

JUSTICE AND THE LEGAL PROFESSION in the Era of New Technologies and AI

Courts
Lawyers
Compliance
Evidence
Professional Responsibility
Governance

Professional training deck — developments through 2025

2025




Neapolis University Pafos, Cyprus

AI-2-TRACE-CRIME

Jean Monnet Center of Excellence



Co-funded by
the European Union



What this session is designed to answer

- How are digital platforms, data analytics, automation and generative AI changing the administration of justice and the everyday work of lawyers, compliance teams and other legal professionals?
- Which uses create genuine efficiency or access-to-justice gains, and which uses create new risks for due process, equality, confidentiality, evidentiary reliability and professional independence?
- What does the emerging governance architecture — especially the EU AI Act, GDPR, Council of Europe standards and professional ethics — require in practice?
- How can an organisation move from ad hoc experimentation with AI to controlled, auditable and professionally responsible adoption?

The central question is not “AI or no AI?” but: which task, which system, which data, which safeguards, and who remains accountable?

A vocabulary for legal professionals

Technology stack

- Digitisation: converting analogue information and processes into digital form.
- Automation: rules-based execution of repetitive tasks, often without “learning”.
- Machine learning: statistical models learn patterns from data to classify, rank or predict.
- Generative AI: models generate new text, images, audio, code or other content from prompts.

Legal significance

- Different technologies create different legal and evidentiary risks; “AI” should not be used as a catch-all label.
- A system may combine several components: OCR, retrieval, predictive models, large language models and workflow automation.
- Risk depends on purpose and context, not only on technical sophistication.
- Legal professionals need enough technical literacy to ask the right questions — not to become data scientists.

From legal databases to generative AI: an accelerated trajectory



The novelty of generative AI is not merely faster search: it can synthesize, draft, transform and interact in natural language across an entire workflow.

Why the legal sector is unusually exposed to AI disruption

01**Text-intensive work**

Law is built around statutes, cases, contracts, correspondence, evidence and reasoned argument — precisely the type of unstructured information modern language models can process.

02**High cost of expertise**

A large share of legal cost reflects professional time spent locating, reviewing, comparing, summarising and drafting information. Automation therefore has an unusually visible economic effect.

03**High stakes**

Errors can affect liberty, property, reputation, regulatory exposure or fundamental rights. The tolerance for plausible-but-wrong output is therefore much lower than in ordinary office work.

04**Trust-based profession**

Confidentiality, independence, candour, competence and accountability are constitutive professional duties. Technology changes the means of work, but does not transfer those duties to the tool.

01

The Digital Transformation of Justice

From e-filing and remote hearings to AI-assisted research, triage, translation, prediction and decision support.

The digital court: a continuum, not a single technology

- Digital identity, e-filing, electronic case files, online payment and service of documents can reduce friction before any AI is introduced.
- Remote hearings and hybrid proceedings increase resilience and accessibility, but raise questions about equality of arms, effective participation, witness assessment and digital exclusion.
- Data infrastructures make case-law and court administration machine-readable, enabling search, analytics, anonymisation, translation and performance management.
- AI enters on top of this infrastructure: document classification, transcription, legal research, summarisation, scheduling, case triage, decision support and online dispute resolution.

Where AI is already being used around courts

Lower-risk / supporting functions

- Speech-to-text transcription and translation.
- Anonymisation or pseudonymisation of judgments.
- Document routing, classification and administrative triage.
- Search and retrieval across legislation and case law.
- Scheduling, workload analysis and court administration.

Higher-impact functions

- Prediction of litigation outcomes or likely case duration.
- Risk scoring, prioritisation or profiling of persons or cases.
- Decision-support systems recommending legal outcomes.
- AI-assisted assessment of evidence or credibility.
- Automated or semi-automated dispute resolution producing legal effects.

What the 2025 CEPEJ picture tells us

125**tools identified**

CEPEJ reported 125 tools in its Resource Centre, primarily from Europe, aimed at judicial efficiency and accessibility.

**Generative AI**

Machine learning and natural-language processing are increasingly important, with generative AI showing a notable rise in justice-sector use.

0**fully autonomous courts**

CEPEJ reported no fully autonomous AI systems capable of operating independently within courts; human oversight remains essential.

!**Accountability remains**

Many tools operate inside internal court IT systems. Lack of public visibility does not remove requirements of transparency, scrutiny and accountability.

Potential gains for justice systems

System-level gains

- Faster handling of routine administrative tasks and large document volumes.
- Improved search, consistency and retrieval across fragmented legal materials.
- Translation, transcription and accessibility tools can lower linguistic and practical barriers.
- Better allocation of scarce judicial resources through workload and process analytics.

Access-to-justice gains

- Lower transaction costs for simple claims and standard procedures.
- More understandable first-line legal information and guided digital pathways.
- Remote participation can reduce geographic and mobility barriers.
- Automation can free professionals to spend more time on complex reasoning and human interaction.

Efficiency is legally valuable only when it strengthens — rather than dilutes — effective access to justice and procedural fairness.

The constitutional risks of “algorithmic justice”

- Opacity can undermine the ability of parties to understand, contest and appeal a decision when the role of an algorithm is unclear or technically inaccessible.
- Historical data can reproduce or amplify structural inequalities; apparently neutral variables may operate as proxies for protected or socially sensitive characteristics.
- Automation bias can cause judges or officials to defer excessively to a quantified recommendation, even when formal human authority remains intact.
- Predictive tools may create feedback loops: past enforcement or adjudication patterns become training data that shape future enforcement or adjudication.
- Private vendors may acquire practical influence over public decision-making through proprietary models, data and technical dependencies that are difficult to audit.

Human oversight: meaningful, not ceremonial

1**Understand**

The professional should know what the system is designed to do, the source and limits of its data, and the types of error it can produce.

2**Interrogate**

Outputs should be treated as contestable inputs. The user must be able to inspect supporting materials, challenge assumptions and seek alternatives.

3**Override**

Human authority must include a practical capacity to reject, correct or depart from the system without organisational penalties for doing so.

4**Record**

Where AI materially affects a decision, organisations should preserve logs, versions, relevant inputs and the human review that occurred.

AI, due process and the right to a fair trial

Questions for courts

- Was AI used in a way that materially affected the handling or outcome of the case?
- Can affected persons know that an AI system was used and understand its role?
- Is there an effective opportunity to contest the data, reasoning or output?
- Can the court itself explain the legally relevant reasons for its decision independently of the model?

Questions for litigants

- Can counsel inspect or challenge AI-generated evidence or expert analysis?
- Does the opposing party have comparable technological capacity, or does technology create a new inequality of arms?
- Are disclosure and discovery rules adequate for algorithmic systems and model-generated records?
- What remedy exists if an automated process materially prejudices procedural rights?

DigitalJustice@2030: the EU policy direction at the end of 2025

- The European Commission adopted DigitalJustice@2030 on 20 November 2025, seeking faster digitalisation of national justice systems, greater resilience and more effective access to justice.
- The strategy envisages digital-by-default public services and highlights AI-supported functions such as speech-to-text, transcription and machine translation, particularly valuable in cross-border proceedings.
- The accompanying European Judicial Training Strategy 2025–2030 treats digital skills and AI awareness as core professional capacity, not a specialist add-on.
- The policy direction is therefore dual: accelerate digital uptake while building the skills, governance and safeguards necessary for trustworthy use.

02

The Legal Profession: Augmentation, Automation and Responsibility

How AI changes research, drafting, contracts, litigation, investigations, compliance and the economics of professional services.

A task-based view of legal AI

Tasks AI can accelerate

- First-pass legal and regulatory research.
- Summarisation of long documents and case files.
- Contract extraction, comparison and clause analysis.
- Drafting and redrafting correspondence, policies, pleadings and memos.
- Chronologies, issue lists, interview preparation and discovery review.

Tasks that remain professionally owned

- Defining the legal question and strategic objective.
- Determining which authorities and facts are legally material.
- Assessing credibility, ambiguity, proportionality and risk appetite.
- Giving advice tailored to the client's circumstances and values.
- Taking responsibility for representations made to clients, courts, regulators and counterparties.

Legal research with generative AI: powerful, but epistemically different

- Traditional legal databases retrieve identifiable sources; generative AI may instead produce a fluent synthesis whose individual propositions are not transparently linked to authoritative materials.
- Retrieval-augmented systems can improve grounding by connecting the model to curated legal sources, but retrieval errors, outdated databases and misinterpretation remain possible.
- The professional workflow should move from “ask and paste” to “ask → inspect sources → verify proposition → check currency → apply legal judgment”.
- The more specific the jurisdiction, procedural posture, date sensitivity or novelty of the issue, the less appropriate it is to rely on an unverified general-purpose answer.

Fluency is not authority. A plausible citation is not a verified citation; a correct citation is not necessarily controlling law.

Drafting: the efficiency gain is real — and so is the review burden

Good uses

- Generating a first structure or alternative formulations.
- Transforming style, tone, length or level of technicality.
- Producing checklists from a verified source document.
- Comparing versions and surfacing inconsistencies.
- Creating first-pass summaries for internal review.

Review obligations

- Verify every legal proposition and citation against authoritative sources.
- Check whether facts have been invented, conflated or subtly changed.
- Ensure the draft reflects the actual client instruction and jurisdiction.
- Remove generic or overconfident language that obscures uncertainty.
- Treat the final document as the professional's own work product.

Contracts and transactional work

- AI can extract clauses, compare agreements against a playbook, identify missing provisions, summarise obligations and generate first-pass redlines across large document sets.
- The principal advantage is scale: due diligence teams can move from manual reading of every document toward machine-assisted prioritisation and exception review.
- The principal risk is false assurance. A model may miss a commercially decisive clause, misunderstand cross-references, or treat unusual drafting as equivalent to a standard provision.
- For high-value transactions, quality assurance should include sampling, escalation thresholds, version control and a clear record of which outputs were machine-generated and how they were reviewed.

Litigation, investigations and e-discovery

Opportunities

- Technology-assisted review can prioritise large evidence sets and identify patterns that manual review may miss.
- AI can build chronologies, map entities, cluster documents and prepare deposition or interview themes.
- Transcription and translation accelerate multi-jurisdictional investigations.
- Generative tools can stress-test arguments by generating counterarguments and alternative narratives.

Risks

- Privilege and confidentiality may be compromised by inappropriate uploads or vendor data practices.
- Synthetic or manipulated evidence increases authentication and provenance challenges.
- Model summaries can omit exculpatory context or overstate weak correlations.
- Counsel remains responsible for discovery representations, evidentiary foundations and candour to the tribunal.

Compliance: from rule mapping to continuous monitoring

- Compliance functions can use AI to map regulatory obligations, compare policies with legal requirements, triage alerts, summarise regulatory updates and identify anomalies across large datasets.
- Generative AI can make compliance knowledge more accessible to business teams through conversational interfaces, but a user-friendly answer can conceal uncertainty or outdated legal content.
- Automated monitoring may process employee, customer or transaction data at scale, creating privacy, labour-law, discrimination and proportionality issues alongside the original compliance objective.
- The strongest operating model combines legal ownership, data governance, model validation, escalation routes and evidence that a human decision-maker reviewed material high-risk outputs.

The economics of legal services may change before the law does



Time per task

Routine research, extraction, summarisation and drafting can require substantially less professional time when workflows are well designed.



Value shifts

Client value moves away from producing text and toward framing, judgment, negotiation, assurance, domain knowledge and accountability.



Billing pressure

Time-based billing becomes harder to reconcile with tasks completed in minutes. Firms may face pressure toward fixed, value-based or technology-inclusive pricing.



Supervision

Efficiency gains do not eliminate review. Firms need new quality controls, training, vendor oversight and governance — costs that should be made visible rather than ignored.

Will AI replace lawyers? A more useful decomposition

- AI is more likely to replace or compress particular tasks than to eliminate the legal role as a whole. Legal work bundles research, factual investigation, persuasion, negotiation, empathy, ethical responsibility and institutional authority.
- Entry-level work is particularly exposed because junior lawyers traditionally learn through research, document review and first drafting — the same tasks most amenable to automation.
- This creates a training paradox: organisations can become more efficient while weakening the apprenticeship pathway through which future senior professionals acquire judgment.
- The strategic response is deliberate redesign of training: juniors should use AI, but must also practise source verification, first-principles analysis, client communication and independent drafting.

03

Professional Ethics in the Age of Generative AI

Competence, confidentiality, supervision, candour, communication, independence and fees remain human obligations.

The emerging professional-ethics consensus

- Using generative AI is not inherently unethical; the ethical question is whether the lawyer uses it consistently with existing professional duties.
- Competence increasingly includes a reasonable understanding of the benefits and risks of the technologies used in representation, while not requiring lawyers to become technical experts.
- Confidentiality requires attention to what data is entered, where it is processed, whether prompts or files are retained, and whether data may be reused for training or exposed to third parties.
- Independent verification is essential because uncritical reliance on AI output can produce inaccurate advice or misleading submissions.
- Supervisory duties extend to staff and technology-enabled workflows: a firm cannot outsource professional responsibility to a vendor or a model.

ABA Formal Opinion 512 (29 July 2024) synthesises duties of competence, confidentiality, communication, supervision, meritorious claims, candour and reasonable fees.

Competence now has a technological dimension

Minimum professional literacy

- Know whether the tool is general-purpose or legal-specific.
- Understand whether it searches sources, predicts text, or combines both.
- Know what information the provider stores and how it may be reused.
- Recognise hallucination, bias, prompt sensitivity and context-window limits.
- Know when the task exceeds the tool's safe use case.

Competent workflow

- Select a tool appropriate to the sensitivity and legal stakes.
- Use authoritative source material whenever possible.
- Verify outputs proportionately to risk.
- Document material review and corrections.
- Update knowledge as tools, terms and regulatory requirements change.

Confidentiality, privilege and professional secrecy

- Before entering client information into an AI service, determine whether the organisation has approved the tool and whether contractual, technical and organisational safeguards are adequate.
- Do not assume that a consumer-facing chatbot is equivalent to an enterprise environment with contractual restrictions on retention, training and access.
- Minimise data: use anonymised, pseudonymised or synthetic material where the task does not require identifiable information.
- Privilege can be affected by disclosure to third parties depending on jurisdiction and circumstances; legal teams should establish clear rules rather than rely on informal assumptions.
- Vendor due diligence should cover subprocessors, hosting location, security controls, incident response, deletion, model training practices and audit rights.

Candour to the court: the hallucinated-authority problem

- Generative AI can fabricate case names, quotations, docket numbers, statutes and factual propositions while presenting them in a professionally convincing form.
- The ethical failure is not “the model hallucinated”; it is the submission of unverified material by a professional who owed duties of competence and candour.
- Verification should occur at source level: open the authority, confirm the proposition, check quotation accuracy, procedural history, jurisdiction, precedential status and current validity.
- Court-specific rules on AI disclosure have developed unevenly. Practitioners should check the applicable judge, court and jurisdiction rather than assume a universal disclosure rule.

A citation checker is useful; a lawyer who understands why the authority matters is indispensable.

Client communication and transparency

When disclosure becomes more important

- AI use materially affects strategy, cost, confidentiality or the means by which the representation is carried out.
- Sensitive client information will be processed by an external system.
- The client has imposed technology restrictions or sector-specific requirements.
- AI is used in an unusual or high-impact manner rather than as ordinary back-office assistance.

What good communication looks like

- Explain the purpose and material risks in understandable language.
- Avoid vague boilerplate that purports to authorise any future AI use.
- Clarify the safeguards, human review and limits of the technology.
- Agree any relevant allocation of costs and treatment of third-party technology expenses.

ABA Formal Opinion 512 provides guidance on communication, confidentiality and fees; local professional rules remain controlling.

Fees, efficiency and the ethics of billing

- AI may reduce the time required to complete a task; professionals should not bill hypothetical hours that were not actually worked merely because the task previously took longer.
- Time spent reviewing and correcting AI output is professional work and may be chargeable when reasonable and consistent with the engagement.
- General learning costs for becoming familiar with a tool may be treated differently from matter-specific work; applicable professional rules and client agreements should be checked.
- Firms should decide whether AI subscriptions are overhead, a separately disclosed expense, or incorporated into alternative fee arrangements — and communicate the approach transparently.

ABA Formal Opinion 512 discusses reasonable fees and distinguishes actual time spent from inflated billing based on pre-AI effort.

A practical “red / amber / green” use policy

G

Green: routine support

Public information; formatting; brainstorming; non-sensitive summaries; translation of non-confidential text. Standard review remains necessary.

A

Amber: controlled use

Client-related drafting, contract analysis, research or internal investigations using approved enterprise tools, with source verification and documented human review.

R

Red: prohibit or escalate

Unapproved tools with confidential data; autonomous legal advice; fabricated evidence; unsupervised court filings; decisions materially affecting rights without approved safeguards.

A good policy is task-based and data-based. A blanket “AI allowed” or “AI prohibited” rule is usually too crude for professional practice.

04

The Regulatory Architecture Through 2025

The EU AI Act, GDPR, Council of Europe standards and sectoral rules create overlapping layers of governance.

EU AI Act: the risk-based architecture

1

Prohibited practices

Certain uses are considered incompatible with EU values and are prohibited, subject to the Regulation's definitions and exceptions.

2

High-risk systems

Specific use cases face extensive requirements on risk management, data governance, documentation, logging, transparency, human oversight, accuracy and cybersecurity.

3

Transparency duties

Certain systems and AI-generated content trigger transparency requirements so that users or affected persons are not misled about AI involvement.

4

GPAI layer

General-purpose AI models are regulated separately, with additional obligations for models presenting systemic risk.

Administration of justice as a high-risk domain

- Annex III classifies as high-risk AI systems intended to be used by or on behalf of a judicial authority to assist in researching and interpreting facts and law and in applying the law to a concrete set of facts.
- Comparable uses in alternative dispute resolution can also fall within the high-risk category where the outcome produces legal effects for the parties.
- The rationale is constitutional: these systems can significantly affect democracy, the rule of law, individual freedoms, the right to an effective remedy and the right to a fair trial.
- The Regulation expressly distinguishes purely ancillary administrative activities — such as certain anonymisation, communication and administrative tasks — from AI that affects the actual administration of justice in individual cases.

The AI Act's premise is clear: AI may support judicial decision-making, but final judicial decision-making must remain human-driven.

High-risk AI: what “compliance” means operationally

System controls

- Continuous risk-management process.
- Appropriate data and data-governance practices.
- Technical documentation and record-keeping/logging.
- Accuracy, robustness and cybersecurity requirements.
- Human-oversight measures designed into deployment.

Organisational controls

- Clear allocation of provider/deployer responsibilities.
- Instructions for use and user competence.
- Monitoring and incident/escalation processes.
- Fundamental-rights considerations where applicable.
- Procurement terms that give the deployer enough information and control to meet legal duties.

AI literacy became a legal governance issue in 2025

- The AI Act's AI-literacy obligation began applying on 2 February 2025, signalling that organisations cannot rely only on technical controls; people who operate or use AI need appropriate knowledge and awareness.
- For legal and compliance organisations, literacy should be role-specific: a litigator, compliance analyst, procurement lawyer and IT administrator face different risks and need different training.
- Training should cover capabilities, limitations, data handling, verification, bias, escalation, approved tools and the organisation's own use policy.
- A defensible programme records training, updates content as tools evolve, and links literacy to actual permissions and responsibilities in the AI workflow.

GDPR does not disappear when AI arrives

- Personal data used to train, configure, prompt or operate AI systems remains subject to GDPR principles including lawfulness, purpose limitation, data minimisation, accuracy, storage limitation, security and accountability.
- Legal teams should distinguish training data, retrieval data, prompt/input data, generated output, logs and telemetry: each may contain personal data and may have a different purpose, retention period and legal basis.
- Automated decision-making and profiling raise additional issues where decisions produce legal or similarly significant effects; the precise application of Article 22 depends on the role of automation and the circumstances.
- Data-protection impact assessment may be required where processing is likely to result in high risk; AI governance and privacy governance should therefore be integrated rather than operated as separate silos.

Council of Europe: human rights, democracy and the rule of law

- The Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law opened for signature on 5 September 2024 and is the first legally binding international treaty specifically addressing AI in this field.
- Its approach is lifecycle-based and rights-centred: AI activities should be consistent with human dignity, autonomy, transparency and oversight, accountability, equality, privacy, reliability and safe innovation.
- The Convention requires risk and impact management and supports procedural safeguards and remedies in relation to AI-related human-rights impacts, within the scope of Parties' obligations.
- For justice professionals, its significance is conceptual as well as legal: AI governance is framed not only as product safety or compliance, but as protection of institutional democracy and the rule of law.

CEPEJ's five enduring principles for AI in justice

01

Fundamental rights

Design and implementation should be compatible with fundamental rights.

02

Non-discrimination

Prevent the development or intensification of discrimination between individuals or groups.

03

Quality & security

Use reliable sources, multidisciplinary design and secure technological environments.

04

Transparency & fairness

Make processing methods accessible and understandable and permit appropriate external scrutiny.

05

User control

Avoid prescriptive automation and ensure users remain informed actors in control of choices.

End-2025 development: CEPEJ guidance on generative AI for courts

- By December 2025, CEPEJ had added Guidelines on the use of generative AI for courts to its cyberjustice toolkit, following its earlier 2024 information note on generative AI for judicial professionals.
- This reflects a shift from debating whether generative AI belongs in justice toward defining conditions for its controlled professional use.
- The direction of travel is consistent across European instruments: preserve human responsibility, assess risks in context, protect confidential and personal data, verify outputs, and maintain transparency and contestability.
- For organisations, the practical lesson is to align internal policies with the highest-impact use case rather than assume that a generic corporate AI policy is sufficient for legal or judicial work.

05

Evidence, Deepfakes and the Integrity of Legal Processes

AI changes not only how professionals work, but also the evidence, communications and identities they must evaluate.

Synthetic media changes the evidentiary baseline

- Generative AI lowers the cost of producing realistic synthetic text, voice, images and video, increasing the possibility that fabricated material enters investigations, negotiations or litigation.
- At the same time, the existence of deepfakes creates the “liar’s dividend”: authentic evidence may be dismissed as synthetic simply because manipulation is technically possible.
- Legal teams increasingly need provenance, chain-of-custody and metadata practices that preserve how digital evidence was acquired, transformed, stored and analysed.
- Authentication becomes a socio-technical exercise: legal rules of evidence interact with forensic methods, platform records, cryptographic provenance, witness testimony and contextual corroboration.

AI-generated evidence: questions before reliance

Authenticity & provenance

- What is the original source and acquisition method?
- Is the original file available, with metadata preserved?
- Has the content been recompressed, edited, translated or enhanced?
- Can provenance or platform records corroborate origin?
- Is expert forensic analysis necessary?

Reliability & admissibility

- What system generated or analysed the material?
- Can the method be explained and tested?
- What error rates or known limitations are relevant?
- Was human review independent and documented?
- Can the opposing party meaningfully challenge the evidence?

Cybersecurity and AI are now inseparable legal risks

- AI systems expand the attack surface through prompts, plugins, connected data sources, model supply chains and automated actions. Legal organisations hold unusually sensitive information and are attractive targets.
- Prompt injection can manipulate an AI system into ignoring instructions or exposing information retrieved from connected sources; this is particularly relevant to agentic and retrieval-based systems.
- Data poisoning, compromised models and insecure third-party integrations can affect integrity as well as confidentiality.
- Legal and compliance teams should require security review, access controls, least privilege, logging, incident response and vendor assurance proportionate to the sensitivity of the workflow.

06

Building a Responsible AI Programme for Legal Work

From experimentation to inventory, risk classification, controls, procurement, training, monitoring and audit.

The legal AI governance lifecycle

1

Inventory

Identify systems, models, vendors, integrations, data sources, users and business purposes — including “shadow AI”.

2

Classify

Assess legal, ethical, data, security and professional risks according to the actual use case and affected persons.

3

Control

Define approved tools, access, data restrictions, human review, verification, disclosure and escalation requirements.

4

Monitor

Track incidents, output quality, model changes, vendor changes, regulatory developments and user behaviour over time.

A practical legal-AI risk assessment

Questions about the use case

- Does the output affect rights, legal position, employment, liberty, financial access or regulatory treatment?
- Is the tool advisory, or is it effectively determinative because humans rarely depart from it?
- Will confidential, privileged, personal or special-category data be processed?
- Is the output externally communicated to a client, court, regulator or counterparty?

Questions about the system

- What model and version are used, and can they change without notice?
- What data sources ground the output and how current are they?
- What logs and audit evidence are available?
- Can users understand uncertainty and limitations?
- What happens if the system fails, becomes unavailable or produces harmful output?

Procurement: clauses legal teams should ask for

- Purpose and data-use restrictions: specify whether customer data, prompts and outputs may be used for training, evaluation or product improvement.
- Security and confidentiality: require appropriate technical and organisational measures, subprocessor transparency, incident notification and deletion/return obligations.
- Change management: obtain notice of material model, hosting, feature or terms changes that could alter risk or compliance.
- Auditability and evidence: secure access to logs, documentation, performance information and cooperation needed for regulatory or internal reviews.
- Liability and exit: address indemnities, service continuity, portability, data deletion and the practical ability to migrate away from the provider.

Human review should be calibrated to consequence

L

Low consequence

Formatting, brainstorming or non-sensitive summaries: ordinary professional review may be sufficient.

M

Moderate consequence

Client drafting, contract analysis or internal compliance advice: source verification, approved tools and named reviewer.

H

High consequence

Court submissions, regulatory filings, rights-affecting decisions or sensitive investigations: independent verification, escalation, documented sign-off and auditable evidence.

“Human in the loop” is not a control unless the human has time, competence, authority and information to challenge the machine.

Scenario workshop: a lawyer uses AI to prepare a court filing

- A lawyer uploads a client chronology and several pleadings to a public generative-AI service, asks for relevant authorities, and uses the resulting draft in a filing after a quick read-through.
- Issue 1 — confidentiality: was the tool approved, and what happens to uploaded client information?
- Issue 2 — competence and candour: were every authority, quotation and factual statement verified at source?
- Issue 3 — supervision: does the firm have a policy, training and review process that makes safe use realistically possible?
- Issue 4 — transparency: do applicable court rules, client instructions or engagement terms require disclosure or consent?
- Issue 5 — governance: could the same efficiency have been achieved through an enterprise tool connected to verified legal sources with stronger data controls?

Scenario workshop: AI-assisted compliance investigation

- A compliance team uses an AI system to prioritise employees for investigation based on communications, expense data and transaction patterns. The model produces a risk score and short explanation.
- Ask whether the system is merely organising information or is materially determining who becomes subject to intrusive investigation.
- Examine lawful basis, necessity, proportionality, employee privacy, labour-law implications, discrimination risk and the reliability of proxy variables.
- Require a review process that can identify false positives and contextual explanations rather than treating a score as evidence of wrongdoing.
- Document the purpose, data sources, model limitations, escalation thresholds and human reasoning that converts an alert into an investigative decision.

Scenario workshop: a court deploys AI decision support

- A court deploys a system that analyses prior decisions and recommends a likely outcome and range of remedies to judges handling high-volume cases.
- The efficiency case is obvious, but so are the risks: historical bias, anchoring, opacity, feedback loops and gradual convergence around past practice.
- The key governance questions are whether the use falls within the AI Act's high-risk justice category, how the system is validated, and whether parties can know and challenge its material role.
- Judges need meaningful authority to depart from recommendations, and the legal reasons in the judgment must remain the court's own reasons rather than a post-hoc rationalisation of an algorithmic output.

A 10-point checklist before using AI on a legal matter

- 1. Define the task precisely and decide whether AI is appropriate for it.
- 2. Use an organisation-approved tool suitable for the sensitivity of the matter.
- 3. Minimise confidential and personal data before it enters the system.
- 4. Understand whether the output is generated, retrieved, predicted or some combination.
- 5. Require sources where the task depends on external factual or legal authority.
- 6. Verify legal propositions, citations, quotations and material facts independently.
- 7. Look actively for bias, omissions, overconfidence and alternative explanations.
- 8. Preserve appropriate records when AI materially affects professional work or decisions.
- 9. Escalate high-impact, novel or rights-affecting uses to legal/AI governance specialists.
- 10. Take personal responsibility for the final advice, filing, decision or communication.

What professionals should learn next

Core professional skills

- AI literacy: capabilities, limitations and terminology.
- Source verification and information provenance.
- Data protection, confidentiality and security by design.
- Risk assessment and documentation.
- Vendor and contract governance.
- Clear communication of uncertainty and limitations.

Higher-value human skills

- Problem framing and legal judgment.
- Cross-domain reasoning and contextual interpretation.
- Negotiation, advocacy and persuasion.
- Ethical judgment and accountability.
- Client empathy and trust.
- Ability to challenge automated recommendations rather than defer to them.

Key conclusions

- AI is becoming part of the infrastructure of legal work and justice systems, but adoption is uneven and fully autonomous adjudication is not the operative model.
- The greatest near-term transformation is task-level: research, review, drafting, translation, discovery, compliance triage and knowledge management can be accelerated substantially.
- The legal risks are not confined to hallucinations. They include confidentiality, bias, opacity, automation bias, evidentiary manipulation, cybersecurity, professional independence and unequal access to technology.
- European governance increasingly treats AI in justice as a rule-of-law issue. The EU AI Act, Council of Europe Convention and CEPEJ framework all reinforce human oversight, accountability and rights protection.
- The durable professional advantage is not competing with AI at producing text; it is combining technological fluency with verified sources, legal judgment, ethical responsibility and human trust.

Use AI to extend professional capacity — never to outsource professional responsibility.

Selected primary and professional sources

- Regulation (EU) 2024/1689 (Artificial Intelligence Act), especially recital 61 and Annex III(8) on administration of justice; European Commission AI Act implementation materials.
- European Commission, DigitalJustice@2030, COM(2025) 802 final, 20 November 2025; European Judicial Training Strategy 2025–2030, COM(2025) 801 final.
- Council of Europe, Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law (CETS No. 225), opened for signature 5 September 2024.
- CEPEJ, European Ethical Charter on the Use of Artificial Intelligence in Judicial Systems and their Environment (2018); Assessment Tool (2023/updated 2025); 1st Report on AI in the Judiciary (2025); Guidelines on the Use of Generative AI for Courts (December 2025).
- American Bar Association, Formal Opinion 512, Generative Artificial Intelligence Tools, 29 July 2024.
- GDPR: Regulation (EU) 2016/679, including principles of processing, DPIAs and rules relevant to automated decision-making and profiling.

Discussion

01

Where would you use it?

Which legal or compliance tasks in your organisation are sufficiently repetitive, document-heavy and reviewable to benefit from AI now?

02

Where would you hesitate?

Which tasks involve rights, sensitive data, evidentiary judgments or strategic decisions where automation creates disproportionate risk?

03

What must change?

Do your policies, engagement terms, procurement processes, training and quality controls reflect actual AI use — including shadow AI?

04

What remains human?

Which aspects of professional value should your organisation deliberately protect and strengthen as routine cognitive work becomes cheaper?