regulator has defined what fairness means in the context of AI-assisted legal work.

The four themes in this domain (bias, fairness, non-discrimination, and accessibility) address whether AI systems treat all users equitably and serve diverse populations effectively. These themes overlap conceptually, and the analysis that follows treats them as distinct but cross-referenced. First, bias concerns the technical sources of systematic error; secondly, fairness concerns the principles and metrics governing equitable treatment; thirdly, non-discrimination concerns legal prohibitions on differential treatment based on protected characteristics; and lastly, accessibility concerns practical barriers to effective use.

Figure 7 shows the number of documents addressing bias, fairness, non-discrimination, and accessibility across different source organisation types.

Figure 7: Theme Coverage by Organisation Type – Fairness and Inclusion



4.2.1 Bias

Bias is the most extensively addressed theme of the four in this domain, reflecting widespread recognition that AI systems can encode and amplify systematic errors from training data, algorithmic design, and deployment contexts.

ISO/IEC TR 24027 provides a comprehensive technical treatment, organising sources of unwanted bias into three high-level groups: human cognitive biases (including automation bias, confirmation bias, in-group bias and societal bias), data bias (including statistical bias, non-representative sampling, missing features and labels, and data aggregation), and bias introduced by engineering decisions (including feature engineering, algorithm selection, hyperparameter tuning and model bias) (ISO/IEC, 2021). A more granular taxonomy of bias sources across the machine learning lifecycle (distinguishing historical bias, representation bias, measurement bias, aggregation bias and evaluation bias) has been developed in the academic literature^{xx}. The NIST (2022) special publication on AI bias complements these frameworks with a practical detection and mitigation framework, identifying computational, statistical, and human-cognitive dimensions of bias and providing structured approaches to measurement.

The FCA has developed one of the most operationally detailed bias analysis within a sectoral regulatory context. Research by Daniel Bogiatzis-Gibbons and colleagues examines how different fairness metrics interact and conflict in financial services contexts. Demonstrating, for instance, that optimising for demographic parity may worsen calibration, and vice versa (Bogiatzis-Gibbons et al., 2024). This level of sector-specific operationalisation has no equivalent in legal services.

The EU AI Act requires providers of high-risk AI systems to examine training, validation, and testing datasets for biases that may affect the health and safety of persons or lead to discrimination prohibited by Union law (European Commission, 2024, Art. 10). This creates a binding obligation to identify and mitigate bias. Though the Act provides limited guidance on how examination should be conducted for generative AI systems where training datasets may be opaque even to providers.

Legal services guidance warns practitioners about bias risks regarding AI in general terms. The Law Society advises that AI tools "may reflect biases present in their training data" (The Law Society, 2025); the SRA similarly instructs solicitors to be "aware of the risk of bias" (SRA, 2026). Neither

source specifies testing methodologies, acceptable bias thresholds, monitoring requirements, or mitigation strategies. There is very little guidance that addresses bias specifically in the context of generative AI or large language model outputs, despite these being the tools most rapidly adopted across the legal sector.

4.2.2 Fairness

Fairness is universally endorsed as a governing principle for AI systems but inconsistently operationalised across sectors and standard setters.

The ICO provides a particularly instructive treatment by distinguishing two dimensions of fairness in the AI context: statistical fairness (whether system outputs are equitable across demographic groups according to specified mathematical criteria) and broader GDPR fairness (whether data processing respects individuals' reasonable expectations and is not unduly detrimental). These dimensions can conflict: a system may satisfy a chosen statistical fairness metric while still processing data in ways that individuals would reasonably consider unfair, and vice versa (ICO, 2023). This distinction is relevant to legal services AI, where fairness encompasses both the statistical properties of system outputs and the broader question of whether consumers are treated in ways consistent with their legitimate expectations.

Financial services regulation has advanced furthest in operationalising fairness requirements. The FCA Consumer Duty requires firms to deliver "good outcomes" for retail customers across four dimensions (products and services, price and value, consumer understanding, and consumer support). It treats AI-driven processes as subject to the same outcome obligations as human-delivered services (FCA, 2024). This outcome-focused approach provides a regulatory model that does not depend on prescribing technical fairness metrics but instead holds firms accountable for demonstrable results.

The EU AI Act requires high-risk AI providers to implement bias examination and mitigation measures (as discussed in subsection 5.2.1 above). It does not though specify which fairness metrics should be applied or how conflicts between competing fairness definitions should be resolved.

Legal services bodies acknowledge fairness as a value but provides little operational detail. No legal services regulator has defined what fairness means in the context of AI-assisted legal work, what metrics might be applied, or how fairness should be tested and monitored. The gap between the financial services sector's maturity (where regulators have engaged substantively with fairness operationalisation) and legal services relative immaturity, where fairness remains an abstract principle, is among the most significant findings of this review.

4.2.3 Non-Discrimination in AI

Non-discrimination is extensively addressed in international and cross-sectoral frameworks but thinly treated in legal-services-specific guidance around AI (though not surprising, as noted before, consumer facing lawtech products are not regulated by the legal regulators).

The EU AI Act classifies AI systems used in the "administration of justice and democratic processes" as high-risk. Subjecting them to the full suite of obligations including bias examination, conformity assessment, and post-market monitoring (European Commission, 2024, Annex III). This classification captures some judicial AI applications but its applicability to broader legal services AI (including consumer-facing tools, document review systems, and legal research platforms) remains subject to interpretation. The EU Fundamental Rights Agency provides a detailed analysis of how AI systems can produce discriminatory outcomes. Through proxy variables, feedback loops, and statistical patterns that correlate with protected characteristics without explicitly referencing them (EU Fundamental Rights Agency, 2022).

The ICO distinguishes direct discrimination (explicit use of protected characteristics as input features) from indirect discrimination (use of proxy variables that correlate with protected characteristics). Noting that the latter is both more common in AI systems and more difficult to detect (ICO, 2023). Financial services regulators have developed one of the most operationally specific anti-discrimination frameworks. The FCA examines how algorithmic decision-making may produce discriminatory outcomes in lending and insurance. While the National Association of Insurance Commissioners provides model guidance on preventing unfair discrimination in AI-driven underwriting and pricing (NAIC, 2024).

Several gaps warrant attention in the context of AI-enabled consumer-facing legal tools. Frameworks treat socioeconomic status inconsistently in discrimination analysis: some treat it as analogous to a protected characteristic, while others do not address it at all. Guidance on intersectional discrimination (the compounding effects of membership in multiple disadvantaged groups) is largely absent. No framework provides distinct treatment for discrimination risks in consumer-facing AI as opposed to professional-intermediated AI, despite the different risk profiles these contexts present. For legal services specifically, the combination of thin sector-specific guidance and the increasing deployment of AI tools that directly serve consumers creates a regulatory gap: the general anti-discrimination frameworks exist but have not been operationalised for the legal services context.

4.2.4 Accessibility and Effective Use

Accessibility and effective use (the extent to which AI systems can be meaningfully accessed and used by diverse populations including those with disabilities, low literacy, or limited digital capability) receives uneven and frequently superficial treatment across the standards landscape.

Most frameworks that address accessibility focus on the output side: ensuring explanations are provided in plain language, that information is layered to accommodate different levels of expertise, and that visual aids supplement textual content. The ICO and Alan Turing Institute guidance on explaining AI decisions provides one of the most detailed sets of practical recommendations on this dimension, proposing multiple explanation formats calibrated to different audience needs (ICO and Alan Turing Institute, 2020). Consumers International calls for AI systems to be accessible to all consumers, including those in vulnerable circumstances, and advocates for universal design principles in AI interfaces (Consumers International, 2024).

The input side of accessibility receives far less attention. No standard in the documents addresses how AI legal tools should accommodate users with low literacy levels, cognitive disabilities, or limited English when those users need to provide information to the system. This is a material gap for consumer-facing lawtech. Where the quality of AI outputs depends directly on the quality of user inputs: a tool that requires precise legal terminology, structured factual narratives, or familiarity with legal concepts will produce poor results for users who cannot meet these implicit requirements, regardless of the system's technical capabilities. A related challenge is that consumers may not know which facts are legally relevant to their problem. They therefore may omit information that is critical to the accuracy of the system's response.

Web Content Accessibility Guidelines (WCAG) standards exist for digital accessibility but were developed for conventional web interfaces. They are not adapted for AI interaction patterns such as conversational interfaces, dynamic questioning, or iterative refinement. The Legal Services Consumer Panel have published wider guidance for accessible service delivery.^{xxi} It should be noted that all legal services users may be considered vulnerable to a greater or lesser extent, given the knowledge imbalance between them and a legal professional and the very reason why they are engaging with a legal professional. Many legal services users may face challenges in interacting effectively with AI tools. Particularly where systems rely on clear inputs or familiarity with legal concepts.

The accessibility gap intersects with the vulnerability identification gap discussed in subsection 5.3.3 below. Arguably, users who may often also be those most in need of the services AI tools provide are least able to interact effectively with AI tools. They may typically be socially and economically less likely to be able to access paid for legal advice.

4.3 User Protection

Key finding: The four user protection themes – transparency, signposting, vulnerability identification, and human oversight – are interconnected but unevenly developed. Transparency is the most extensively addressed theme across the entire landscape and vulnerability identification the weakest.

Four themes in this domain address how AI systems protect the interests of people who interact with them. These are:


- transparency;
- signposting and referral;
- vulnerability identification; and
- human oversight)

These themes are interconnected:

- transparency provides the information basis for informed decision-making;
- signposting directs users to appropriate support;
- vulnerability identification identifies users who need additional protection; and
- human oversight ensures that consequential decisions are not delegated entirely to automated systems.

Figure 8 shows the number of documents addressing transparency, signposting and referral, vulnerability identification, and human oversight across different source organisation types.

Figure 8: Theme Coverage by Organisation Type – User Protection



4.3.1 Transparency

Transparency is the most extensively addressed theme across the entire standards landscape, appearing in substantive form in the majority of sources reviewed.

The EU AI Act establishes one of the most structurally detailed transparency regimes. Article 50 mandates that providers of AI systems designed to interact with natural persons ensure that individuals are informed they are interacting with an AI system. Unless this is "obvious from the point of view of a reasonably well-informed, observant and circumspect natural person." For high-risk systems, the Act further requires disclosure of accuracy levels, known limitations, and the categories of natural persons and scenarios for which the system is intended (European Commission, 2024). The National Telecommunications and Information Administration characterises the relationship between transparency and accountability as constitutive: transparency enables accountability, and accountability requires transparency (NTIA, 2024). AlgorithmWatch (2025) documents how algorithmic transparency is emerging as a distinct governance domain with its own institutional infrastructure, reporting frameworks, and civil society oversight mechanisms.

The ICO and Alan Turing Institute guidance provides one of the most practically detailed frameworks, identifying six types of explanation that organisations should consider providing:

- rationale explanations (why a particular decision was reached);
- responsibility explanations (who is accountable);
- data explanations (what data was used);
- fairness explanations (what steps were taken to ensure equitable treatment);
- safety and performance explanations (what steps were taken to ensure reliability); and,
- impact explanations (what effect the decision has on the individual) (ICO and Alan Turing Institute, 2020).

This multi-dimensional approach recognises that transparency is not a single disclosure but a structured communication challenge.

Legal services guidance advises disclosure of AI use but does not specify the minimum content, format, or timing of such disclosure (SRA, 2026; The Law Society, 2025; CCBE, 2025). Three gaps are particularly significant. First, there is no defined minimum disclosure content for consumer-facing legal AI – for instance, what a lawtech tool should tell users about how it works, what data it uses, and what its limitations are. Second, no guidance addresses the communication of uncertainty or confidence levels, a critical omission for legal AI, where users may reasonably but incorrectly assume that system outputs are definitive. Third, the distinction between transparency to professionals (who have the training to interpret technical disclosures) and transparency to consumers (who typically do not) is not set out sufficiently.

4.3.2 Signposting and Referral

Signposting and referral (the mechanisms by which AI systems direct users to appropriate human assistance, alternative services, or relevant information) is acknowledged across the standards landscape but rarely developed with operational specificity i.e. what this means in practice.

Most frameworks do not emphasize signposting and treat it obliquely, as a subordinate component of transparency or human oversight obligations rather than as a distinct regulatory requirement. This is a missed opportunity: signposting serves a different function from transparency (it directs action rather than providing information) and from human oversight (it operates at the user-system interface rather than the organisational governance level).

Healthcare regulation is most developed in this area. The WHO guidance on the ethics and governance of AI for health emphasises human control and the need to identify situations where human assistance is needed and communicate this to users in a timely and actionable manner (WHO, 2021). The Royal Australian College of General Practitioners (RACGP) requires that GPs check and verify all AI-generated outputs and stipulates that AI should never be used as the sole source for clinical decision-making, placing responsibility on the clinician to identify when AI assistance is insufficient and direct professional involvement is needed (RACGP, 2025). These healthcare models demonstrate that signposting can be operationalised through defined communication standards and human-based escalation protocols.

Three gaps are particularly consequential for legal services. First, no standard provides a methodology for communicating AI confidence levels to lay users in a way that enables informed decisions about whether to proceed with AI-generated guidance or seek professional assistance. Second, there is no consistent standard for what should trigger a referral to a human professional. Whether based on matter complexity, potential consequences, user characteristics, or system confidence levels. Third, signposting to complaints mechanisms and regulatory bodies is almost entirely absent from the standards reviewed: users who receive poor outcomes from AI legal tools are given no systematic direction toward redress. Interview evidence underscores this concern: as one stakeholder observed, "any general AI tool should have to tell people upfront: I'm not a lawyer" (Consumer/Advice Body interview, 2026). The principle is sound; the challenge is translating it into operationally specific requirements that go beyond a single disclaimer.

4.3.3 Identifying vulnerability users

The weakest area across the entire standards landscape is on AI systems inferring a user's ability to protect their own interests. No existing standard provides a workable framework for detecting vulnerable users in real time during AI interactions.

Standards acknowledge that vulnerability matters but stop well short of operationalisation into practical actions. The EU AI Act prohibits AI systems that "exploit vulnerabilities" of specific groups due to age, disability, or social or economic situation, but provides no detection methodology. The prohibition assumes vulnerability has already been identified (European Commission, 2024). Consumers International calls for particular attention to consumers in vulnerable circumstances but does not specify how AI systems should identify such circumstances (Consumers International, 2024).

Healthcare offers one of the most developed approaches, though these are designed for clinical rather than legal contexts. The WHO (2021) guidance suggests assessing whether AI tools are appropriate for particular patient populations, including consideration of digital literacy, cognitive capacity, and situational factors. The RACGP specifies that clinicians should assess patient suitability for AI-assisted pathways before directing them to such pathways (RACGP, 2025). These models rely on professional intermediaries conducting vulnerability assessments, a model that does not translate to consumer-facing AI legal tools where no professional intermediary is present.

The gap is particularly acute for legal services. People who access legal services often do so in circumstances of distress or urgency.^{xxii} An additional source of increased risk of vulnerability at this time is the knowledge imbalance between the client and professional (similar to what is seen in medicine). Another is the simple cognitive overload of managing new information and the enormity of the issue while still dealing with everything else in daily life. This is the case for even positive legal situations such as conveyancing as well as more negative situations such as eviction, relationship breakdown, involvement with the criminal justice system, or debt or welfare crisis. These circumstances impair people's decision-making capacity precisely when the stakes of poor decisions are highest. Interview evidence from a lawtech provider is candid about the current state of play: "we

don't have a good way of identifying vulnerability in real time" and existing approaches rely on "self-declaration, which is the worst possible mechanism for the most vulnerable users" (Lawtech Company interview, 2026).

The absence of any standard addressing this gap represents perhaps the most significant finding of this review for the LSB's regulatory consideration. Even at the level of principles, similar to the Design Council's principles for inclusive design or research priorities.

4.3.4 Human Oversight

Human oversight (the requirement that a real person can effectively supervise, intervene in, and override AI system outputs) is among the most commonly addressed themes in the standards landscape.

The EU AI Act provides the strongest binding framework. Article 14 mandates that high-risk AI systems be "designed and developed in such a way that they can be effectively overseen by natural persons". With overseers empowered to fully understand the capacities and limitations of the system, correctly interpret its outputs, decide not to use the system. Also, to "intervene on the operation of the system or interrupt it" (European Commission, 2024). UK data protection law provides a complementary mechanism through Article 22 of the UK GDPR. This grants individuals the right not to be subject to decisions based solely on automated processing that produce legal effects or similarly significant effects, with the right to obtain human intervention (European Parliament and Council, 2016). The Dutch judiciary provides a particularly instructive model of principled limitation, restricting AI to well-defined, auditable tasks such as anonymisation and explicitly excluding AI from judicial reasoning and decision-making (De Rechtspraak, 2024).

Healthcare and financial services regulators have translated oversight principles into operationally specific requirements. MHRA guidance indicates that AI medical devices should provide clear information about intended use, the role of the user, and how the device should be used so that users can exercise appropriate oversight. (MHRA, 2024). The FCA asks firms to maintain human oversight over algorithmic decision-making affecting consumers.


Some legal services bodies (the SRA, the Law Society, and the CCBE) note that the practitioner remains responsible for checking and verifying AI outputs (SRA, 2026; The Law Society, 2025; CCBE, 2025). This position is sound in principle but leaves several operational questions unresolved. No standard specifies what qualifications or training overseers need to exercise effective oversight of AI legal tools. There is no consensus on review frequency or intensity: should every output be reviewed, or is sampling acceptable? Most significantly, the assumption of a professional intermediary breaks down entirely for consumer-facing lawtech. Where the "overseer" is the consumer themselves, a person who, by definition, lacks the legal expertise needed to assess whether the system's output is correct.

4.4 Governance and Assurance

Key finding: Quality assurance and accountability have clear principles but unresolved operational questions. Redress and traceability are among the thinnest areas. The accountability gap is sharpest for unregulated consumer-facing tools where no professional-client relationship exists.

The four themes in this domain (quality assurance, accountability, redress and complaints, and traceability) address the institutional and procedural mechanisms through which AI systems are governed and through which harms can be identified, attributed, and remedied. Figure 9 shows the number of documents addressing quality assurance, accountability, redress and complaints, and traceability across different source organisation types.

Figure 9: Theme Coverage by Organisation Type – Governance and Assurance



4.4.1 Quality Assurance

Quality assurance (the systematic processes for ensuring AI systems meet defined performance standards) is extensively addressed. The most granular frameworks emerge from ISO/IEC standards and healthcare regulators.

ISO/IEC provides one of the most structured technical approaches. The SQuaRE quality model (ISO/IEC 25010) defines measurable quality characteristics across eight dimensions: functional suitability, performance efficiency, compatibility, usability, reliability, security, maintainability, and portability (ISO/IEC, 2023). ISO/IEC TR 29119-11 provides AI-specific testing guidelines addressing test design for machine learning systems, including considerations relating to training data, model evaluation, and monitoring across the system lifecycle (ISO/IEC, 2020). These frameworks are technically rigorous but domain-agnostic: they provide the tools for quality assurance without specifying how those tools should be applied in legal services contexts.

Healthcare regulators demonstrate how domain-specific quality assurance operates in practice. The FDA's pre-market review process requires manufacturers to submit evidence of analytical and clinical validity before AI medical devices can be marketed, with ongoing post-market surveillance obligations (US FDA, 2025). The MHRA's approach to conformity assessment for software and AI medical devices increasingly accommodates predetermined change control plans for certain software updates, setting out in advance how specified modifications will be managed and assessed. (MHRA, 2023). The EU AI Act mandates that providers of high-risk systems implement quality management systems covering design specification, development processes, testing procedures, and post-market monitoring (European Commission, 2024, Art. 17).

The legal services sector has no sector-specific quality assurance requirements for AI systems. The SRA and the Law Society frame quality assurance as a dimension of professional responsibility (firms should satisfy themselves that AI tools are fit for purpose) but provide no operational detail on how such assurance should be conducted (SRA, 2026; The Law Society, 2025). Three gaps are particularly significant. First, quality assurance frameworks for generative AI outputs remain underdeveloped across all sectors; the existing frameworks were designed primarily for predictive and

classificatory AI. Second, requirements for independent third-party audit of AI system performance are sparse. Third, no standard provides a methodology for validating AI performance specifically in legal reasoning tasks. This is a domain where "correct" outputs may be contested, context-dependent, and resistant to automated evaluation.

4.4.2 Accountability

Accountability (when the person responsible for the AI system's outcomes is identified clearly as well as the mechanisms to enforce that responsibility) is among the most consistently ignored principles. The EU AI Act creates one of the most structurally detailed accountability architecture. It establishes a chain of responsibility that distinguishes between providers (who develop or place AI systems on the market), deployers (who use AI systems under their authority), importers, and distributors. Each role carries specified obligations, and providers of high-risk systems must demonstrate compliance through conformity assessments before market placement (European Commission, 2024). The UK government's pro-innovation regulatory framework takes a deliberately different approach, delegating accountability for AI governance to existing sectoral regulators rather than creating new cross-cutting obligations, with the expectation that each regulator will apply five core principles (safety, transparency, fairness, accountability, and contestability) through its existing regulatory tools (DSIT, 2024).

Financial services regulation provides a model of individual accountability. The FCA's Senior Managers and Certification Regime holds named individuals personally responsible for the conduct of functions within their remit, including functions that involve or are affected by AI systems (FCA, 2024). This creates a direct line of accountability from organisational AI deployment decisions to identifiable human decision-makers, a feature that distinguishes it from more diffuse accountability frameworks.

In legal services, the SRA, the Law Society, and the CCBE are clear and consistent: the practitioner retains full professional responsibility for work product regardless of whether AI tools were used in its production (SRA, 2026; The Law Society, 2025; CCBE, 2025). This principle is well-established and uncontested for regulated legal practice. The accountability gap emerges for direct-to-consumer AI legal tools where no legal professional intermediates: when an unregulated lawtech product provides incorrect legal information to a consumer, the chain of accountability is unclear. The consumer has no professional-client relationship, may not be able to identify the responsible entity, and may face practical barriers to pursuing a claim. In some cases, consumer-facing AI tools also include prominent disclaimers stating that outputs do not constitute legal advice and that the provider accepts no liability for users' reliance on the information provided. This gap (between clear practitioner/firm accountability within the regulated sphere and diffuse or absent accountability outside it) is perhaps the most pressing issue that needs to be addressed to ensure that consumers using B2C lawtech products are protected.

4.4.3 Redress and Complaints

Redress and complaints (the mechanisms through which individuals can challenge AI-driven outcomes and obtain remedies for harm) is another one of the thinner areas in the standards landscape.

Article 22 of the UK GDPR provides one of the most established individual rights, enabling data subjects to contest decisions based solely on automated processing that produce legal or similarly significant effects. As well as to obtain human intervention, an explanation (Art. 13-15), and the opportunity to challenge the decision (European Parliament and Council, 2016). However, Article 22's coverage weakens in two important respects. First, it applies to "solely automated" decisions, creating uncertainty about its applicability to AI-assisted decisions where a human nominally reviews the output but may lack the capacity or inclination to override it. Second, it was not designed for

generative AI outputs (legal information, document drafts, or guidance generated by large language models) which do not fit neatly into the “decision” framing that Article 22 addresses.

The EU AI Act introduces some additional obligations, including requirements for providers to implement post-market monitoring and to report serious incidents. However, the implementing detail relevant to individual redress remains limited (European Commission, 2024). Consumers International calls for effective, accessible and affordable redress mechanisms for consumers affected by AI systems, noting that existing consumer protection frameworks may not adequately address AI-specific harms (Consumers International, 2024).

Both the legal and financial sectors share the same structural issue: if the provider is unregulated, the consumer does not have access to a redress scheme. In legal services the Legal Ombudsman (LeO) only covers authorised persons. But AI legal B2C tools are unregulated so the consumer has no redress through the LeO. In financial services, the FOS only covers Financial Conduct Authority-regulated firms. If a consumer uses an unauthorised fintech/AI tools they also have no FOS redress pathway. The key difference is the scale of the gap between the two sectors. In financial services, most meaningful consumer-facing activity must be FCA-regulated, so unregulated firms are a smaller part of the landscape. In legal services, the majority of legal issue types fall within unreserved and therefore often unregulated activities. Where this is so the redress gap is much bigger and affects far more people.

Legal services guidance is sparse on what happens when AI tools produce erroneous or harmful outputs for consumers. Within the regulated sphere, the LeO provides an established complaints mechanism for consumers of regulated legal services. Also, the SRA’s own disciplinary processes apply to AI-assisted work products, though their practical adequacy for AI-specific complaints has not been tested. The critical gap is that these mechanisms do not extend to unregulated B2C lawtech providers. For consumers using unregulated lawtech products directly, the position is starkly different. As interview evidence illustrates: “consumers have no idea who to complain to if an AI legal tool gives them bad information” (Consumer/Advice Body interview, 2026). The absence of clear, accessible, and proportionate redress mechanisms for this growing category of AI-legal interaction represents a significant regulatory gap.

4.4.4 Traceability

Traceability (the capacity to reconstruct the inputs, processing steps, and reasoning chains that produced a given AI output) is recognised as foundational to accountability, auditability, and redress, but receives uneven treatment.

The EU AI Act provides the strongest binding requirements. Article 12 mandates that high-risk AI systems incorporate automatic event logging capabilities that record system operation over the system’s lifetime. This includes input data, system actions, and output data, at a level of detail sufficient to enable monitoring of system operation and post-market surveillance (European Commission, 2024). ISO/IEC 23894 provides structured guidance for AI risk management documentation, including data lineage, model provenance, and decision audit trails (ISO/IEC, 2023).

Healthcare and financial services regulators have the most mature traceability regimes in practice. The FDA requires manufacturers of AI-enabled medical devices to provide documentation and information supporting the evaluation of safety and effectiveness across the device total product life cycle, including risk management, design validation, design changes, and performance monitoring (US FDA, 2025). The MHRA similarly requires manufacturers of software and AI medical devices to produce documented evidence supporting intended use, safety, performance, and regulatory compliance within the UK medical device framework (MHRA, 2023; 2025). The NAIC expects insurers

using AI in underwriting, pricing, claims, and other regulated insurance practices to maintain governance, validation, audit, and documentation records sufficient to support regulatory oversight and to produce them when requested (NAIC, 2023).


Legal services guidance advises practitioners to check sources (particularly to verify that AI-generated legal citations are real) but does not address the technical infrastructure required to make such checking possible. No legal services standard specifies audit trail requirements, log retention periods, reasoning chain documentation, or data lineage records. For consumer-facing AI legal tools, traceability requirements are entirely absent: users have no right to request an account of how a particular output was generated, what data informed it, or what processing steps were involved. The challenge is compounded for large language models, where the relationship between training data and specific outputs is not straightforwardly traceable even in principle. This is a technical limitation that the standards landscape has not yet addressed in operationally useful terms.

4.5 Data and Privacy

Key finding: Data protection is the most thoroughly addressed theme across the landscape, benefiting from the maturity of the GDPR framework. However, several AI-specific questions remain unaddressed for legal services, including training on case data, inferential privacy, and accountability for privacy when unregulated tools serve consumers directly.

The two themes in this domain (privacy and data protection) are closely related but analytically distinct. Privacy concerns the individual's right to control information about themselves and to be free from intrusive processing. While data protection concerns the legal and technical framework governing the collection, processing, storage, and sharing of personal data. Figure 10 shows the number of documents addressing privacy and data protection across different source organisation types.

Figure 10: Theme Coverage by Organisation Type – Data and Privacy



4.5.1 Privacy

Privacy (the protection of individuals' informational autonomy and freedom from intrusive data processing) is extensively addressed across the standards landscape, though with significant gaps in legal services-specific application.

The ICO provides one of the most comprehensive UK regulatory treatments, addressing the full range of privacy considerations relevant to AI: lawful basis for processing; necessity and proportionality assessments; data protection impact assessments (DPIAs); data minimisation; purpose limitation; and, the rights of data subjects including the Article 22 automated decision-making rights discussed in subsection 5.14 above (ICO, 2023). However, the Data (Use and Access) Act 2025 replaced the earlier Article 22 framework with a more permissive regime for significant solely automated decisions, subject to safeguards. The EU AI Act creates parallel privacy-relevant obligations for high-risk AI systems. This includes data governance requirements that mandate attention to the representativeness, accuracy, and completeness of training datasets, and transparency obligations that require disclosure of what data categories are processed (European Commission, 2024).

Healthcare regulators demonstrate more operationally specific privacy frameworks than those available for legal services. The RACGP guidance addresses clinical data privacy in the AI context directly, specifying requirements for patient consent, data de-identification, and restrictions on secondary use of clinical data for AI training purposes (RACGP, 2025). The MHRA does not appear to frame privacy as a standalone element of conformity assessment, but related UK guidance recognises data privacy and cybersecurity as important issues in the regulation of software and AI medical devices (Regulatory Horizons Council, 2022; MHRA, 2024).

Several gaps are relevant to legal services. First, no guidance addresses the privacy implications of training AI systems on legal case data (raising questions about client confidentiality, legal professional privilege, and the risk of re-identification from anonymised case records). Second, the concept of inferential privacy (the capacity of AI systems to derive sensitive personal information from seemingly innocuous data inputs), receives almost no attention. Despite being directly relevant to AI tools that process user narratives about legal problems. Third, accountability for privacy protection is unclear when unregulated AI tools serve consumers directly. Neither the tool provider nor the absent legal professional bears clearly defined privacy obligations equivalent to those that apply within the solicitor-client relationship. Fourth, guidance on cross-border data flows in the context of AI services (where processing may occur in multiple jurisdictions, and training data may have been collected globally) remains thin despite the inherently international nature of major AI platforms.

4.5.2 Data Protection

Data protection is the most thoroughly addressed theme across the entire standards landscape. Reflecting the maturity of the legal and regulatory framework governing personal data processing in the UK and EU.

The ICO provides one of the most granular requirements. It covers lawful bases for processing personal data through AI systems and the obligation to conduct data protection impact assessments for high-risk processing. The specific requirements of Article 22 concerning automated decision-making and the principle of data protection by design and default (ICO, 2023). The ICO guidance is notable for its operational specificity: it does not merely state principles but provides detailed expectations regarding how organisations should identify lawful bases, what DPIAs should contain, and how data subject rights should be facilitated in the AI context.

The EU AI Act creates a parallel and complementary data governance regime specifically for AI systems. Article 10 requires that training, validation, and testing datasets be subject to "appropriate data governance and management practices". To address training methodologies, data collection

processes, representativeness, bias examination, identification of data gaps, and measures to detect and correct errors (European Commission, 2024). This creates obligations that go beyond conventional data protection law: they are concerned not only with protecting individuals whose data is processed but with ensuring that the datasets themselves are of sufficient quality to support reliable and fair AI system performance.

Legal sector sources engage with data protection at a high level, directing practitioners to their existing obligations under the UK GDPR and Data Protection Act 2018 rather than providing AI-specific operational guidance (SRA, 2026; The Law Society, 2025). This approach has merit (data protection obligations are well-established and practitioners should indeed apply them) but it leaves several AI-specific questions unaddressed. How should firms conduct DPIAs for AI tools whose internal processing is opaque? What lawful basis supports the processing of client data through third-party AI platforms? How should the obligation to facilitate data subject rights (access, rectification, erasure) be met when personal data has been used to fine-tune or prompt an AI model?

The gap is most acute for consumer-facing tools. Within the regulated sphere, solicitors bear professional obligations regarding client data that broadly align with data protection requirements. Outside the regulated sphere, the consumer's data protection depends entirely on the tool provider's compliance with general data protection law. This compliance may be difficult for consumers to verify, enforce, or even understand. The intersection of data protection with the accountability and redress gaps identified in subsections 5.13 and 5.14 creates a compounding risk: consumers who cannot identify the person or organisation responsible for an AI tool's outputs are unlikely to be able to exercise data protection rights effectively against that same entity.

5. Taxonomy Analysis: Cross-Cutting Dimensions

The thematic analysis presented in the preceding section examined *what* standards address. This section examines *how* existing standards differentiate (or fail to differentiate) across six structural dimensions that shape the risk profile of AI tools in legal services. Using this taxonomy to identify where standards align with the realities of consumer-facing legal technology and where important mismatches remain.

5.1 System Type

Key finding: General-purpose AI tools (ChatGPT, Claude, Gemini) present one of the most acute consumer protection risks: no domain-specific expertise, uncertain training data, and outputs delivered with a fluency that masks their limitations. Yet existing standards do not distinguish between these and purpose-built legal tools with domain expertise.

The taxonomy distinguishes three AI-powered B2C lawtech types:

- General-purpose AI tools (such as ChatGPT, Claude, Gemini),
- Legal information platforms (such as legal databases and knowledge retrieval systems), and
- Issue-specific legal tools designed for particular practice areas (such as housing, debt, employment, immigration).

Each presents a distinct risk profile, yet the existing standards landscape tends to treat them the same.

General-purpose AI tools present one of the most acute consumer protection risks. These systems possess no domain-specific legal expertise, draw on training data of uncertain provenance and accuracy, and deliver outputs with a fluency and confidence that masks their epistemic limitations. A consumer querying ChatGPT about their housing rights receives a response indistinguishable in tone from professional legal advice, yet it's been generated through statistical pattern matching rather than legal reasoning. The system cannot distinguish between jurisdictions unless the user specifies where they are located. It may not identify when legislation has been superseded, or recognise that a consumer's specific circumstances engage exceptions or qualifications to a general legal principle. The EU AI Act classifies AI deployed in the "administration of justice and democratic processes" as high-risk, triggering mandatory conformity assessments, risk management systems, and human oversight requirements (European Commission, 2024). However, this classification was designed primarily for AI used within judicial and administrative decision-making, and its application to consumer-facing general AI tools providing incidental legal information remains ambiguous.

Legal information tools occupy an intermediate position. These platforms typically draw on verified, curated legal databases and present information within structured frameworks. Their risk profile differs from general AI tools in that their knowledge base is, in principle, bounded and auditable. However, the landscape review did not identify standards that specifically address this category of tools. Despite heightened expectations around their accuracy because users reasonably assume that a platform marketed as a legal information service maintains current, up to date, geographic jurisdiction-specific content. The integration of generative AI capabilities into these platforms (enabling natural language queries and synthesised responses rather than simple search results) introduces new risks that existing quality assurance frameworks have not fully addressed. Because they were not designed with generative AI in mind.

Issue-specific legal tools represent the category that most directly substitutes for professional legal services. A tool designed to assess eligibility for housing disrepair claims, guide consumers through

employment tribunal procedures, or help individuals navigate debt relief options performs activities that are commonly carried out by a regulated lawyer (despite not all of these being reserved legal activities).^{xxiii} These tools trigger the sharpest debates about authorisation requirements. Consumer advocates argue that tools providing personalised legal guidance in specific practice areas should be subject to pre-approval or certification before reaching consumers (Consumer/Advice Body interview). Lawtech providers counter that mandatory pre-approval processes would impose delays and costs that effectively prevent innovative access-to-justice tools from reaching the consumers who need them most (Lawtech Company interview). Professional guidance points to a middle path: rather than insisting on a rigid credential-based model for providers, it emphasises legal competence, provider transparency, and effective oversight by qualified lawyers when AI is used in legal practice (Law Society of Scotland, 2024; Law Society, 2025). This formulation acknowledges that legal domain competence matters while leaving open *how* that competence is demonstrated.

The practical significance of this dimension is substantial. Standards that apply uniformly across all three system types will likely be either too restrictive for lower-risk applications (legal information retrieval) or insufficiently protective for higher-risk ones (issue-specific guidance tools). The current landscape overwhelmingly adopts the former approach: general principles applied without calibration to system type.

5.2 User Function

Key finding: The distinction between legal information and legal advice is increasingly unstable in practice. The same AI model can produce both depending on the user's query. Consumers do not think in these regulatory terms, and regulatory frameworks premised on the distinction face a fundamental challenge.

This concerns what the AI tool does for the user. Two distinctions are relevant to how standards address user-facing functions. The first is between legal information (general statements about the law) and legal advice (the application of law to a specific individual's circumstances). While this is not a formal regulatory boundary in legal services (unlike in financial services), it is central to understanding consumer risk. Because users of tools providing personalised outputs face different risks from those receiving general information. The second, and legally operative, distinction is between reserved and unreserved activities under the Legal Services Act 2007. Reserved activities may only be carried out by authorised persons (such as solicitors or barristers); unreserved activities may be provided by anyone. Both distinctions are pertinent: reserved/unreserved determines regulatory powers and oversight, while information/advice determines the risk profile and the safeguards required for AI systems that may blur these boundaries.

Across the documents there is broad agreement that personalised outputs (those that take account of an individual's specific circumstances and provide tailored recommendations) warrant stronger safeguards than general information provision. Multiple sources specify that personalised AI outputs should be subject to mandatory human review before reaching consumers.^{xxiv} The Regulatory Response Unit has noted that regulations governing reserved legal activities are means-agnostic, applying regardless of whether the activities are performed with or without the use of technology (LawtechUK, 2024). This carries significant regulatory consequences: providing reserved legal activities without authorisation is a criminal offence under the Legal Services Act 2007.

However, this binary distinction is increasingly unstable in practice. The same large language model can produce both general legal information and highly personalised guidance depending solely on how the user frames their query. A consumer asking "what are my rights if my landlord hasn't returned my deposit?" will receive a response that blends general legal information with what appears to be personalised advice, without any clear demarcation between the two. As one interviewee observed,

“people don’t necessarily know the difference between legal information and legal advice” (Consumer/Advice Body interview).

A notable gap in the existing standards landscape concerns intermediate tools that sit between pure information delivery and personalised advice. Guided triage systems that assess which legal pathway a consumer should pursue, eligibility checkers that determine whether a consumer qualifies for a particular remedy, and document assembly tools that generate legal documents based on user inputs all occupy this intermediate space. These tools provide outputs that are personalised to the user’s circumstances and that the user is likely to rely upon, yet they may not involve the “application of law to facts” in the sense traditionally understood by legal regulators. The LSB’s own research into access to justice has suggested that these intermediate tools hold particular promise for underserved consumers who currently receive no assistance at all.^{xxv} One of the regulatory challenges is to develop standards that protect consumers using these tools without foreclosing their potential to narrow the access to justice gap.

5.3 Legal Scope

Key finding: Risks differ materially by legal domain: AI errors in family law carry different consequences from errors in commercial contract review. Yet the landscape overwhelmingly applies general principles without calibration to specific legal domains.

The third dimension examines whether standards address AI in legal services broadly or target narrow, function-specific applications. The distribution across the documents is heavily skewed toward breadth, with most coded findings addressing AI in general terms rather than focusing on specific legal applications or practice areas.

This breadth reflects a familiar regulatory approach, particularly in the UK, where frameworks tend to be principles- or outcomes-based rather than prescriptive. Core principles (transparency, accountability, data protection, human oversight) apply regardless of whether an AI tool assists with conveyancing, employment disputes, or family law. There are strong arguments for horizontal standards that establish baseline requirements across all legal AI applications.

However, the evidence from both the thematic analysis and interviews demonstrates that risks differ materially by legal domain. As one interview participant noted, “the risks of AI in family law are completely different from AI in commercial contract review” (Industry Association interview). Family law involves emotionally vulnerable individuals making decisions with profound personal consequences; errors carry risks that extend beyond financial loss to child welfare and personal safety. Commercial contract review involves sophisticated parties with access to professional advice; errors carry primarily financial consequences and are more likely to be identified through existing quality assurance processes.

The Dutch judiciary has taken an explicitly scope-limited approach, permitting AI only for “well-defined, auditable tasks” rather than open-ended legal reasoning (De Rechtspraak, 2024). This reflects a recognition that the tractability of AI oversight varies by application: structured, rule-based legal tasks are more amenable to quality assurance than open-textured legal reasoning requiring the weighing of competing considerations.

The structural tension is clear. Regulatory frameworks designed to apply across all areas of law provide useful baseline principles. However, they risk being too general to address the specific risks that arise in particular legal domains and contexts. Regulators may argue that the role of guidance is to show how principles apply in different contexts and domains. Domain-specific calibration is needed to translate general principles into operational requirements that reflect the actual risk profile of AI tools in different areas of law. The current landscape provides the general principles but largely lacks the domain-specific translation layer.

5.4 Oversight Model

Key finding: The most detailed oversight standards apply exclusively within regulated legal practice. A consumer using an unregulated AI chatbot has no equivalent protection – the chatbot may serve thousands of users daily with no human review of any output.

The fourth dimension concerns how human oversight is structured. Legal profession bodies internationally strongly agree on the requirement for meaningful human oversight. The SRA, Law Society, CCBE, and Queensland Law Society all mandate that AI outputs be reviewed by a competent legal professional before reaching clients. This convergence reflects the profession's established model of individual practitioner responsibility: a solicitor who deploys AI remains personally accountable for the accuracy and appropriateness of the output. Just as they would be accountable for the work of a junior colleague or paralegal.

Autonomous AI operation (systems that generate and deliver legal outputs without human review) is treated in the current standards landscape as only for “information” delivery, and even then, with qualifications. The EU AI Act requires that high-risk AI systems be designed to allow “effective oversight by natural persons” during their period of use (European Commission, 2024). ISO/IEC standards address the controllability of automated AI systems, including mechanisms for human intervention and override (ISO/IEC, 2024). The FCA's Consumer Duty represents a distinct approach: rather than prescribing specific oversight mechanisms, it requires firms to demonstrate that their products deliver good outcomes for consumers. Leaving the means, including the nature and extent of human oversight, to the firm's judgement (FCA, 2022).

The critical unresolved challenge is that the most detailed oversight standards apply exclusively within regulated legal practice. A solicitor using AI to draft client advice is bound by SRA standards requiring competent supervision. A consumer using an unregulated AI chatbot to understand their legal rights has no equivalent oversight or protection requirement. The chatbot may serve thousands of users simultaneously, generating thousands of individual legal outputs, with no human review of any of them. Even within regulated practice, the practical challenge of maintaining oversight quality at scale is significant. A firm deploying AI to generate first drafts of client correspondence may technically review each output, but the cognitive dynamics of reviewing AI-generated text (the tendency to defer to fluent, apparently competent output) raise questions about whether such review constitutes meaningful oversight or ‘tick box’ exercise.

5.5 Standards Approach

Key finding: The dominant orientation is principles-based. The legal services sector has principles in abundance; what it lacks is the infrastructure to operationalise them into verifiable, enforceable standards that provide consumers with meaningful protection.

The fifth dimension classifies instruments by their regulatory mode: rules-based, principles-based, or outcomes-based. The dominant orientation across the 234 sources is principles-based. The UK Government has organised its AI governance framework around five cross-cutting principles (safety, transparency, fairness, accountability, and contestability) to be applied by existing regulators within their respective domains (DSIT/OAI, 2023). This approach reflects a deliberate choice to avoid prescriptive technology-specific rules and reflects the practical recognition that principles can accommodate rapid technological change and tend to date less quickly than rules).

The EU AI Act represents the clearest rules-based instrument in the body of documents, establishing mandatory conformity assessments, technical documentation requirements, quality management systems, and post-market monitoring obligations for high-risk AI systems (European Commission,

2024). Its rules-based approach provides greater certainty for providers about what compliance requires. However, it risks obsolescence as technology evolves and imposes compliance costs that may be disproportionate for smaller providers.

The FCA's Consumer Duty exemplifies an outcomes-based approach, requiring firms to act to deliver good outcomes for retail customers across four outcome areas: products and services, price and value, consumer understanding, and consumer support (FCA, 2022). This approach focuses regulatory attention on what matters (whether consumers are, in fact, protected) rather than on process compliance. However, it requires robust measurement of outcomes, which is challenging when the harms from inadequate legal AI may be diffuse, delayed, or difficult to attribute.

The key gap across all three approaches is the translation from articulated principles to verifiable, enforceable standards that provide consumers with meaningful protection. The CDEI roadmap called for the development of an AI assurance ecosystem (including technical standards, audit methodologies, and certification schemes) to bridge this translation gap (CDEI, 2021). While progress has been made in some sectors, operational standards for legal AI remain substantially underdeveloped. The legal services sector has principles in abundance; what it lacks is the infrastructure to operationalise them.

5.6 Reserved and Non-Reserved Activity

Key finding: The result is a two-tier system of consumer protection: robust standards for AI used by professionals within regulated practice, and minimal protection for consumers using AI Lawtech tools, where consumer-facing AI tools proliferate most rapidly.

The sixth dimension examines whether standards differentiate between AI tools used for reserved legal activities (rights of audience, conduct of litigation, reserved instrument activities, probate activities, notarial activities, administration of oaths) and those used for non-reserved legal services in England and Wales. This dimension is specific to the Legal Services Act 2007 framework and has no direct equivalent in other jurisdictions or sectors.

Most findings in this dimension fail to specify whether they apply to reserved or non-reserved legal activities. The majority of standards were drafted without proper consideration of the Legal Services Act 2007. Understandably so, given that most originate from international, cross-sector, or technology-specific contexts. Even standards produced by legal services bodies tend to frame requirements in terms of "legal services" generically rather than distinguishing between specific legal domains and contexts,

Within regulated practice, there is full convergence that existing professional obligations apply to AI use. A solicitor conducting a reserved legal activity using AI remains subject to the same professional standards, competence requirements, and accountability frameworks as if they had performed the work manually.

The critical structural gap lies outside the regulatory perimeter. Most enforceable consumer protections exist within that perimeter, applying to reserved activities performed by authorised persons. Consumer-facing AI tools proliferate most rapidly in non-reserved areas: legal information provision, document assembly, guided triage, eligibility checking, and general legal guidance. These activities fall outside the reserved activities framework and may be provided by anyone, including unregulated technology companies with no connection to the legal profession. The result, as multiple sources and interviewees identified, is a two-tier system of consumer protection: robust standards for AI used within regulated practice, and minimal sector-specific protection for consumers using AI tools outside it.

6. Cross-Theme Synthesis

The preceding sections examined the thematic landscape and the taxonomic structure of existing standards. This section draws the analysis together, identifying where standards converge into clusters, where they create unresolved tensions, how they distribute along a maturity spectrum, and what lessons emerge from cross-sector comparison. Figure 11 presents a heatmap showing coverage (number of documents) containing content related to each standard and taxonomy element.

Figure 11: Heatmap showing coverage (number of documents) by theme and taxonomy

| | System Type | User-Facing Function | Legal Scope | Oversight Model | Standards Approach | Reserved / Non-Reserved | |
|------------------------------|-------------|----------------------|-------------|-----------------|--------------------|-------------------------|------------------------|
| Data Protection | 61 | 52 | 27 | 73 | 73 | 15 | Data & Privacy |
| Privacy | 55 | 49 | 24 | 69 | 69 | 11 | |
| Transparency | 101 | 92 | 47 | 126 | 128 | 21 | Governance & Assurance |
| Accountability | 94 | 88 | 46 | 115 | 117 | 19 | |
| Human Oversight | 84 | 78 | 44 | 101 | 101 | 20 | |
| Quality Assurance | 72 | 63 | 33 | 95 | 96 | 14 | |
| Traceability | 71 | 62 | 31 | 90 | 89 | 12 | |
| Bias | 71 | 67 | 33 | 90 | 89 | 15 | |
| Fairness | 71 | 65 | 32 | 88 | 89 | 14 | Fairness & Inclusion |
| Non-discrimination | 60 | 56 | 29 | 77 | 76 | 11 | |
| Accessibility | 43 | 43 | 24 | 46 | 46 | 12 | |
| Accuracy | 68 | 63 | 33 | 84 | 84 | 13 | Output Quality |
| Reliability | 66 | 60 | 32 | 86 | 86 | 15 | |
| Consistency | 33 | 32 | 17 | 45 | 46 | 9 | |
| Signposting & Referral | 54 | 51 | 28 | 59 | 60 | 12 | User Protection |
| Vulnerability Identification | 66 | 62 | 29 | 78 | 79 | 12 | |
| Redress & Complaints | 41 | 39 | 24 | 48 | 48 | 12 | |

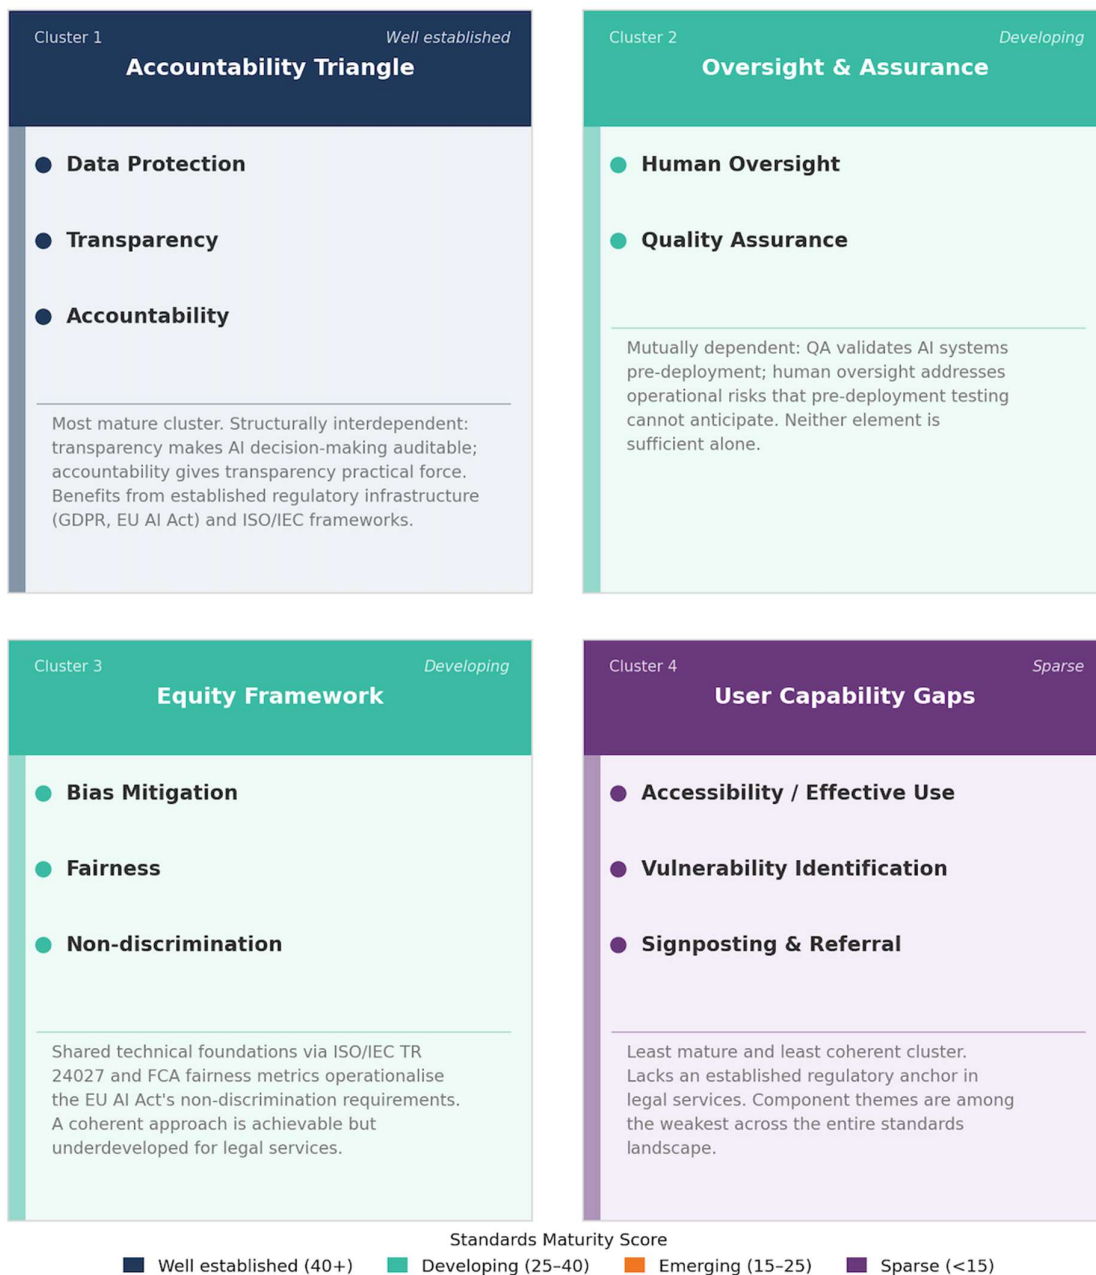
Taxonomy dimensions: System Type classifies AI tools as general-purpose (e.g. ChatGPT, Claude), legal information platforms (e.g. legal databases), or issue-specific legal tools (e.g. housing, debt, employment). User-Facing Function distinguishes between legal information provision, personalised legal advice, and intermediate functions such as triage and eligibility checking. Legal Scope indicates whether standards address AI in legal services broadly or target narrow, function-specific applications. Oversight Model captures how human oversight is structured: human-in-the-loop review, autonomous operation, or hybrid arrangements. Standards Approach classifies instruments by regulatory modality: rules-based, principles-based, or outcomes-based. Reserved / Non-Reserved indicates whether standards differentiate between AI used for reserved legal activities (e.g. rights of audience, conduct of litigation, etc.) and non-reserved legal services.

6.1 Convergences

Key finding: Standards cluster into four interdependent groupings: (1) data protection, transparency, and accountability; (2) human oversight and quality assurance; (3) fairness, bias, and non-discrimination; and (4) accessibility, vulnerability, and signposting. Of these, the first is relatively mature and well-supported, while the latter, most directly linked to consumer protection, remains fragmented and underdeveloped.

The analysis reveals four clusters of themes that reinforce one another and function as interdependent elements rather than standalone requirements, as shown in Figure 12. These clusters were identified through qualitative analysis of coding patterns, grouping themes that tended to appear together in the documents.

Figure 12: Cross-Theme Convergence Clusters



The first cluster links data protection, transparency, and accountability into a mutually reinforcing triangle. The ICO requires data protection impact assessments for high-risk AI processing; the EU AI Act mandates technical documentation and transparency obligations; ISO/IEC 23894 integrates risk management across the AI system lifecycle. These are not parallel requirements that happen to coincide. They are structurally interdependent: transparency enables accountability by making AI decision-making processes auditable, and accountability gives transparency its practical force by creating consequences for failures. The NTIA observed that transparency enables accountability, and accountability typically requires transparency (NTIA, 2024). This cluster is amongst the most mature in the standards landscape, benefiting from established regulatory infrastructure (GDPR, sector regulators) and well-developed technical standards (ISO/IEC frameworks).

The second cluster pairs human oversight with quality assurance. The EU AI Act requires "effective oversight by natural persons" for high-risk AI systems. ISO/IEC quality evaluation standards specify testing and validation protocols before deployment. The RACGP framework demonstrates this pairing in clinical practice: AI systems are validated through quality assurance processes before deployment, and individual clinicians review AI outputs before acting on them. Neither element is sufficient alone. Quality assurance without human oversight assumes that pre-deployment testing can anticipate all operational risks, an assumption contradicted by the emergent behaviour of large language models. Human oversight without quality assurance places the entire burden of safety on individual reviewers, whose capacity to identify AI errors is limited by their own knowledge, cognitive biases, and the volume of outputs requiring review.

The third cluster connects fairness, bias mitigation, and non-discrimination through shared technical foundations. ISO/IEC TR 24027 provides a taxonomy of bias sources that informs the FCA's fairness metrics, which in turn operationalise the EU AI Act's non-discrimination requirements. The ICO's guidance on fairness metrics selection provides the measurement methodology. This cluster demonstrates how international technical standards, sector-specific regulatory requirements, and data protection frameworks can be integrated into a coherent approach, though its operationalisation in legal services remains underdeveloped.

The fourth cluster addresses user capability gaps through accessibility, vulnerability identification, and signposting. The ICO and Alan Turing Institute recommend layered transparency approaches calibrated to user capability. The WHO recommends mechanisms to identify patients with reduced decision-making capacity. The RACGP requires that clinicians are accountable for care decisions supported by AI tools, and documents processes to support their clinical oversight. This cluster is the least mature and least coherent. Unlike the other three clusters, it lacks an established regulatory anchor in legal services, and its component themes are among the weakest in the landscape.

6.2 Tensions and Trade-Offs

Key finding: The standards landscape is shaped by four inherent and unresolved tensions – between transparency and technical complexity, data minimisation and bias detection, human oversight and scalability, and accuracy and legal uncertainty – which reflect fundamental trade-offs in AI governance.

Four structural tensions emerge that cannot be resolved through better drafting alone, as they reflect genuine conflicts between legitimate regulatory objectives.

First, the tension between transparency and the inherent complexity of generative AI explanations. Standards consistently require that AI decision-making processes be explicable to users. Large language models produce outputs through billions of probabilistic transformations across neural network layers, and no faithful explanation of this process is accessible to a lay consumer. Simplified explanations risk being misleading; technically accurate explanations are incomprehensible. The legal

services context sharpens this tension because consumers need to calibrate their reliance on AI outputs, a task that requires some understanding of the system's reliability and limitations, which in turn requires some form of explanation. Where meaningful explanation cannot realistically be provided, alternative mechanisms for signalling trustworthiness, such as independent certification or accreditation schemes, may play an important complementary role.

Second, data minimisation conflicts with bias detection. The GDPR requires that personal data processing be limited to what is necessary for the specified purpose. Detecting and mitigating bias in AI outputs, however, frequently requires collecting and analysing data about protected characteristics (ethnicity, gender, disability, age) that would not otherwise be necessary for the service being provided. This tension is not unique to legal services, but it is particularly acute in contexts where bias could affect access to justice.

Third, the requirement for meaningful human oversight conflicts with the scalability of consumer-facing AI services. Standards developed within the framework of professional legal practice assume a model in which an individual practitioner reviews AI outputs before they reach a specific client. Consumer-facing AI tools operate at a fundamentally different scale: a legal chatbot may serve thousands of users daily, generating individualised outputs for each. Meaningful human review of every output is economically infeasible at this scale, yet the standards offer no framework for determining what level of sampling, spot-checking, or automated quality monitoring might constitute adequate oversight.

Fourth, accuracy requirements sit uneasily with the irreducible uncertainty of legal reasoning. Technical standards for AI accuracy presuppose the existence of ground truth against which outputs can be evaluated. Legal reasoning frequently involves the weighing of competing considerations where reasonable lawyers would disagree, where the law itself is uncertain, and where the "correct" answer depends on judicial interpretation that has not yet occurred. AI accuracy metrics developed for classification tasks or factual retrieval do not translate straightforwardly to contexts of legal uncertainty.

6.3 Standards Development Spectrum

Key finding: The level of development across themes is highly uneven: data protection, transparency, and fairness are the most developed, supported by established regulatory and technical frameworks, while areas most critical to consumer protection (particularly vulnerability identification, signposting, redress, accessibility, and system reliability) remain the least developed and lack operational standards.

Mapping the 17 areas along a degree of maturity spectrum (from most to least developed) reveals a clear pattern in the level of development of standards across themes. Figure 13 illustrates the level of development of individual themes.^{xxvi} Figure 14 aggregates related themes into broader policy lenses, illustrating how standards coverage clusters across key governance functions.

Figure 13: Standards Maturity Spectrum by Theme (weighted score: substantive coverage = 2, partial = 1)





Figure 14: Standards Coverage by Policy Lens (average document coverage per theme within each lens)



The most developed themes are data protection, transparency, and bias and fairness. These benefit from established regulatory infrastructure (the GDPR and its enforcement by the ICO), well-developed international technical standards (the ISO/IEC 42001 series), and cross-sector regulatory attention (the DRCF's collaborative work programme). Standards in these areas have moved beyond principles to include specific requirements, metrics, and enforcement mechanisms.

Moderately developed themes include accountability, quality assurance, human oversight, and privacy. These have clear principles articulated across multiple sources and some operational guidance but lack the specificity of the most mature themes. Accountability, for example, benefits from established professional responsibility frameworks within regulated legal practice but has no equivalent framework for unregulated providers. Quality assurance has well-developed technical standards in healthcare and manufacturing that could be adapted but have not yet been translated for legal services.

The least developed themes are vulnerability identification, signposting and referral, redress and complaints, accessibility, consistency, and reliability. These themes have the thinnest coverage in the documents, the fewest operational requirements, and the weakest enforcement mechanisms. Vulnerability identification has no workable framework despite healthcare models that could serve as templates. Signposting and referral mechanisms are largely absent from standards despite being fundamental to ensuring that consumers who need human professional assistance are directed toward it.

6.4 Cross-Sector Lessons

Key finding: Both healthcare and financial services are substantially ahead of legal services in operationalising AI governance. Healthcare has pre-market validation; financial services has named accountability and outcomes-based conduct rules. Legal services has principles but lacks the infrastructure to apply them.

Comparison with healthcare and financial services reveals that both sectors are substantially ahead of legal services in operationalising AI governance principles.

Healthcare has developed the most comprehensive pre-market validation infrastructure. The MHRA's regulatory framework for AI as a medical device requires clinical evidence demonstrating safety and efficacy before market access. As well as ongoing post-market surveillance to identify emerging risks, and predetermined change control plans that specify which modifications to an AI system trigger renewed regulatory review (MHRA, 2023). The FDA's approach in the United States parallels this framework. The underlying principle (that AI tools affecting health outcomes must demonstrate safety before reaching patients, not after) has no equivalent in legal services, where AI tools reaching consumers face no pre-market validation requirement.

The financial services sector has developed robust accountability and conduct-of-business infrastructure. The FCA's Senior Managers and Certification Regime creates personal accountability for named individuals within regulated firms. The Consumer Duty establishes outcomes-based conduct standards and emphasises bias monitoring for AI used in credit decisions (FCA, 2022). The Financial Ombudsman Service provides a well-established complaints pathway with binding adjudication powers. Legal services has analogous professional accountability within regulated practice (the SRA's Principles and Code of Conduct, the Legal Ombudsman) but no equivalent for consumer-facing AI tools provided by unregulated entities.

The lesson is not that legal services should replicate healthcare or financial services frameworks wholesale. The nature of legal services (characterised by legal uncertainty, adversarial processes, and the irreducibility of professional judgement in many contexts) presents distinctive challenges. The lesson is rather that legal services possesses the principles but lacks the regulatory infrastructure to

operationalise them: the pre-market validation protocols, the conduct-of-business rules for consumer-facing tools, the sector-specific complaints pathways, and the named accountability for AI governance.

7. Summary Gap Analysis

The thematic analysis, taxonomy analysis, and cross-theme synthesis collectively reveal a standards landscape that is broad in its coverage of principles but uneven in its operational depth. This section provides an overview of the gaps, distinguishing between thematic gaps (areas where substantive standards are weak or absent), structural gaps (architectural features of the regulatory landscape that create systematic blind spots), and taxonomy-specific gaps (failures of differentiation within existing standards).

7.1 Thematic Gaps

Key finding: The most serious and immediate thematic gaps are in vulnerability identification, signposting, and redress, with further deficiencies in accessibility, consistency, reliability, and accuracy.

Figure 15 provides a visual overview of thematic coverage across the standards landscape, highlighting areas where standards attention is strongest and where significant gaps remain. Themes assessed as better covered (data protection, transparency, bias, accountability, human oversight, and privacy) benefit from existing regulatory frameworks that provide enforceable requirements. Even these, however, lack legal-services-specific detail and do not address the distinctive characteristics of consumer-facing legal AI.

Figure 15: Standards Coverage Gap Analysis by Theme^{xxvii}




Table 1 summarises priority thematic gaps identified in the standards landscape by cross-referencing coverage rates with policy relevance for consumer-facing legal services. Level of priority is assessed according to the severity (potential for consumer harm), urgency (consumers are already exposed to the risk), and tractability (feasible solutions exist or can be adapted from other sectors).

Table 1: Priority Thematic Gaps in the Standards Landscape

| Level of Priority | Gaps Identified |
|--------------------|--|
| Urgent | <ol style="list-style-type: none"> Vulnerability identification is one of the most significant gap in the standards landscape. No workable framework exists for identifying consumers who may be particularly susceptible to harm from AI-generated legal information or guidance. Vulnerability in the legal services context includes cognitive impairment, emotional distress (particularly prevalent among individuals dealing with family, housing, welfare, or debt issues), low literacy, limited digital skills, and language barriers. AI tools currently treat all users identically regardless of capability or circumstance. Healthcare has developed transferable models: the WHO framework recommends mechanisms to identify patients with reduced decision-making capacity, and the RACGP mandates escalation protocols when AI systems encounter users who may require additional support. Adapting these models for legal services is feasible and should be a priority. Signposting and referral have similar gaps in coverage. No standard methodology exists for AI B2C legal tools to communicate the limits of their competence, assess when a user's situation exceeds those limits, or direct users toward appropriate professional assistance. Consumers querying a general AI tool about a complex legal matter receive a response regardless of whether the matter is within the system's competence. There is no obligation to flag uncertainty, no trigger for escalation, and no mechanism for referral. This gap is particularly acute given evidence that consumers are using general AI tools for legal queries precisely because they do not know where else to turn (Consumer/Advice Body interview). Redress and complaints lack any AI-specific pathway. A consumer who suffers harm as a result of relying on AI-generated legal information has no clear route to redress. Within regulated practice, the Legal Ombudsman provides a complaints mechanism, but its jurisdiction is limited to complaints about authorised persons. For unregulated AI tools, general consumer protection law provides theoretical remedies, but the practical barriers (identifying the responsible entity, demonstrating causation, bearing the costs of pursuing a complaint) render these remedies largely inaccessible to the consumers most likely to need them. |
| Significant | <ol style="list-style-type: none"> Accessibility and effective use is largely unaddressed on the input side: standards focus on making AI outputs accessible but rarely address whether consumers can effectively formulate queries, understand the limitations of the system, or evaluate the reliability of responses. Consumers with low literacy or limited English proficiency face particular barriers. Consistency lacks any metrics or benchmarks; there are no standards addressing the well-documented tendency of large language models to generate different responses to the same query depending on phrasing, session context, or model version. Reliability has no testing protocols or acceptable error thresholds for legal AI. Accuracy has no hallucination prevention standards and no pre-deployment testing requirements specific to legal content. |
| Moderate | <ol style="list-style-type: none"> Fairness and non-discrimination principles exist but lack legal-services-specific operationalisation. Quality assurance has no sector-specific requirements despite well-developed frameworks in healthcare and manufacturing. Traceability has no standardised audit trail requirements for legal AI outputs. |

7.2 Structural Gaps

Key finding: Structural gaps in the regulatory landscape leave unregulated providers, cross-border use, consumer protection, and professional oversight largely unaddressed, with additional challenges around the information–advice boundary,

proportionality for smaller providers, and the unclear scope of legal regulation in AI contexts.

Seven structural gaps emerge from the analysis. These are not gaps in individual standards but architectural features of the regulatory landscape that create systematic blind spots, as follow:

1. **Unregulated providers** of consumer-facing AI tools operate with no sector-specific standards. Professional regulation, by definition, can only reach authorised persons and regulated firms. General-purpose AI tools (ChatGPT, Claude, Gemini) that consumers use for legal queries are developed by technology companies with no connection to the legal profession and have no obligation to comply with legal services regulation (or understanding of its regulation). Platform usage policies, while sometimes addressing legal use cases, are unilaterally set, inconsistently enforced, and provide no consumer remedies.
2. **Cross-border challenges** are largely unaddressed. AI tools operate across jurisdictions, and a consumer in England and Wales may use a tool developed in the United States, trained on data from multiple legal systems, and hosted in a third country. This raises both regulatory and substantive risks: outputs may draw on legal rules or precedents from other jurisdictions unless systems are specifically designed and validated for England and Wales. No harmonised standards exist for determining which jurisdiction's regulatory requirements apply, how conflicts between regulatory frameworks should be resolved, or how enforcement should be coordinated.
3. The **compliance burden** of formal standards may be disproportionate for small and emerging providers, precisely the firms or developers most likely to be developing innovative access-to-justice tools. If compliance requires investment in conformity assessments, technical documentation, and quality management systems designed for large organisations, smaller providers may be excluded from the market, reducing rather than enhancing consumer access to legal services.
4. **Consumer protection standards** are almost entirely absent from the standards landscape. No standard requires AI legal tools to ensure that users understand what they are receiving (information, not advice; AI-generated, not professionally reviewed; general, not jurisdiction-specific) and what they are not receiving (professional legal advice, a confidential solicitor-client relationship, access to professional indemnity insurance).
5. The **lack of professional oversight** is perhaps the most consequential structural feature. Professional regulation assumes a solicitor intermediary who exercises judgement, assumes responsibility, and provides a point of accountability. Consumer-facing AI tools eliminate this intermediary. Standards designed for a model of professional intermediation do not function when the consumer interacts directly with the AI system with no professional gatekeeper.
6. The **boundary between reserved and unreserved legal activities** remains operationally undefined in AI contexts and many Lawtech developers we spoke to had no understanding of reserved activities and the regulation around them.
7. **While the regulatory boundary lies between reserved and unreserved legal activities, the distinction between legal information and personalised advice is central to consumer risk.** Existing standards rarely engage clearly with either. They do not differentiate between tools providing general information and those offering personalised guidance, despite the significantly different risks. This distinction is increasingly blurred by generative AI, which combines both without clear demarcation, while many tools frame their outputs as “information only” to avoid regulatory exposure even where they resemble advice.

7.3 Taxonomy-Specific Gaps

Key finding: Existing standards apply a one-size-fits-all approach, failing to differentiate by system type, user function, or regulatory scope, and therefore misalign requirements with actual levels of risk.

The taxonomy analysis reveals that existing standards fail to differentiate where differentiation is most needed. Standards overwhelmingly treat all AI tools identically regardless of system type: no distinction between the oversight requirements appropriate for a general-purpose chatbot and those appropriate for a narrow, domain-specific eligibility checker. Standards rarely distinguish between the requirements appropriate for information delivery and those appropriate for personalised guidance. Most critically, the majority of standards were not drafted with attention to the reserved/non-reserved distinction that structures the Legal Services Act 2007 framework. Resulting in standards that are either inapplicable to the highest-risk tools (those outside the regulatory perimeter) or unnecessarily burdensome for lower-risk applications within It. By applying uniform governance expectations that do not reflect differences in risk, user reliance, or the presence of professional oversight.”

8. Interview Findings

Interviews provided a distinct evidence stream within this landscape review. Nine interviews were conducted with representatives of lawtech companies, consumer advocacy organisations, industry associations, and regulatory bodies. Interviewees included founder-practitioners developing and operating AI tools directly, as well as policy and regulatory specialists with sector-wide perspectives. This section summarises key insights from interview data. Where interview findings reinforce or complicate the documentary analysis, those connections are made explicit.

8.1 Awareness of Existing Standards

Key finding: Most lawtech providers, particularly those without a legal background, could not identify specific standards applicable to AI in legal services. The gap between the volume of standards identified in this review and providers' awareness of them points to a fundamental dissemination failure. In the absence of clear guidance, providers have improvised their own compliance approaches, drawing on adjacent regulatory frameworks. Regulators were also perceived as behind the curve in their understanding of the technology.

One of the most striking findings from the interviewees is the low level of awareness of existing standards particularly any legal standards and especially from lawtech developers who do not have a legal background. Most interviewees were unable to identify specific standards applicable to AI in legal services. A small number cited forthcoming EU legislation (principally the EU AI Act) as a relevant framework. None referenced the ISO/IEC standards on AI governance, risk management, or bias mitigation that constitute a significant portion of the documents identified through this review.

This finding carries significant implications. Standards can influence behaviour only if those whose behaviour they seek to influence are aware of their existence, understand their requirements, and perceive consequences for non-compliance. The gap between the volume of standards identified in this review (160 documents) and many interviewees' awareness of them suggests a fundamental disconnect. Standards that exist on paper but not in practice provide no consumer protection. It also reflects two further structural features of the landscape. First, many of the documents identified were not designed with AI specifically in mind; they are general technology governance standards, professional conduct frameworks, or data protection instruments applied by extension to AI contexts. Second, many were not conceived by their creators primarily as 'standards' but as guidance documents, policy statements, or voluntary frameworks that do not carry the authority or recognisability of formal standards. Developers seeking to understand what standards apply to them face a fragmented, unlabelled landscape rather than a coherent framework, which partly explains why awareness is low even among those who have made genuine efforts to comply.

The headline implication is this: in the absence of clear, accessible, and AI-specific standards, lawtech providers have arrived at their own working frameworks through a combination of adjacent regulatory requirements, professional instinct, and trial and error. Several interviewees had drawn on the SRA professional conduct rules and the Advertising Standards Authority requirements as the nearest available analogues. One developer reported applying:

"The SRA rules, code of conduct, advertising standard agency recommendations, but there's nothing enshrined that I've seen currently which would apply." (Lawtech Company interview)

This improvised compliance is characteristic of the sector and depends on individual awareness and good faith rather than systematic application. Section 8.6 below explores in more detail how providers have developed de facto standards through product design choices in the absence of formal requirements.

A related finding is that regulators themselves were perceived as behind the curve. One interviewee observed that regulators' working knowledge of AI tended to extend little beyond familiarity with ChatGPT. Without a deeper understanding of how large language models can be configured, constrained, and applied in specific contexts (Lawtech Company interview). This points to a mutual awareness gap: low awareness of standards among developers corresponds to low awareness of the technology among regulators, creating conditions in which productive dialogue about appropriate standards is difficult to initiate.

8.2 What Interviewees Prioritise

Key finding: Interviewees consistently identified three priorities for standards development: data protection and privacy, transparency about AI use, and safeguarding vulnerable users. These reflect where providers have invested most heavily in design choices and where regulatory support is weakest. Providers face the greatest direct exposure in precisely these areas.

When asked which areas most urgently require standards development, interviewees consistently identified three priorities: data protection and privacy, transparency about AI use, and safeguarding vulnerable users. These priorities align with the thematic analysis but with a notable difference in emphasis. The formal standards landscape prioritises data protection and transparency (themes with established regulatory infrastructure). Legal professionals and consumer representatives placed comparatively greater weight on vulnerability and safeguarding, themes that the formal landscape addresses least adequately.

Why these three? These areas reflect where providers had invested most heavily in design choices, suggesting they perceived them as the areas of greatest live risk. None identified them as gaps in what they were doing; rather, they emerged as areas where providers had had to develop their own approaches in the absence of clear guidance, and where the regulatory environment offered least support. The three priorities also reflect the practical exposure of providers operating without regulatory authorisation: data breaches, misleading communications, and harm to vulnerable users are the most direct routes to reputational damage and legal exposure for unregulated providers.

- **Data protection and privacy:** Privacy emerged not merely as a compliance obligation but as a design principle across several products. All described deliberate effort to minimise the data footprint of consumer interactions, driven by the sensitivity of legal queries. Specific design choices (including technical architecture, data handling arrangements, and data minimisation) are discussed in detail in section 8.6 below.

"People may seek legal advice on things that are incredibly private to them... it is incredibly important to make sure that the systems are secure and that privacy of the people who input that information is safe." (Lawtech Company interview)

This sensitivity is compounded by nascent evidence of users sharing personal identifiers with AI tools without understanding the consequences, a pattern already observed with general-purpose AI tools and likely to extend to AI legal tools.

- **Transparency:** Interviewees emphasised transparency as both a consumer-facing and a structural requirement. Consumer-facing transparency was interpreted consistently as: being explicit that users are interacting with an AI; being clear about what the tool can and cannot do; and providing appropriate caveats. Several interviewees went further, proposing that disclosure should extend beyond a generic disclaimer to identify specific areas of uncertainty in outputs, empowering users to seek verification where outputs are less reliable.

"We make it very clear from the offset that they're speaking to an AI and what the tool can and cannot do." (Lawtech Company interview)

- **Safeguarding vulnerable users:** Multiple interviewees placed safeguarding at the centre of their product design, encompassing interface design for neurodivergent users, data minimisation for users in distress, and deliberate counter-measures against confirmation bias. Specific approaches are discussed in detail in section 8.6 below.

Lawtech providers described their approach to managing risk in the absence of specific standards. One representative explained that their tool is "deliberately narrow in scope" (confined to a specific legal domain and a specific use case) as a quality control strategy (Lawtech Company interview). This approach reflects an implicit risk-proportionality calculus: by limiting the scope of the tool, providers can maintain the domain expertise necessary for quality assurance without the resources required for a broad-scope system. Another provider described designing their tool as an adviser-support tool rather than a consumer-facing system:

"We're not giving advice to the end user, we're giving information to the adviser." (Lawtech Company interview).

This architectural choice transfers the oversight responsibility to the human adviser, effectively reintroducing the professional gatekeeper that consumer-facing tools eliminate.

Consumer representatives emphasised that vulnerability is not an abstract concern:

"People are going to ChatGPT because they don't know where else to go." (Consumer/Advice Body interview)

This observation reframes the regulatory challenge. Consumers using AI for legal guidance include individuals who cannot access or afford professional legal services. They may be the same individuals who are more vulnerable to harm from inaccurate outputs and at risk of harm with no legal advice. However, there is also a segment of users who use AI tools for reasons of cost, convenience, or accessibility (for example, as a 24/7 resource or as an initial step before seeking professional help). Both groups require protection, though the nature and degree of risk differs.

One interviewee challenged what they characterised as a paternalistic tendency in parts of the sector; that is, the assumption that B2C users:

"Need an expert and the expert services are going to cost more than they spend on food in a month." (Lawtech Company interview)

If the realistic alternative to an imperfect AI tool is no legal support at all, the risk-benefit calculus is not straightforward.

- **Accuracy and quality assurance:** Interviewees consistently identified accuracy of AI outputs as a priority and described a range of practical quality assurance approaches (discussed in detail in section 8.6 below). One interviewee put the question of what accuracy should be measured against directly:

"Accuracy of advice is hard to measure; how are human professionals regulated to that end? It's not a uniform, easy-to-parse definition...it needs to be the baseline." (Lawtech Company interview)

This framing aligns with the documentary finding that the standards landscape lacks agreed definitions and measures for accuracy in AI legal outputs.

8.3 The Case For and Against Standards

Key finding: There is broad support for standards, but significant concerns about speed of development, cross-border inconsistency, and enforceability. The pace of AI capability improvement means prescriptive technical standards risk obsolescence before publication; outcomes-based frameworks are more durable. The appropriate quality comparator is the realistic alternative available to consumers, often no legal assistance at all, not a theoretical standard of perfection. Unenforced standards protect no one.

Interviewees expressed broad support for the development of standards but with significant caveats about feasibility and unintended consequences.

The principal concern was the speed of technological change. Standards development processes (whether through ISO consensus procedures, regulatory consultation, or legislative processes) operate on timescales of years. AI capabilities are advancing on timescales of months. Standards that prescribe specific technical requirements risk being obsolete before they are published. This concern favours principles-based and outcomes-based approaches over prescriptive rules, though it does not resolve the challenge of translating principles into operational requirements. It is worth noting that advocates of voluntary and principles-based standards argue that these instruments can be developed and revised far faster than hard regulation requiring primary legislation, where timescales can run to a decade or more. The speed-of-change concern is most acute for prescriptive technical standards; outcomes-based frameworks are more durable. One interviewee illustrated the pace of change by noting that hallucination (widely cited in policy discourse as a primary AI risk) had in practice receded as a concern as model capability improved, and that regulators were working from an outdated risk picture (Lawtech Company interview).

Standards were seen positively by some - meeting industry standards or having accreditation help build consumer confidence and trust and they also provide product differentiation in a busy market. The balance is between standards being an enabler that give consumers trust and a burden on developers or choking innovation in a developing market.

Several interviewees raised the cross-border dimension. AI is inherently borderless: the same model serves users across jurisdictions, and the same provider can operate from any location with internet connectivity. Overly restrictive UK standards risk driving lawtech companies out of England and Wales rather than improving access to justice and consumer protection. One interviewee described this concretely, noting that jurisdictions with more agile regulatory processes (such as the UAE, where ministerial direction can translate into law almost immediately) were attracting AI investment. Precisely because regulatory frameworks could adapt at the pace of technological development (Lawtech Company interview).

A more nuanced point concerned the appropriate comparator for AI quality. Several interviewees observed that the quality of existing human legal advice is variable and legal professionals make mistakes. AI Lawtech, however, can reach far more consumers than an individual legal professional or firm and consequently, its potential harm (or benefit) is greater. One interviewee argued that:

"AI standards should be compared against the realistic alternative, not a perfect ideal."
(LawTech Company interview).

For consumers whose realistic alternative is no legal assistance at all, an AI tool that provides reasonably accurate general guidance (even with known limitations) may represent an improvement in access to justice. This argument does not justify the absence of standards, but it does suggest that standards should be calibrated to realistic alternatives rather than theoretical perfection.

An industry association representative observed that:

"Current regulatory frameworks do not adequately distinguish between AI tools used by lawyers and those used directly by consumers." (Industry Association interview)

This observation identifies one of the key regulatory challenges: the regulatory framework was designed for a model of professional intermediation and has not been adapted for a world in which consumers interact directly with AI systems. The interviewee added that this had real commercial consequences. Some providers who would in principle prefer to operate within a regulated perimeter were deterred from seeking authorisation. Not by unwillingness to comply, but by the disproportionate cost of requirements designed for a different model of legal services provision. This creates a perverse outcome in which the regulatory framework discourages responsible actors from entering the regulated space.

A further concern raised by interviewees was enforceability. Regulatory bodies operate with limited resources relative to the scale of the B2C AI market. One interviewee offered a structural critique: AI companies operating at scale may calculate compliance risk on a cost-benefit basis, factoring in the expected cost of enforcement as a manageable business cost rather than a deterrent (Lawtech Company interview).

8.4 The Information-Advice Distinction

Key finding: The information-advice distinction, while useful as a regulatory organising principle, does not map onto how consumers experience AI legal tools. Users treat personalised AI outputs as advice regardless of how providers characterise the service. A more useful test is whether a tool provides bespoke, tailored output that a consumer relies on to make a decision. Tools covering non-reserved activities operate in a regulatory gap with no consumer protections equivalent to those in regulated practice.

The blurred information-advice boundary in AI contexts was a recurring theme across interviews – though it needs to be acknowledged that consumers often don't know the difference between information and advice. Several interviewees described deliberate product design choices to position their services as legal "explainers" or "information" tools rather than advice services. These choices are partly pragmatic (avoiding the regulatory requirements that attach to legal advice) and partly substantive, reflecting genuine differences in what the tools are designed to do.

However, interviewees on all sides of the debate acknowledged that consumers do not think in these terms. The information-advice distinction is a useful regulatory organising principle, but it primarily serves providers and regulators rather than users. Consumers experience a continuum of guidance and act on whatever they receive; the distinction does not map onto their reality. This asymmetry is itself a significant finding:

"The technology doesn't respect the information-advice distinction in the way that regulation does." (Consumer/Advice Body interview)

A consumer who receives a personalised response from an AI tool about their specific legal situation will treat it as advice regardless of how the provider characterises the service.

An industry association representative offered a functional test: the key regulatory question should be whether a tool provides:

"Bespoke, tailored output that a consumer relies on to make a decision." (Industry Association interview)

This formulation shifts focus from the provider's characterisation of the service to the consumer's experience of it, an approach more consistent with outcomes-based regulation and with the consumer protection objectives of the Legal Services Act 2007.

Interviewees also noted that in practice, users rarely present with a single, neatly defined legal issue. One provider described users arriving with intertwined legal, financial, housing, and personal problems that they may not recognise as legal at all (Lawtech Company interview). The question is not just whether an output counts as advice; it is whether AI tools can identify when a user needs professional help and direct them to it.

A related point concerns non-reserved legal activities. Tools covering employment, housing, or family matters short of court proceedings face no regulatory requirement under the Legal Services Act 2007. Some providers in these areas have not sought SRA authorisation, often because the costs outweigh the perceived benefit. As a result, consumers using these tools have none of the protections that come with professional regulation (Lawtech Company interview).

8.5 Consumer Perspectives

Key finding: Consumer-facing evidence reinforces the urgency of minimum standards for transparency, redress, and domain expertise. There is no clear complaint pathway for consumers who receive poor guidance from an AI legal tool outside the regulated perimeter. AI legal tools are already deeply embedded in consumer information-seeking behaviour – with or without standards to govern their quality.

Interview evidence reinforces the urgency of several of the gaps identified through the document review:

- **Transparency/user protection:** *"Any general AI tool should have to tell people upfront: I'm not a lawyer, I can't give you legal advice, you need proper legal help"* (Consumer/Advice Body interview)
This points to a baseline expectation of consumer-facing transparency that is not consistently reflected in current standards.
- **Redress:** *"Consumers have no idea who to complain to if an AI legal tool gives them bad information"* (Consumer/Advice Body interview)
This highlights the absence of a clear complaint pathway. Within the existing architecture, the Legal Ombudsman (LeO) covers only authorised persons, meaning most B2C AI legal tools fall outside its remit. The parallel with financial services (where the FOS similarly covers only FCA-regulated firms) is instructive, but the gap in legal services is wider: many common consumer legal problems fall within unreserved activities. This is a structural feature of the system that standards alone may not address.
- **Expertise:** *"If you're building an AI tool specifically for housing possession cases, you need qualified housing lawyers involved in building and checking it"* (Consumer/Advice Body interview)
This reflects a broader expectation that tools performing legal functions should demonstrate appropriate domain expertise. One interviewee noted that more effective B2C lawtech providers often combine legal and technical expertise, or lived experience of the issues addressed:
"That experience and expertise is very difficult to replicate by a pure technologist that just sees an opening in the market"(Legal Sector Body interview).
- **Consumer demand:**
Evidence from providers and advice organisations indicates that AI tools are already widely used by consumers. One provider described users signing up late at night when traditional advice is unavailable; another reported significant uptake driven by unmet need among parents of children with special educational needs. Citizens Advice advisers also report clients arriving having already consulted general AI tools, with results of variable quality. This

highlights a mismatch between the level of consumer reliance on AI tools and the current level of standards and safeguards governing them.

8.6 Design Choices and Quality Assurance in Practice

Key finding: In the absence of formal standards, lawtech providers have developed their own informal quality frameworks through product design and quality assurance practice. Consistent patterns include deliberate scope limitation, trusted source curation using retrieval-augmented generation, data minimisation beyond GDPR requirements, and interaction design tailored to vulnerable users. Quality assurance practice ranges from no formal legal review to sophisticated AI-assisted monitoring. These informal standards reveal both the priorities providers themselves consider important and where formal standards are most needed.

A significant finding from the interviews is the extent to which lawtech providers have, in the absence of formal standards, developed their own approaches to risk management through product design choices and quality assurance practice. These choices function as informal standards and reveal both what providers themselves consider important and where they perceive the most significant risks. The documentary analysis identifies quality assurance as an area with significant gaps in existing standards (section 4.4); the interview evidence confirms that practice is varied and, in the absence of agreed standards, has developed through individual provider choices. Several consistent patterns emerged across interviewees, suggesting an emerging informal consensus about what responsible AI lawtech looks like:

- **Scope limitation:** Multiple providers had deliberately constrained their tools to specific legal domains or use cases, treating narrow scope as a quality control strategy. One described their approach as "deliberately narrow in scope, confined to a specific legal domain and a specific use case" as a means of maintaining domain expertise without the resources required for a broad-scope system (Lawtech Company interview). Another limited their tool to a single area of education law, trained exclusively on authoritative government and statutory sources (Lawtech Company interview). A third focused exclusively on small debt recovery proceedings, exploiting the fixed, court-mandated process as a quality constraint (Lawtech Company interview). This pattern connects to the documentary finding on the relationship between system type and oversight model (section 5.1): specialist, narrow-scope tools are more amenable to quality assurance and domain expertise requirements than broad-scope general legal tools.
- **Trusted source curation:** All interviewees whose products provide substantive legal information described using retrieval-augmented generation or equivalent approaches. Drawing only on curated authoritative sources rather than open-web content. Sources used included gov.uk, legislation and case law from official archives, and statutory procedural guidance. One provider maintained a database of over 10,000 UK-specific documents, carefully selected through a defined quality process and updated regularly (Lawtech Company interview). Another trained exclusively on government and statutory materials (Lawtech Company interview). A third used a three-step verification process in which the system parses the user's question to ensure contextual understanding, checks the response against trusted sources, and conducts internal fact-checking before human review (Consumer/Advice Body interview). One interviewee argued that this should be a standard requirement: "responses should be grounded in trusted legal sources rather than general internet information. Some of the AIs at the moment, they'll crawl the net to get their information" (Lawtech Company interview). The curation of source material is an area where informal practice appears to be ahead of formal standards. This approach limits but does not eliminate risk of error:

authoritative sources can themselves be outdated or ambiguous. No interviewee described a mechanism for systematic identification of outdated source material.

- **Data minimisation as design principle:** Several providers made deliberate choices to minimise data collection and retention beyond GDPR compliance requirements. One product stores no data, requires no login, and does not permit document upload (Lawtech Company interview). Another redacts all personally identifiable information before data reaches a large language model (Consumer/Advice Body interview). A third built its entire infrastructure on Microsoft Azure (a deliberate choice based on compliance credentials), using only one additional third-party data processor (Pinecone), with all other processing in-house to minimise points of data exposure (Lawtech Company interview). These choices reflect a considered view of appropriate data handling in a legal context, and represent a standard of data protection above the regulatory minimum.
- **Interaction design for vulnerable users:** Several interviewees invested significantly in tone and interaction design to meet the specific needs of their user base. One designed their interface to be low-stimulation for neurodivergent users, with a communication style calibrated to be "compassionate and understanding" (Lawtech Company interview). Another deliberately built against what they described as sycophancy – the tendency of AI to tell users what they want to hear - designing for a more measured, honest communicative style (Lawtech Company interview). This dimension of product design is entirely absent from the formal standards landscape. Yet it may be highly consequential for user outcomes, particularly for vulnerable users who may be more susceptible to over-reliance on AI outputs.
- **AI-assisted monitoring:** One interviewee had built real-time AI monitoring tools to detect anomalous outputs, with alerts delivered immediately to staff:

"Ironically you probably get to a world where the AI monitoring system is better than a human monitoring system." (Lawtech Company interview)

This was a view supported by direct experience of cases where the monitoring AI had identified errors missed by human review. This represents an emerging quality assurance model in which AI quality is itself monitored by AI: offering the possibility of continuous, systematic monitoring at scale.

As regards quality assurance, several interviewees noted the absence of any agreed external standard against which to measure accuracy, and the difficulty of defining accuracy in a context where outputs are necessarily context-dependent. One argued for explicit uncertainty communication as a standard requirement:

"It often gives answers that sound really too certain and you almost believe too much in the outcome they're saying and it's not always correct. So the AI should be explaining uncertainty if it's there." (Lawtech Company interview)

One provider had specifically designed their tool to resist confirming what users want to hear (Lawtech Company interview). These approaches address a gap in the documentary analysis: current standards do not address the communication of uncertainty, which may be as important as the accuracy of outputs themselves.

A further challenge was the absence of formal legal review in some products. One provider confirmed that outputs had not been formally reviewed by qualified solicitors, though they had been exposed to a large group including legal professionals for informal feedback (Lawtech Company interview). The range across interviewees (from no formal legal review, through informal exposure to practitioners, to

ongoing oversight by qualified lawyers) illustrates the gap between the domain expertise requirements recommended by bodies such as the Law Society and current practice. The use of mandatory human review as a quality assurance mechanism (and the question of how much oversight is appropriate and in what form) is discussed in section 8.7 below.

8.7 Human Oversight and Phased Automation

Key finding: Human oversight is currently considered necessary by all interviewees, but this position is explicitly transitional. A phased, evidence-based reduction in oversight represents one emerging model. The critical challenge is ensuring oversight requirements specify quality and depth, not merely formal presence: perfunctory sign-off provides accountability without substantive protection.

The question of human oversight – how much, of what kind, for how long – emerged as one of the most substantively important themes across the interviews. The documentary analysis identifies human oversight as a cross-cutting concern in the standards landscape (section 4.3.4). The interview evidence adds significant texture.

The consensus across interviewees was that human oversight is currently necessary, but that this position is explicitly transitional. Multiple providers described their current oversight arrangements as interim – reflecting the current state of AI capability and available evidence, rather than a permanent design feature.

One provider described this most explicitly: a phased reduction in human oversight, with evidence-based review triggers and quarterly regulatory dialogue. The aspiration was a fully automated system; the path to that endpoint was governed by accumulated quality data:

"I think we would end up in a world where there will always be some degree of oversight from a human, but it will be very much risk-based" (Lawtech Company interview)

This represents a sophisticated engagement with the oversight question: rather than a blanket requirement for human review, an evidence-based, graduated reduction as confidence in system performance develops - a model that standards frameworks could usefully adopt.

Other providers described human oversight as a non-negotiable design constraint:

"Human in the loop is non-negotiable." (Consumer/Advice Body interview)

Their system included active counter-measures against automation bias, including a quality score against a national quality framework generated alongside each output, and randomised confirmation prompts to prevent advisers from approving without genuine scrutiny. The same interviewee emphasised audit trail requirements: even where AI significantly augments the process, "the critical judgement is still a person who's made that and we've got the audit trail of who made that judgment."

A nuanced challenge raised by one interviewee was the question of what human oversight is for. If oversight is intended to catch errors, its value depends entirely on the knowledge and attention of the human reviewer. If advisers are approving AI outputs perfunctorily, the oversight requirement provides formal accountability without substantive protection. This connects to the documentary finding on accountability standards (section 4.4.2): existing standards typically require human oversight without specifying its quality, depth, or documentation requirements.

The interviews also highlighted an emerging divergence between products where human oversight is structurally embedded and products that function effectively as information retrieval systems without individual output review. For the latter category, oversight takes the form of source curation, output monitoring, and user-accessible verification rather than human review of each interaction. Whether

this constitutes adequate oversight for consumer protection purposes is a question the standards landscape does not currently address.

8.8 Consumer Education and Expectation Management

Key finding: Consumer education is as important as formal standards in reducing the risks of AI legal tools. Where users understand what AI can and cannot do, they use it more effectively and are better protected. Integrating AI legal literacy into education from an early age was proposed as a near-term necessity. Transparency standards should extend beyond generic disclaimers to cover the specific limitations of AI in different legal contexts.

A theme that emerged with notable consistency, and which is not adequately addressed in the current documentary analysis, was the importance of consumer education alongside or as an alternative to formal standards.

Multiple interviewees argued that the risks of AI legal tools are substantially mediated by whether users understand what they are interacting with and how to use it appropriately. One provider described users who had developed sophisticated complementary strategies, consulting the AI tool alongside a solicitor and using outputs from one to validate the other (Lawtech Company interview). This level of informed usage is the analogue of the "informed consumer" that consumer protection frameworks are designed to produce – but it was achieved through user experience and personal research, not through product design or regulatory requirement.

As noted in section 8.5, clients are arriving at advice services like Citizens Advice having already consulted ChatGPT, with results of variable quality. This illustrates the practical consequences of the consumer education gap: where users do not understand the limitations of AI tools or how to verify outputs, AI does not supplement professional advice; it pre-empts it, sometimes with harmful effects. Standards requiring transparency and disclosure address part of this gap but do not substitute for broader public understanding of what AI tools can and cannot do.

One interviewee made the case for integrating AI legal literacy into public education from an early age:

"If we can help people to understand what access to justice means at 14 and how they can get it, it's better than trying to access it and having no idea how to do it when you're 18 or 20."
(Legal Sector Body or Professional Organisation interview)

By the time young people currently in secondary school become adults, "everything is going to have an AI footprint." Starting the conversation about AI and legal rights in schools is not a long-term aspiration but a near-term necessity.

A related point concerned expectation management. One interviewee distinguished between what AI tools are good at (operational, document-oriented, process-structured tasks) and where they are likely to produce poor results (complex, contested, strategic legal matters with determined opposition). The risk is not only that users rely on AI for things it cannot do, but that they commit to courses of action – including litigation – on the basis of AI assessments that are operationally competent but strategically inadequate (Legal Sector Body or Professional Organisation interview). Transparency standards should extend to transparency about the limitations of AI in different legal contexts, not just a generic disclaimer.

One interviewee proposed "first aid law" as a conceptual organising principle for consumer-facing AI legal standards: the equivalent of a domestic first aid kit, providing accessible guidance for common legal situations, with clear signposting to professional help when the situation exceeds what first aid can address (Legal Sector Body or Professional Organisation interview). This framing connects AI

standards to the broader access-to-justice agenda: the goal is not to replace professional services but to provide a first layer. Of accessible, trustworthy guidance that helps consumers understand their situation and navigate toward appropriate support.

8.9 Beyond Standards

Key finding: Regulatory pilots and sandboxes are the most consistently endorsed mechanism for improving quality and accountability beyond formal standards, generating first-party evidence about how AI tools operate in practice. Cross-sector collaboration and honest public communication about the purpose and limits of AI legal tools are seen as essential complements to standards development.

Several interviewees proposed or endorsed mechanisms beyond formal standards as means of improving the quality and accountability of AI legal tools:

- **Regulatory pilots and sandboxes:** The most consistently endorsed alternative mechanism was the use of regulatory pilots and sandboxes, allowing regulators and developers to engage in a structured environment. The logic is that pilots generate first-party evidence about how AI tools operate in consumer interactions, rather than hypothetical assessments:

"The best way to make products safer is to learn about it. And the best way to learn about it is through pilots, because through pilots that's where you're going to get first-party information about how technology works, how consumers are using it."(Lawtech Company interview)

Standards that emerge from observed practice are more likely to be accurate, proportionate, and enforceable than those derived from advance specification. An international precedent cited was the Ontario law society's "Access to Innovation" sandbox, in which companies engage with the regulator in a collaborative, lighter-touch environment (Lawtech Company interview).

One provider had independently proposed a sandbox approach to the SRA, including an application for an insurance waiver commensurate with actual risk rather than the regulatory minimum. The SRA indicated it had never granted such a waiver. A flexible sandbox framework that allowed providers to trial innovative approaches under regulatory observation – with tailored rather than blanket obligations – would make it more feasible for well-intentioned actors to enter the regulated perimeter, increasing the proportion of the market operating to professional standards (Lawtech Company interview).

- **Convening and cross-sector collaboration:** One interviewee emphasised the convening power of the LSB as distinct from its standard-setting function:

"This is too big an issue for one organisation to consider." (Legal Sector Body or Professional Organisation interview).

Many of the risks associated with AI legal tools are shared with AI tools providing financial advice, medical information, and other high-stakes consumer guidance. Standards developed in the legal sector would benefit from dialogue with the FCA (consumer duty), potentially the CQC, and the Advertising Standards Authority. The consumer protection principles at stake are common across regulated advice services; joined-up standards would reduce compliance burden for providers and provide consumers with consistent expectations across their different legal, financial, and health needs.

- **Government-embedded AI services:** One interviewee noted that the Ministry of Justice was exploring embedding AI tools within government-facing services, including gov.uk interfaces for specific legal processes. This approach would place AI within a publicly accountable infrastructure, with government taking direct responsibility for accuracy and consumer

protection (Legal Sector Body or Professional Organisation interview). It does not address the broader B2C lawtech market, but it illustrates an alternative model: rather than regulating a private market through standards, government could provide a publicly trusted alternative that anchors quality expectations for the wider market.

- **Communication and public narrative:** One interviewee argued that the most important non-regulatory intervention was a more honest public conversation about the purpose of AI tools:

"We need to have a really transparent conversation with the people who work in the industry about what we're trying to do and what we're not trying to do. This isn't trying to replace expertise or skills...it's ensuring that the people who need our help can get it."

(Consumer/Advice Body interview).

The absence of this narrative creates conditions in which public anxiety about AI displacement co-exists with under-regulated proliferation of AI tools of variable quality. Standards development is most effective when situated within a broader public conversation about the purpose and limits of AI in legal services.

9. Conclusions and Implications

9.1 Answering the Research Questions

RQ1: Are there any standards that apply to B2C lawtechs?

Yes. This review assessed 234 documents from 70 organisations and identified 160 of these to contain insights relevant to consumer-facing legal technology. However, the vast majority of these standards are general professional, regulatory, or technical instruments that apply to AI broadly or to sectors other than legal services. Standards specifically designed for consumer-facing lawtech are almost entirely absent. The protections available to consumers of AI legal tools are derivative, inherited from general data protection law, from horizontal AI governance frameworks, or from professional regulation that applies to the regulated practitioner rather than to the consumer-facing tool. The consumer using an unregulated AI tool for legal guidance is protected primarily by the GDPR and general consumer protection law, neither of which addresses the specific risks of AI-generated legal content.

RQ2: Who sets these and how are they developed?

Standards are set across multiple levels of governance. International bodies (ISO/IEC, NIST, IEEE) develop technical standards through consensus-based processes involving national standards bodies and expert committees. Supranational regulators, principally the European Union, develop binding legislation through parliamentary procedures. UK cross-sector regulators, the ICO, CMA, FCA, Ofcom, develop guidance and enforceable requirements within their respective mandates, increasingly coordinated through the DRCF. Legal services bodies, the SRA, the Law Society, the BSB, develop professional guidance applicable to authorised persons. AI platform providers, OpenAI, Google, Microsoft, Anthropic, develop usage policies applicable to users of their platforms. Each category operates through different processes with different levels of stakeholder engagement, democratic legitimacy, and enforcement power. The result is a fragmented landscape in which no single body has comprehensive responsibility for AI in legal services, let alone B2C lawtech.

RQ3: What format do they take and how are they communicated?

The standards take predominantly soft forms, with guidance documents accounting for 58% of the documents. Communication is fragmented: formal ISO/IEC standards are behind paywalls, regulatory guidance is scattered across multiple organisational websites, and platform usage policies are updated unilaterally without notification to affected users. The interview programme confirmed that most practitioners are unaware of existing standards, a finding that fundamentally undermines the standards' capacity to influence behaviour. A standard that is not known cannot be followed.

RQ4: What do the standards cover?

Coverage is uneven across the 17 thematic areas identified in this review. Data protection and transparency receive most attention, benefiting from established regulatory infrastructure and extensive cross-sector attention. Bias, accountability, human oversight, and privacy are moderately well covered, with clear principles but variable operational specificity. Vulnerability identification, signposting and referral, redress mechanisms, accessibility, consistency, and reliability have the thinnest coverage. This pattern has great consequences for users: the themes with the weakest coverage are precisely those most directly relevant to protecting consumers who use AI tools for legal purposes without professional intermediation.

RQ5: Are there any obvious gaps?

The gaps are substantial and come with great consequences. At the thematic level, vulnerability identification, signposting and referral, and redress mechanisms represent critical gaps. These are

areas where consumer harm is already occurring, where feasible solutions exist in other sectors, and where the standards landscape provides no adequate response. At the structural level, the information-advice boundary remains operationally undefined for AI contexts. Unregulated providers operate outside the reach of professional regulation, and the assumption of professional intermediation that underpins most existing protections does not hold for consumer-facing AI tools. The gap between the maturity of standards in healthcare and financial services and the relative underdevelopment of standards in legal services is significant and represents an area where cross-sector learning could accelerate progress.

9.2 Potential implications

The following observations emerge from the evidence gathered in this review:

1. This review identifies key gaps (namely vulnerability identification frameworks, signposting and referral protocols, mechanisms for redress of AI-related complaints, and clear consumer-facing transparency requirements). Because these areas currently present the greatest consumer risk while workable models already exist in other sectors.
2. There are existing standards in other sectors that could be adapted to the needs of AI-powered lawtech. Operational frameworks provide tested structures that can be tailored to legal services rather than invented anew. For example, from healthcare (such as pre-market validation, escalation pathways for vulnerable users, and controlled AI modification processes). Also, from financial services (including conduct-of-business rules, senior manager accountability, and outcomes-based consumer protection duties).
3. The evidence highlights a “gatekeeper gap” by protecting consumers who interact directly with AI tools outside regulated professional practice, not only those using AI through lawyers. This could include voluntary accreditation schemes through which providers demonstrate compliance with recognised standards.
4. Risks vary significantly across different types of B2C lawtech tools. Tools that provide personalised legal guidance directly to consumers, particularly in higher-risk or sensitive contexts, present different risks from those offering general legal information or operating under professional supervision. This points to the potential value of approaches that differentiate expectations according to use case and risk.
5. The traditional distinction between legal information and legal advice is difficult to apply in AI contexts. Many tools generate outputs that are tailored to a user’s specific situation, regardless of how the service is labelled. This suggests a need for a more practical way of distinguishing between the two, potentially based on whether users are likely to rely on the output when making decisions.

Annexes

Annex A: Glossary

Artificial Intelligence (AI): Technology that enables computer systems to perform tasks that normally require human intelligence, including natural language processing, pattern recognition, and decision-making. In this report, AI encompasses machine learning, deep learning, and large language models.

B2C Lawtech: Business-to-consumer legal technology, AI-powered tools that deliver legal information, guidance, or advice directly to consumers without requiring the intermediation of a regulated legal professional.

Formal Standard: A document established by consensus and approved by a recognised standardisation organisation (e.g., ISO, IEEE, BSI), providing rules, guidelines, or characteristics for activities or their results.

High-Risk AI System: Under the EU AI Act, an AI system that falls within categories specified in Annex III, including systems used in the administration of justice and democratic processes, subject to mandatory conformity assessment and enhanced governance requirements.

Large Language Model (LLM): A type of AI model trained on large volumes of text data that can generate, summarise, translate, and reason about natural language. Examples include GPT-4 (OpenAI), Claude (Anthropic), and Gemini (Google).

Lawtech: Technology applied to legal services delivery, encompassing tools used by legal professionals and those used directly by consumers.

Reserved Legal Activity: Under the Legal Services Act 2007, activities that may only be carried out by authorised persons: rights of audience, conduct of litigation, reserved instrument activities, probate activities, notarial activities, and administration of oaths.

Voluntary Standard: A standard that organisations may choose to adopt but are not legally required to follow, though adoption may be incentivised through market expectations, procurement requirements, or regulatory recognition.

Annex B: References

AI Safety Institute (2024) *Inspect: A framework for large language model evaluations*. AI Safety Institute. Available at: <https://www.aisi.gov.uk>.

AlgorithmWatch (2025) *Automating Society Report 2025*. AlgorithmWatch. Available at: <https://automatingsociety.algorithmwatch.org>.

Bogiatzis-Gibbons, D. et al. (2024) *A literature review on bias in supervised machine learning*. Financial Conduct Authority.

Brancati, C.U. (2025) *AI Governance in the UK: Regulatory Landscape and Emerging Frameworks*. The Alan Turing Institute.

British Standards Institution (2024) *PAS 2816:2024 Responsible AI Management*. London: BSI.

British Standards Institution (n.d.) *BS 30440:2024 Regulatory technology (RegTech) (AI-enabled regulatory compliance) Guidelines*. London: BSI.

Brownsword, R. (n.d.) *The Regulation of New Technologies in Professional Service Sectors in the United Kingdom: Key Issues and Comparative Lessons*. Legal Services Board.

Byrom, N. (2024) *Protecting consumers and access to justice in the age of AI: Mapping the emerging policy, regulatory and advocacy landscape around the use of Generative AI Applications and case outcome predictive technologies in legal services*. The Law Society of England and Wales.

Campbell, F. (2025) *Generative AI in legal disclosure: a practical guide*. The Law Society. Available at: <https://www.lawsociety.org.uk/topics/civil-litigation/generative-ai-in-legal-disclosure-guide>.

CCBE (2025) *CCBE guide on the use of generative AI by lawyers*. Council of Bars and Law Societies of Europe.

CDEI (2020) *Review into bias in algorithmic decision-making*. Centre for Data Ethics and Innovation.

CDEI (2021) *The roadmap to an effective AI assurance ecosystem*. Centre for Data Ethics and Innovation.

CEN-CENELEC (n.d.) *Standardisation request to support the AI Act*. European Committee for Standardization / European Committee for Electrotechnical Standardization.

CMA (2021) *Algorithms: How they can reduce competition and harm consumers*. Competition and Markets Authority.

Consumers International (2024) *Our vision for Fair and Responsible AI for Consumers*. Consumers International

De Rechtspraak (2024) *Responsible and innovative: AI for a fair Dutch judicial system*. De Rechtspraak. Available at: <https://www.rechtspraak.nl>.

de Souza, S. (2018) *Digital technology in the realm of law*. International Council for Online Dispute Resolution.

Department for Business and Trade (2025) *Backing your business: our plan for small and medium sized businesses*. London: DBT.

Department for Science, Innovation and Technology (2023) *A pro-innovation approach to AI regulation*. London: DSIT.

Department for Science, Innovation and Technology (2024) *Implementing the UK's AI Regulatory Principles: Initial Guidance for Regulators*. London: DSIT.

Department for Science, Innovation and Technology and Office for Artificial Intelligence (2023) *A pro-innovation approach to AI regulation, white paper*. London: DSIT.

EDPB and EDPS (2026) *Joint Opinion on simplification of the AI Act*. European Data Protection Board / European Data Protection Supervisor.

EDPS (2024) *Guidelines on the use of AI by EU institutions*. European Data Protection Supervisor.

EDPS (2025) *Annual Report on AI and data protection*. European Data Protection Supervisor.

European Commission (2019) *Policy and investment recommendations for trustworthy AI*. European Commission.

European Commission (2022) *Proposal for a Directive on adapting non-contractual civil liability rules to artificial intelligence (AI Liability Directive)*. European Commission.

European Commission (2024) *Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*. Official Journal of the European Union.

European Commission (2025) *Guidelines on prohibited AI practices*. European Commission.

European Commission (2026) *Guidelines on the definition of an AI system*. European Commission.

European Parliament and Council (2016) *Art. 22 GDPR, Automated individual decision-making, including profiling*. General Data Protection Regulation.

EU Fundamental Rights Agency (2022) *Bias in algorithms, Artificial intelligence and discrimination*. European Union Agency for Fundamental Rights.

Federation of Small Businesses (2024), *Tied Up: Unravelling the Dispute Resolution Process for Small Firms*.

Financial Conduct Authority (2022) *Consumer Duty*. London: FCA. Available at: <https://www.fca.org.uk/firms/consumer-duty>.

Financial Conduct Authority (2024) *AI in financial services: Framework and guidance*. London: FCA.

GDS (2023) *Data Ethics Framework*. Government Digital Service. Available at: <https://www.gov.uk/government/publications/data-ethics-framework>.

Groves, L. et al. (2023) *Going beyond the AI Ethics principles*. Centre for Data Ethics and Innovation.

Harkavy, R. (2025) *Public wary of robot lawyers without human oversight, survey finds*.

High-Level Expert Group on AI (2019) *Ethics guidelines for trustworthy AI*. European Commission.

High-Level Expert Group on AI (2020) *Assessment List for Trustworthy AI (ALTAI)*. European Commission.

ICO (Information Commissioner's Office) (2023) *Guidance on AI and data protection*. Available at: <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/>.

ICO (n.d.) *Rights related to automated decision making including profiling*. Available at: <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/individual-rights/>.

ICO and The Alan Turing Institute (2020) *Explaining decisions made with AI*. Available at: <https://ico.org.uk/for-organisations/guide-to-data-protection/key-dp-themes/explaining-decisions-made-with-ai/>.

ISO/IEC (2021) *ISO/IEC TR 24027:2021 Information technology (Artificial intelligence) Bias in AI systems and AI aided decision making*. Geneva: ISO.

ISO/IEC (2023) *ISO/IEC 23894:2023 Information technology (Artificial intelligence) Guidance on risk management*. Geneva: ISO.

ISO / IEC (2023) *ISO/IEC 25010:2023 Systems and software engineering – Systems and software Quality Requirements and Evaluation (SQuaRE) – Product quality model*. Geneva: ISO.

ISO/IEC (2024) *ISO/IEC 22989:2024 Information technology (Artificial intelligence) Artificial intelligence concepts and terminology*. Geneva: ISO.

ISO/IEC (2024) *Controllability of automated AI systems*. Geneva: ISO.

ISO/IEC (2021) *ISO/IEC 25059:2023 Software engineering (Systems and software Quality Requirements and Evaluation (SQuaRE)) Quality model for AI systems*. Geneva: ISO.

ISO/IEC TR 29119-11 (2020) *Software and systems engineering (Software testing) Part 11: Guidelines on the testing of AI-based systems*. Geneva: ISO.

Law Society (2025) *Introduction to lawtech*. London: Law Society.

Law Society (2021) *Lawtech and ethics principles report*. The Law Society.

Law Society of Scotland (2024) *Guide to Generative AI*. Law Society of Scotland

LawtechUK (2024) *The Regulatory Response Unit*, LinkedIn, 4 November.

Law Society of Scotland (n.d.) *Guide to Generative AI*. Available at: <https://www.lawscot.org.uk/members/business-support/lawscottech/resources/guide-to-generative-ai/>.

Legal Services Board (2020) *Perspectives on Lawtech and Regulation*. Legal Services Board.

Legal Services Board (2023) *Legal needs survey 2023*. Legal Services Board.

Legal Services Board (2024) *Delivering a pro-innovation approach to AI regulation, an outline of the LSB's approach*. Legal Services Board.

Legal Services Consumer Panel (2019) *Lawtech and consumers*. Legal Services Board.

Levene, M. et al. (2024) *A Lifecycle Approach to AI Trustworthiness*. National Physical Laboratory.

LTUK Ecosystem Tracker (2024) *An Exciting Insight into The Present and Future of Lawtech*. London: Lawtech UK.

Medicines and Healthcare products Regulatory Agency (2023) *Software and AI as a Medical Device*. London: MHRA.

Medicines and Healthcare products Regulatory Agency (2023) *Predetermined change control plans for machine learning-enabled medical devices*. Medicines and Healthcare products Regulatory Agency.

Medicines and Healthcare products Regulatory Agency (2024) *Impact of AI on the regulation of medical products*. Medicines and Healthcare products Regulatory Agency.

Medicines and Healthcare products Regulatory Agency, U.S. Food and Drug Administration and Health Canada (2024) *Transparency for machine learning-enabled medical devices: guiding principles*. Medicines and Healthcare products Regulatory Agency.

Mello, M.M. (2025) *AI regulation and the health sector*. New England Journal of Medicine.

Ministry of Justice (2025) *AI Action Plan for Justice*. London: Ministry of Justice.

NAIC (2024) *Model Bulletin on the Use of AI by Insurers*. National Association of Insurance Commissioners.

Natalie Le (2025) *Data protection in AI contexts*. Data Protection Commissioner.

National Institute of Standards and Technology (2023) *AI Risk Management Framework (AI RMF 1.0)*. Gaithersburg, MD: NIST.

National Institute of Standards and Technology (2024) *NIST AI 600-1: Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*. Gaithersburg, MD: NIST.

National Institute of Standards and Technology (2025) *Adversarial Machine Learning: A Taxonomy and Terminology of Attacks and Mitigations (NIST AI 100-2e2024)*. Gaithersburg, MD: NIST.

NIST (2022) *Towards a Standard for Identifying and Managing Bias in Artificial Intelligence*. Gaithersburg, MD: NIST.

NTIA (2024) *AI Accountability Policy Report*. National Telecommunications and Information Administration.

OAIC (2024) *Guidance on privacy and AI*. Office of the Australian Information Commissioner.

OECD (2024) *OECD AI Principles (updated)*. Organisation for Economic Co-operation and Development.

OAI/CDEI (2019) *Initial guidance for regulators*. Office for Artificial Intelligence / Centre for Data Ethics and Innovation.

PDPC Singapore (2020) *Model AI Governance Framework*. Personal Data Protection Commission of Singapore.

Regulatory Horizons Council (2022) *The Regulation of Artificial Intelligence as a Medical Device*. Regulatory Horizons Council.

Queensland Law Society (2024) *No.37 Artificial Intelligence in Legal Practice*. Queensland Law Society.

RACGP (2025). *Artificial intelligence (AI) scribes*. Royal Australian College of General Practitioners.

RACGP (2025). *Conversational artificial intelligence (AI)*. Royal Australian College of General Practitioners.

RACGP (2025). *Artificial intelligence in primary care*. Royal Australian College of General Practitioners.

Solicitors Regulation Authority (2019) *SRA Code of Conduct for Firms*. Available at: <https://www.sra.org.uk/solicitors/standards-regulations/code-conduct-firms/>.

Solicitors Regulation Authority (2023) *Risk Outlook report: The use of artificial intelligence in the legal market*. SRA.

Solicitors Regulation Authority (2026) *Compliance tips for solicitors regarding the use of AI and technology*. Available at: <https://www.sra.org.uk/solicitors/resources/innovate/compliance-tips-for-solicitors/>.

The Law Society (2025) *Generative AI: the essentials*. Available at: <https://www.lawsociety.org.uk/topics/ai-and-lawtech/generative-ai-the-essentials>.

Tulk, S. et al. (2020) *AI and accessibility*. Microsoft Research.

UK Jurisdiction Taskforce (n.d.) *Public consultation, Liability for AI Harms under the private law of England and Wales*. LawtechUK.

Unwilderred editorial team (2026) *AI Legal Tools comparison for UK Consumers*. Unwilderred.

U.S. Food and Drug Administration (2025) *Artificial Intelligence-Enabled Device Software Functions: Lifecycle Management and Marketing Submission Recommendations*; Draft Guidance for Industry and Food and Drug Administration Staff. Silver Spring, MD: FDA.

US Food and Drug Administration (2026) *Artificial Intelligence-Enabled Medical Devices*.

VDI (2022) *Ethical principles for the development and application of AI*. Verein Deutscher Ingenieure.

Wawrzyszczuk, A. (2025) *AI B2C Blueprint: Transforming Consumer Legal Services*. LawtechUK.

Webley, L. (2020) *Ethics, Technology and Regulation*. Legal Services Board.

WHO (World Health Organization) (2021) *Ethics and governance of artificial intelligence for health*. Geneva: WHO.

Winfield, A. et al. (2021) *IEEE P7001: A proposed standard on transparency*. IEEE.

ⁱ AI-powered business-to-consumer (B2C) lawtech refers to AI tools that give legal information or support directly to members of the public, helping them understand a problem, explore options, or take action without needing a lawyer at the outset. While these tools are primarily designed for consumers, some may signpost to legal professionals or intermediaries such as advice agencies or charities. This is different from most AI tools in the legal sector, which are business-to-business (B2B) products aimed at professional firms and legal practitioners, which are designed to make legal teams more efficient.

ⁱⁱ For the purposes of this report, the term 'standards' is used broadly to encompass formal standards (issued by recognised standards bodies), regulations, guidance, codes of practice, platform policies, and assurance schemes that establish expectations for the development, deployment, governance, and quality of AI-powered tools.

ⁱⁱⁱ <https://www.gov.uk/government/publications/ai-action-plan-for-justice/ai-action-plan-for-justice>

^{iv} <https://www.gov.uk/government/publications/ai-regulation-a-pro-innovation-approach>

^v ICO, Legal Services Operational Privacy Certification Scheme, available at <https://ico.org.uk>.

^{vi} There are other definitions. The Law Society expands the definition of lawtech as "technology that helps improve or automate legal work" where a piece of technology "supports, supplements or replaces traditional methods for delivering legal services, or improves the way the justice system operates" (The Law Society, 2025. Available at: <https://www.lawsociety.org.uk/topics/ai-and-lawtech/introduction-to-lawtech>).

^{vii} A rating of "substantive" (Yes) was assigned where a document engages meaningfully with the theme, for example by providing dedicated sections, detailed analysis, or specific guidance; these instances were scored as 2 points. A "partial" rating was applied where the theme is mentioned or briefly addressed without sustained treatment (e.g. passing references or high-level statements), and was scored as 1 point. Documents with no relevant engagement were scored 0.

^{viii} <https://www.iso.org/standard/42001>

^{ix} <https://www.iso.org/standard/77304.html>

^x <https://www.iso.org/standard/81118.html>

^{xi} <https://ieeexplore.ieee.org/document/9536679>

^{xii} <https://www.nist.gov/itl/ai-risk-management-framework>

^{xiii} <https://doi.org/10.6028/NIST.AI.600-1>

^{xiv} <https://www.gov.uk/government/publications/ai-regulation-a-pro-innovation-approach/white-paper>

^{xv} <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/guidance-on-ai-and-data-protection/>

^{xvi} <https://www.fca.org.uk/firms/innovation/ai-approach>

^{xvii} For full details see <https://legalservicesboard.org.uk/wp-content/uploads/2024/04/Technology-and-innovation-guidance-for-publication.pdf>

^{xviii} Examples include:

CCBE (2025) *CCBE guide on the use of generative AI by lawyers*. Council of Bars and Law Societies of Europe. Law Society (2025) Introduction to lawtech. London: Law Society.

LTUK Ecosystem Tracker (2024) 'An Exciting Insight into The Present and Future of Lawtech'. London: Lawtech UK.

^{xix} It should be noted that a few B2C AI lawtech products are regulated because they have been developed by legal professionals or firms (for example, Garfield AI).

^{xx} Suresh, H. and Gutttag, J., 2021, 'A Framework for Understanding Sources of Harm throughout the Machine Learning Life Cycle', EAAMO '21, ACM). This paper was not included in the scope of this review beyond the bias classification it presents.

^{xxi} LSCP. 2025. 'Service Delivery Research'. Available at: <https://www.legalservicesconsumerpanel.org.uk/wp-content/uploads/2025/12/05-LSCP-Service-Delivery-Guideline-Development-Final-Report.pdf>

^{xxii} Previous work on vulnerability in legal services includes: The Law Society, 'Meeting the needs of vulnerable clients' (available at <https://www.lawsociety.org.uk/topics/client-care/meeting-the-needs-of-vulnerable-clients>); SRA, 'Consumer vulnerability in the legal market' (available at <https://www.sra.org.uk/sra/research-publications/consumer-vulnerability-legal-market/>); and the LSB's own rapid literature review on vulnerability conducted in 2014.

^{xxiii} It should be noted that the provision of legal advice is not restricted, and may be offered by regulated and unregulated providers alike. However, only certain reserved legal activities, such as advocacy before the courts, the conduct of litigation, notarial services, and probate activities, may be carried out by authorised individuals.

^{xxiv} These documents include UK GDPR Article 22, ICO Guidance on Automated Decision-Making and the EU AI Act.

^{xxv} The State of Legal Services 2020. Legal Services Board 2020 at https://legalservicesboard.org.uk/wp-content/uploads/2020/11/The-State-of-Legal-Services-Narrative-Volume_Final.pdf

^{xxvi} Maturity Score = $(\text{Yes} \times 2 + \text{Partial} \times 1) / (\text{Total Documents} \times 2) \times 100$. “Yes” (substantive coverage) is scored as 2, “Partial” as 1, and “No” as 0. The denominator represents the theoretical maximum ($234 \times 2 = 468$), producing a percentage scale where 100% indicates all documents substantively address the theme, 50% indicates universal partial coverage (or equivalent), and 0% indicates no coverage. Scores are grouped into four bands based on natural breaks in the distribution: “Well established” ($\geq 40\%$), “Developing” (25–40%), “Emerging” (15–25%), and “Sparse” ($< 15\%$).

^{xxvii} Figure 15 shows the coverage rate per theme i.e. the percentage of all 234 documents that address each theme either substantively or partially. Formula: Coverage Rate = $(\text{Yes} + \text{Partial}) / \text{Total Documents} \times 100$. Unlike the maturity score (Figure 14), this metric does not weight by depth — it simply measures how many documents engage with the theme at all. This means a document that briefly mentions bias counts the same as one with a detailed bias framework.