



PROGRAM MATERIALS

Program #36218

October 22, 2026

Artificial Intelligence and the Practice of Law

Copyright ©2026 by

- **Joseph J Bambara, Esq. CIPP/US - Withers
Worldwide**

**All Rights Reserved.
Licensed to Celesq®, Inc.**

Celesq® AttorneysEd Center
www.celesq.com

**5301 North Federal Highway, Suite 150, Boca Raton, FL 33487
Phone 561-241-1919**

CONTINUING LEGAL EDUCATION · TECHNOLOGY & THE LAW

Artificial Intelligence and the Practice of Law

A Competence Primer for the Practicing Lawyer

From LLMs to liability, IP, contracts, and global regulation **what Model Rule 1.1 now expects you to know.**

Presented by: Joseph J Bambara WithersWorldwide.com Date: October 2026

Educational materials only not legal advice. Case law and statutes are current through September 2026 and evolving rapidly.

Joseph Bambara, CIPP/US

Attorney Most recently, Joseph is Of Counsel with withersworldwide.com developing their AI and Blockchain/Digital Assets practice. Specifically providing guidance on regulatory and governance issues regarding use of AI and digital assets for large enterprises, decentralized startups and crypto clients. He handles privacy policy and risk assessment, as well as mobile and email and website practices and procedures to self-regulate and maintain compliance with applicable legislation including: AI and digital asset regulation and HIPAA, CAN/SPAM, TCPA, Gramm-Leach-Bliley ECPA, COPPA and all related US privacy legislation CCPA as well as the EU AI Act , European Union's Data Protection Directive and GDPR.. He has handled all interactions, inquiries with FTC. He is experienced with eDiscovery and legal tech software and procedures.

Entrepreneur/Technologist with architectural and software development experience with cyber security, banking, financial, manufacturing, health care and legal applications for Goldman Sachs, Bank of America Merrill Lynch, Federal Reserve Bank, Viacom, Cablevision, Pfizer, Merck and Sloan Kettering Cancer Care.

Educator/Speaker who has produced and delivered over 50 CLE courses and corporate presentations covering the intersection of technology and the law including: artificial intelligence, blockchain, cybercrime, cybersecurity, eDiscovery, Privacy, mobile devices, social media and the internet of things.

Author of ten plus books for McGraw Hill and other publishers on intersection of law and technology: mobile, internet, blockchain development and related technologies. Most recently, in 2019, he authored "*Law and Technology: The Future is Now*" and "*AI, IoT and the Blockchain Using the Power of Three to create Business, Legal and Technical Solutions*". In 2018 he authored "*Blockchain: A Practical Guide to developing Business, Law and Technology Solutions*".

Admissions NY State 1981, US Supreme Court 1993

Courses/Lectures

- Lectured as adjunct faculty at CCNY School of Engineering , Computer Science Department Instructor for over 50 CLE's (not all courses are listed)
- What You Need to Know About DeFi: Decentralized Finance and the Law NYCLA, September 2020
- AI, IoT, Blockchain: Business/Legal Ramifications NYCLA, April 2020
- Law and Technology: The Future is Now NYCLA, October 2019
- What Lawyers Need to Know About Cryptocurrency National Constitution Center, April 2019
- Ethics and Technology: What Lawyers MUST Know to Protect Client Data NYCLA, December 2018
- The Law and Cybersecurity, Cybercrime and Cryptocurrencies Lawline, March 2017
- eDiscovery: Trends and TAR Lawline, April 2016.

Education

CCNY School of Engineering: BS Computer Science and MS Computer Science
Brooklyn Law School Juris Doctorate

Joseph Bambara CIPP/US

Digital Assets/AI/Media/Privacy

WithersWorldwide

Of Counsel

joseph.bambara@withersworldwide.com

UCNY, Inc.

New York, New York

jbambara@ucny.com

Web: <http://www.ucny.com>

Twitter: <http://www.twitter.com/jbesquire7>

Amazon.com/author/josephbambara

Google Scholar

<https://tinyurl.com/y6yzt9qg>


**Chair New York County Law and
Technology Committee**

AI, IoT and the Blockchain: Using the Power of Three to create Business, Legal and Technical Solutions

Joseph J Bambara and Paul R Allen

<https://www.amazon.com/AI-IoT-Blockchain-Technical-Solutions-ebook/dp/B081TMHN5D>

This comprehensive guide teaches organizations and their developers to prepare to incorporate the three main emerging technology trends: blockchain, Internet of Things (IoT), and Artificial Intelligence (AI). While each technology is powerful on its own, their functionality is transformational when properly combined.



Six Modules, One Objective: Competent AI Counsel

1

Core AI Concepts

LLMs, generative & agentic AI the vocabulary
ABA Rule 1.1 competence requires.

2

Theories of Liability

Negligence, products liability & vicarious
liability across the value chain.

3

Intellectual Property

Training data, output ownership, and
inventorship/authorship exposure.

4

Contract Drafting

Provenance, IP indemnity, SLAs, and allocating
hallucination risk.

5

Global Regulation

EU AI Act vs. UK, China, and the U.S. state-law
patchwork.

6

The Horizon

Pending litigation, new statutes, and treaty
obligations through 2026+.



A Hypothetical to Anchor the Analysis


MERIDIAN HEALTH SYSTEMS

Your client is a U.S.-headquartered health company with operations in the EU and China. It is deploying **“AskMeridian,”** a customer-facing generative assistant built on a third-party foundation model, augmented with agentic tools that can schedule appointments and draft correspondence.

Recurring questions we will answer:


- Who bears the loss when the assistant hallucinates?
- Is Meridian exposed for the vendor's training data?
- Which regimes govern the same deployment at once?
- What must the contract with the vendor say?

THREE PARTIES, THREE RISK POSTURES




Developer

The foundation-model vendor that trained the model.



Deployer

Meridian integrates and puts the system into service.



User

Patients and staff who rely on the outputs.

MODULE 1

Core AI Concepts & Architectures

Building the technical fluency ABA Rule 1.1 competence now demands.

1

Rule 1.1: Competence Now Includes Technology

MODEL RULE 1.1, COMMENT [8]

“To maintain the requisite knowledge and skill, a lawyer should keep abreast of changes in the law and its practice, including the benefits and risks associated with relevant technology ...”

Adopted in some form by 40+ U.S. jurisdictions.

§

Duty to understand

Enough grasp of how AI systems work to evaluate their benefits and risks for a matter.

§

Duty of candor & verification

Sanctions have followed fabricated, AI-'hallucinated' citations Rule 3.3 and 11 exposure is real.

§

Duty of confidentiality

Rule 1.6 client data entered into third-party models may forfeit privilege or confidentiality.


§

Duty to supervise

Rules 5.1 & 5.3 extend to non-lawyer 'assistants,' including AI tools used by the firm.

Sources: ABA Model Rule 1.1 cmt. [8]; ABA Formal Op. 512 (2024) on generative AI; Rules 1.6, 3.3, 5.1, 5.3.

The Taxonomy: AI $\supset$ Machine Learning $\supset$ Generative AI



Why the nesting matters

Not all 'AI' is generative. A fraud-scoring model and ChatGPT raise very different legal issues precision matters in pleadings and contracts.

"The model" is a set of weights. Numbers learned during training; running it (inference) is a separate act from training it.

Foundation vs. fine-tuned. A general 'foundation' model is often adapted downstream changing who controls, and is liable for, behavior.

How a Large Language Model Actually Works



THE LEGAL PUNCHLINE

An LLM is a probabilistic next-word predictor, not a database of facts. It has no built-in notion of truth, no citation index, and no memory of individual training documents which is exactly why it can 'hallucinate' fluent, confident, and false output.

Generative AI and the Hallucination Problem

What 'generative' spans

- **Text:** Drafting, summarizing, Q&A, code.
- **Image / audio / video:** Synthetic media and deepfakes.
- **Structured output:** Extraction, classification, tool calls.
- **RAG:** Retrieval-augmented generation grounds output in supplied documents mitigates but does not eliminate error.

Why it hallucinates and why it matters



Confident fabrication Invented cases, citations, quotes, and figures presented as fact.



Plausible but wrong Fluency masks error; users over-trust polished output ('automation bias').




Context sensitivity Small prompt changes flip the answer hard to reproduce or audit.



Agentic amplification When the model acts (books, sends, transacts), an error becomes a real-world harm.

The Training Pipeline and Who Controls Each Stage



Why it matters: liability tracks the stage. The vendor owns pre-training and its data provenance; the deployer owns fine-tuning data, guardrail choices, and deployment context. Diligence and contract terms should follow the same seams.

Four Ways to Customize & Ground a Model

1

Prompting / system prompts

Instructions steer behavior at run time no retraining.

Legal footprint: Cheapest lever; but prompt-injection is an attack surface.

2

Context & RAG

Retrieval-augmented generation feeds source docs into the window.

Legal footprint: Grounds answers in YOUR data; reduces (not eliminates) hallucination.

3

Fine-tuning

Further training on curated examples bakes in behavior.

Legal footprint: Your data shapes the model provenance & confidentiality risk.

4

Embeddings / vector DB

Text encoded as vectors for semantic search over a knowledge base.

Legal footprint: Where confidential corpora live secure and access-control them.

Distinguish these in diligence and drafting: each puts different data at risk and shifts control between vendor and client.

Agentic AI: When the Model Takes Action

From answering to acting

An agentic system pairs an LLM with **tools, memory, and autonomy** so it can plan and execute multi-step tasks calling APIs, querying databases, sending communications, and transacting with limited human oversight.

Loop: perceive → plan → act → observe → repeat

For Meridian: AskMeridian doesn't just suggest an appointment it books one, sends the email, and updates the record.

A new legal risk surface



Attribution Harm results from a chain of autonomous decisions whose act was it?



Authority & agency Can an agent bind the principal to a contract it 'negotiated'? Actual vs. apparent authority.



Foreseeability Emergent, multi-step behavior strains the negligence foreseeability inquiry.



Security Prompt injection and tool misuse turn the agent into an attack vector.

The Competent Lawyer's AI Vocabulary



Distinguish training vs. inference; foundation vs. fine-tuned; discriminative vs. generative.



Explain that an LLM predicts tokens probabilistically and does not store or retrieve facts.



Anticipate hallucination, automation bias, and reproducibility as inherent, not incidental.



Interrogate an agent's tools, autonomy, guardrails, and human-in-the-loop checkpoints.



Know your limits when to retain a technical expert and how to frame the questions.



Protect client confidences: understand a tool's data-retention and training-on-input terms.


MODULE 2

Theories of Liability for AI Harm

Negligence, products liability, and vicarious liability and who ends up holding the risk.

2

Three Theories, One Value Chain



Negligence: Fault in a Black Box

1

Duty

Does the developer/deployer owe a duty to the end user or third parties?
Special relationships (medical, financial) raise the bar.

2

Breach / Standard of Care

What would a reasonable developer or deployer do testing, guardrails, disclosures, human oversight? An emerging 'reasonable AI' standard.

3

Causation

The 'black box' problem: proving a specific model behavior caused the harm, and that it was foreseeable, is evidentiarily hard.

4

Damages

Economic loss doctrine may bar pure-economic-loss claims absent physical injury or property damage.

Watch: professional-negligence framing where AI supplements licensed judgment (medicine, law, finance); disclaimers rarely defeat a duty grounded in a special relationship.

The 'Reasonable AI' Standard Is Being Written Now

1 Voluntary frameworks

NIST AI Risk Management Framework (AI RMF 1.0) and ISO/IEC 42001 are becoming the reference points courts and regulators cite for 'reasonable' AI governance.

2 Sectoral standards

FDA (medical devices), FINRA/SEC (finance), and state medical boards supply domain-specific care benchmarks that raise the bar.

3 Proving the breach

Expert testimony (subject to Daubert) will define the standard; plaintiffs may invoke *res ipsa loquitur* where the harm 'speaks for itself.'

4 Building the defense

Documented testing, red-teaming, bias audits, human-in-the-loop, and disclosures are the record that shows reasonable care.

NIST AI RMF 1.0 (2023) & Generative AI Profile (2024); ISO/IEC 42001:2023; sector regulators. No AI-specific tort standard is settled these fill the gap.

Products Liability: Is Software a 'Product'?

THE THRESHOLD FIGHT

Is an AI model a 'product' or a 'service'?

Traditionally, courts hesitate to treat intangible software and information as 'products' under strict liability. **But that line is eroding** the 3rd Restatement's tangible-goods focus is under pressure, and the EU's revised Product Liability Directive expressly brings software and AI within scope.

Three classic defect theories, applied to AI



Manufacturing defect

Hard to frame every inference is 'as designed'; there is no deviation from spec.



Design defect

The core battleground: risk-utility and whether a reasonable alternative design (better guardrails/training) existed.



Failure to warn

Inadequate disclosure of known limitations, error rates, and appropriate-use boundaries.

Compare: EU Product Liability Directive (revised 2024) expressly covers software & AI and eases the burden of proof; U.S. law remains state-by-state and unsettled.

Vicarious Liability & the Agency Question

Respondeat superior

Employers answer for employees' torts in the scope of employment. Does deploying an AI 'agent' invoke the doctrine or is the AI merely a tool?

Non-delegable duties

Some duties (safety, care of patients/clients) cannot be outsourced to a vendor or an algorithm the deployer stays on the hook.

Apparent authority

If AskMeridian appears to speak for Meridian, patients' reasonable reliance can bind Meridian to its statements and 'agreements.'

Contractual risk-shifting

Indemnities and 'independent contractor' framing reallocate but do not always defeat third-party vicarious claims.

Beyond the Big Three: Other Theories in Play

Breach of warranty

UCC express & implied warranties (merchantability, fitness) and marketing claims that become express warranties.

Misrepresentation / fraud

Overstated accuracy or capability claims can support fraud and deceptive-practices theories.

Privacy & biometric

BIPA, CCPA/CPRA, wiretap and VPPA class actions over data used to train or run the system.

Discrimination

Title VII, ECOA, FHA disparate-impact exposure when AI drives hiring, credit, or housing decisions.

Defamation

Hallucinated false statements about real people e.g., *Walters v. OpenAI* test libel-by-AI.

§ 230 limited shield

Immunity likely does not cover content the AI itself generates, unlike third-party content it merely hosts.

Walters v. OpenAI (Ga. 2024–25) first notable AI-defamation suit. § 230 application to generative output remains unsettled.

Who Holds the Risk? An Allocation Matrix

Theory	Developer	Deployer	User
Negligence	Design/testing conduct	Deployment, oversight & warnings	Misuse; ignoring warnings
Products liability	Highest the 'manufacturer'	Seller/distributor exposure	Generally protected (bystander)
Vicarious / agency	Limited (no relationship)	Highest the principal	Low
Contract / indemnity	Caps, disclaimers, IP indemnity	Negotiates protection both ways	Terms of use, acceptable-use limits

Illustrative, not jurisdiction-specific. 'Highest' flags the party most naturally targeted under each theory; contract terms redraw these lines.

MODULE 3

Intellectual Property Exposure

Training-data inputs, generated outputs, and the authorship/inventorship questions no statute cleanly answers.

3

Three Zones of AI Intellectual-Property Risk

INPUTS

Training data

Was ingesting copyrighted works to train the model infringement or fair use? Exposure runs to developers and, by indemnity, to deployers.

OUTPUTS

Generated content

Can AI output be owned? Does it infringe? Substantial-similarity and 'regurgitation' risk.

OWNERSHIP

Authorship & inventorship

Human-authorship and human-inventorship requirements leave gaps in copyright and patent protection.

Training Data & the Fair-Use Battleground

The four fair-use factors, AI-style

1 • Purpose & character

Is training 'transformative'? Courts in 2025 largely said yes but commercial use cuts against.

2 • Nature of the work

Creative works get more protection than factual ones.

3 • Amount used

Whole-work copying is common in training but courts weigh it against transformativeness.

4 • Market effect

The decisive factor: lost licensing markets and 'market dilution' from AI-generated substitutes.

U.S. COPYRIGHT OFFICE PART 3 (May 2025)

“Generative AI Training”

Training can be transformative but **commercial use that produces content competing in the same market, especially via pirated access, exceeds fair use.**

- Factors 1 and 4 carry the most weight; the analysis is case-by-case.

Endorses **market-based licensing** over compulsory licenses.

Flags “**market dilution**” as a cognizable harm.

Four Rulings Every AI Lawyer Should Know

Case / Court	Holding	Why it matters
Thomson Reuters v. Ross <i>D. Del. Feb. 2025</i>	Training on Westlaw headnotes NOT fair use (non-generative AI).	First major fair-use loss for an AI defendant; now on interlocutory appeal to the 3d Cir.
Bartz v. Anthropic <i>N.D. Cal. June 2025</i>	Training on lawfully bought books = fair use; retaining pirated copies = infringement.	Split holding; led to a ~\$1.5B settlement (~\$3,000/work, ~500k works).
Kadrey v. Meta <i>N.D. Cal. June 2025</i>	Fair use but only because plaintiffs failed to prove market harm.	Chhabria signals 'market dilution' could win on a better record.
NYT v. OpenAI/Microsoft <i>S.D.N.Y. pending</i>	Infringement claims survived dismissal; SJ briefed 2026.	Landmark discovery: order to produce ~20M ChatGPT logs (aff'd Jan. 2026).


Holdings summarized as of Sept. 2026; several are on appeal or pre-judgment. Verify current posture before relying.

Output-Side Infringement & Adjacent Claims


HOW OUTPUTS INFRINGE

- **Memorization / regurgitation** Models can reproduce training text or images near-verbatim (the NYT exhibits) direct-infringement risk.
- **Substantial similarity** Even non-verbatim output may infringe if it copies protected expression, not just unprotected style or ideas.
- **Style vs. expression** Imitating an artist's 'style' is generally not infringement but courts probe where style ends and copying begins.

Adjacent claims & mitigations




DMCA § 1202 (CMI) Removing/altering copyright-management info SDNY courts are split on its reach in AI cases.



Trademark & publicity Outputs may infringe marks or misappropriate a person's name, voice, or likeness.



Vendor output indemnity Leading vendors now indemnify paid users for output IP claims read the conditions and caps.



Technical guardrails Output filters, dedup, and 'copyright shields' reduce regurgitation risk diligence them.

DMCA § 1202 CMI claims split in S.D.N.Y. AI cases; Getty v. Stability (trademark + copyright). Output indemnities vary widely verify scope.

Who Owns AI Output? The Human-Authorship Rule

Copyright Office & courts: no human, no copyright



Thaler (Creativity Machine) D.D.C. / D.C. Cir. affirmed 2025 a work authored solely by AI is not copyrightable.



“Zarya of the Dawn” Office cancelled registration of AI-generated images; kept human-authored arrangement & text.



USCO Part 2 (Jan. 2025) Copyright protects human contribution; prompts alone generally don't confer authorship.

Practical consequences


Protectability gap Pure AI output may be free for anyone to copy a business-model risk for content clients.

Human-in-the-loop Selection, arrangement, and meaningful edits can create protectable authorship document it.

Infringement risk Outputs can still infringe inputs (regurgitation / substantial similarity) even if unprotectable.

Contract fills the gap Where copyright fails, allocate rights and use restrictions by contract instead.

Patents: A Natural Person Must Invent



Counsel's task: identify and document the human conception behind an AI-assisted invention prompt design, problem framing, and output selection and name only natural persons as inventors.

Thaler v. Vidal, 43 F.4th 1207 (Fed. Cir. 2022); *USPTO Inventorship Guidance* (Feb. 2024); *Revised Guidance*, 90 Fed. Reg. (Nov. 28, 2025).

Protecting AI-Assisted Work in Practice

1

Copyright registration

Disclose AI-generated portions and disclaim them; register the human-authored selection, arrangement & edits. Misstatements risk cancellation.

2

Patent prosecution

Name only natural persons; document the human conception (problem framing, prompts, output selection); honor the duty of candor.

3

Trade secret

Where copyright/patent fail, secrecy can protect models, prompts & outputs but entering data into third-party tools can destroy it.

4

Contract & licensing

Allocate whatever rights exist by agreement; use license grants, assignments, and use restrictions to fill statutory gaps.

USCO registration guidance (2023) requires disclosure of AI-generated content; USPTO duty of disclosure applies to AI-assisted inventions.

MODULE 4

Drafting & Negotiating AI Contracts

Turning the risk map into enforceable terms: provenance, indemnity, SLAs, and liability allocation.

4

The AI Contract: Eight Provisions to Get Right

1

Data provenance & rights Reps that training and input data were lawfully sourced and licensed.

2

IP indemnification Vendor defends third-party IP claims arising from the model and its outputs.

3

Service levels (SLAs) Uptime plus accuracy, latency, and error-rate commitments with remedies.

4

Liability & limitations Caps, carve-outs, and who bears loss from erroneous output.

5

Confidentiality & training use Bar the vendor from training on your data; define retention.

6

Security & compliance Standards (SOC 2, ISO 42001), incident notice, regulatory cooperation.

7

Acceptable use & human oversight Allocate misuse risk; require human-in-the-loop where law demands.

8

Audit, exit & transition Model/version change control, explainability support, wind-down rights.

Data Provenance & IP Indemnification

DATA PROVENANCE

Ask for representations that:

- training data was lawfully obtained and licensed;
- no known infringing or improperly scraped content;
- compliance with privacy law for any personal data;
- disclosure sufficient to meet your own transparency duties (e.g., CA AB 2013);
- audit or documentation rights to verify.

Vendor pushback: “as-is,” no training-data warranties. Trade for indemnity + higher fees.

IP INDEMNIFICATION

Negotiate on five axes:

Scope: Inputs, model, AND outputs not just the software.

Trigger: Third-party IP claims, defense + indemnity duties.

Carve-outs: Watch exclusions for customer prompts / fine-tuning / modifications.

Caps: Resist a cap that swallows the indemnity; seek super-cap or uncapped IP.

Conditions: Prompt notice, control of defense, duty to mitigate.

Caps, Carve-Outs & Insurance

A TYPICAL CAP STRUCTURE

- 1 General cap** Fees paid over trailing 12 months × a multiple (1×–3×).
- 2 Super-cap** Higher, separate cap for data breach & confidentiality (e.g., 3×–5× or a fixed floor).
- 3 Uncapped bucket** IP indemnity, gross negligence, willful misconduct, death/bodily injury resist any cap.

Watch mutual vs. one-way caps, and whether the cap swallows the indemnities.

Insurance & indemnity procedure



Coverage to require Tech E&O, cyber, and emerging affirmative 'AI' endorsements confirm AI isn't excluded.



Additional insured Name the client; require certificates and notice of cancellation.



Defense & control Who controls defense and settlement; consent-to-settle; duty to cooperate.



Notice & mitigation Prompt-notice conditions and the client's duty to mitigate losses.

SLAs & Allocating the Hallucination Risk

Move beyond uptime

- **Accuracy / error rate** Define, measure, and remedy benchmarked on agreed test sets.
- **Groundedness** For RAG: outputs traceable to source documents.
- **Latency & availability** Classic SLA metrics, with service credits.
- **Change control** Notice + testing window before model/version swaps.

SAMPLE LIABILITY FOR ERRONEOUS OUTPUT

“Provider acknowledges the Service may generate inaccurate output (‘Hallucinations’). Provider shall [maintain error rates below X% on the Benchmark] and [indemnify Customer for Losses arising from Hallucinations relied upon in accordance with the Documentation]. Customer shall maintain human review of [high-risk] outputs. Nothing herein limits Provider’s IP indemnity or liability for [gross negligence / willful misconduct / breach of confidentiality].”

Bracketed terms are negotiation levers, not settled market terms.

Data, Privacy & Confidentiality Terms

1

No training on your data

Bar the vendor from using your inputs/outputs to train or improve models absent explicit opt-in.

2

Data processing agreement

GDPR Art. 28 / state-law DPA: roles, purposes, sub-processors, security, transfer mechanisms.

3

Data residency

Where data is stored and processed; localization needs for EU, China, and regulated sectors.

4

Output ownership & license

Who owns generated output; the license grant back to the vendor; use restrictions.

5

Retention & deletion

Retention limits, deletion on termination, and proof of deletion.

6

Privilege & work product

Protect confidential/privileged material entered into the tool; access & disclosure limits.

Align these with your regulatory duties: EU/UK GDPR, U.S. state privacy laws, sectoral rules (HIPAA, GLBA), and professional confidentiality obligations.


MODULE 5

Comparative Global Regulation

One deployment, many regimes: the EU AI Act against the UK, China, and the U.S. state patchwork.

5

The EU AI Act: Risk Tiers, Dates & Fines



Phased application

- Aug 2024**: Act enters into force
- Feb 2025**: Prohibited-practice bans apply
- Aug 2025**: GPAI (general-purpose model) obligations
- Aug 2026**: Transparency rules & AI Office enforcement
- 2027–28**: High-risk obligations phase in

Fines: up to **€35M** or **7%** of global turnover for prohibited practices (extraterritorial reach).

Regulation (EU) 2024/1689. Dates approximate; GPAI Code of Practice deemed adequate Aug 2025.

High-Risk Obligations and the GPAI Tier

PROVIDER DUTIES FOR HIGH-RISK SYSTEMS (Arts. 9–17)

- | | |
|--|---|
| ● Risk-management system (Art. 9) | ● Human oversight by design (Art. 14) |
| ● Data governance & bias control (Art. 10) | ● Accuracy, robustness & cybersecurity (Art. 15) |
| ● Technical documentation (Art. 11) | ● Quality-management system (Art. 17) |
| ● Automatic logging / traceability (Art. 12) | ● Conformity assessment + CE mark + EU registration |
| ● Transparency & instructions (Art. 13) | ● Serious-incident reporting (Art. 73) |

Deployers: use per instructions · assign human oversight · monitor & keep logs (6 mo) · FRIA for public bodies & credit/insurance uses.

GENERAL-PURPOSE AI (GPAI)

All GPAI models:

- technical documentation;
- a copyright policy;
- a public summary of training data.

Systemic-risk models (> 10²⁵ FLOPs) add:

- model evaluations & adversarial testing;
- systemic-risk assessment & mitigation;
- serious-incident reporting & cybersecurity.

EU vs. UK vs. China vs. United States

	Instrument	Posture	Enforcement
EU	Comprehensive AI Act (2024)	Prescriptive, risk-tiered, rights-based	AI Office + national authorities; heavy fines
UK	No single AI statute; sector regulators + principles	Pro-innovation, principles-based; targeted AI/copyright reform	Existing regulators (ICO, CMA, FCA)
China	Interim GenAI Measures (2023) + labeling rules (Sept 2025)	State-directed; content control & 'core socialist values'	CAC registration, security review, algorithm filing
U.S.	No comprehensive federal law; state patchwork + EOs	Market-driven; sectoral; deregulatory federal turn	FTC, state AGs; preemption push (Dec 2025 EO)

High-level comparison as of Sept. 2026. China: CAC = Cyberspace Administration of China. U.S. posture shifting toward federal preemption of state AI laws.

State AI Laws and the Federal Preemption Fight

Colorado

AI Act (SB 24-205): 'reasonable care' vs. algorithmic discrimination in high-risk AI. Amended; effective delayed to June 30, 2026.

Texas

Responsible AI Governance Act (TRAIGA): bans enumerated harmful uses; AG enforcement. Effective Jan. 1, 2026.

California

AB 2013 (training-data transparency); SB 53 (frontier-model safety); SB 942 (AI transparency, delayed to Aug 2026).

Illinois / Utah

IL: AI in employment decisions (HRA amendment). UT: AI Policy Act disclosure duties.

FEDERAL COUNTERCURRENT

Dec. 11, 2025 Executive Order establishes an AI Litigation Task Force to challenge 'onerous' state AI laws on preemption/commerce grounds, conditions certain federal funding, and directs FTC/Commerce review injecting real uncertainty into state compliance planning.

Effective dates and preemption posture are in flux; confirm current status before advising.

No AI Statute but Plenty of Federal Reach

FTC

FTC Act § 5 polices deceptive AI claims & unfair practices 'Operation AI Comply' enforcement sweep.

EEOC

Title VII & ADA reach AI hiring and assessment tools that screen out protected groups.

CFPB

ECOA requires specific adverse-action reasons even for complex 'black-box' credit models.

FDA

AI/ML-enabled medical devices regulated as software (SaMD), incl. predetermined change-control plans.

NIST

Not a regulator but its AI RMF is the de facto federal governance benchmark.

Executive orders

Biden EO 14110 (2023) rescinded; 2025 EOs pivot to pro-innovation + state-law preemption.

Existing statutes apply to AI today: FTC Act § 5, Title VII, ADA, ECOA, FD&C Act. 'No AI law' ≠ 'no law.'

Advising the Multinational: A Compliance Strategy

1

Map the footprint Where are users, data, and outputs? Extraterritorial laws (EU AI Act) reach in.

2

Classify by risk tier Fit the deployment to the strictest regime's categories (usually the EU's).

3

Build to the ceiling Design one governance program to the highest common denominator; localize deltas.

4

Paper the chain Push obligations to vendors by contract (Module 4); keep provenance & oversight records.

5

Govern & document Adopt NIST AI RMF / ISO 42001; log risk assessments, testing, and human oversight.

6

Monitor & adapt Assign ownership to track statutes, cases, and preemption through 2026+.

MODULE 6

Tracking the Horizon

Pending litigation, new statutes, and treaty obligations shaping client risk through 2026 and beyond.

6

Litigation & Regulatory Developments to Watch

Copyright docket

- NYT v. OpenAI summary judgment & the ChatGPT-logs discovery precedent
- Anthropic settlement administration (~\$1.5B)
- Thomson Reuters v. Ross 3d Cir. appeal on fair use
- Getty v. Stability; Authors Guild & music-industry suits

Statutes & agencies

- State AI laws taking effect 2026 (CO, TX, CA) and preemption litigation
- Federal AI executive order implementation & FTC/Commerce actions
- USPTO conception-based inventorship guidance in practice
- Copyright Office future guidance amid leadership change

International

- EU AI Act high-risk obligations phasing in (2026–28)
- Council of Europe AI Convention first binding AI treaty
- China labeling-rule enforcement
- UK AI/copyright reform outcome

The Litigation Map by Legal Question

Legal question	Representative matters	What a ruling will settle
Training = fair use?	<i>Bartz, Kadrey, Thomson Reuters v. Ross, NYT v. OpenAI</i>	Whether ingesting copyrighted works to train is lawful and the weight of 'market dilution.'
Output infringement / CMI	<i>Getty v. Stability; DMCA § 1202 claims (SDNY split)</i>	When outputs infringe, and whether stripping copyright-management info is actionable.
Defamation by AI	<i>Walters v. OpenAI</i>	Whether a provider is liable for hallucinated false statements about real people.
Privacy & biometrics	<i>BIPA / CCPA / wiretap class actions</i>	Liability for personal & biometric data used to train or operate models.
Deception / consumer	<i>FTC 'Operation AI Comply' actions</i>	Where AI marketing and capability claims cross into unlawful deception.

Organize your docket-watching by the question each case answers not just by party name. Status current to Sept. 2026; verify before relying.

What to Do Monday Morning for Your Clients



Inventory AI use Catalog where clients build, buy, and deploy AI and which hat they wear.



Assess the risk tier Classify each system; flag high-risk and customer-facing agentic uses.



Audit the contracts Check provenance reps, IP indemnity scope, caps, and training-on-input terms.



Stand up governance Adopt NIST AI RMF / ISO 42001; require documented human oversight.



Fix the confidentiality gap Set policy on entering client/PII data into third-party tools.



Assign a monitor Own the docket and the statute pipeline; brief clients on changes.

International Instruments Shaping the Field



For multinational clients: soft-law principles converge (transparency, accountability, human oversight), but binding force and enforcement still vary sharply by jurisdiction.

CoE Framework Convention on AI (CETS 225), opened for signature 5 Sept 2024; G7 Hiroshima Process (2023); OECD AI Principles (2019, updated 2024).

KEY TAKEAWAYS

Six Modules, Six Things to Remember

1

Competence is technical now

You must be able to explain LLMs, generative, and agentic AI Rule 1.1 requires it.

2

Liability follows the chain

Negligence, products, and vicarious theories target developers, deployers, and users differently.

3

IP law is unsettled but moving

2025 fair-use rulings and human-authorship/inventorship rules reshaped the landscape.

4

The contract is the control

Provenance, indemnity, SLAs, and caps allocate the risk the law leaves open.

5

Regulation is plural

One deployment can face the EU AI Act, China's rules, and a shifting U.S. patchwork at once.

6

Monitoring is the job

Pending cases, new statutes, and treaties will keep moving the line through 2026+.

Selected Authorities & Disclaimer

Selected authorities

- ABA Model Rule 1.1 cmt. [8]; ABA Formal Op. 512 (2024).
- Thomson Reuters v. Ross Intelligence, No. 1:20-cv-613 (D. Del. Feb. 11, 2025).
- Bartz v. Anthropic, No. 3:24-cv-05417 (N.D. Cal. June 2025); settlement (2025–26).
- Kadrey v. Meta Platforms, No. 3:23-cv-03417 (N.D. Cal. June 25, 2025).
- NY Times Co. v. Microsoft & OpenAI, No. 1:23-cv-11195 (S.D.N.Y.).
- Thaler v. Vidal, 43 F.4th 1207 (Fed. Cir. 2022); Thaler v. Perlmutter (D.C. Cir. 2025).
- U.S. Copyright Office, Copyright & AI, Parts 1–3 (2024–2025).
- USPTO Inventorship Guidance (Feb. 2024); Revised Guidance (Nov. 28, 2025).
- Regulation (EU) 2024/1689 (EU AI Act); EU Product Liability Directive (2024).
- Colorado SB 24-205; Texas TRAIGA; California AB 2013, SB 53, SB 942.
- China Interim Measures for Generative AI (2023) & AI-labeling rules (2025).
- Exec. Order on National AI Policy Framework (Dec. 11, 2025).

DISCLAIMER

These materials are for continuing legal education and general information only. They are not legal advice, do not create an attorney-client relationship, and may be incomplete or out of date.

AI law is evolving rapidly; many cited matters are pending or on appeal. Independently verify authorities and confirm the current posture before advising a client or relying on any point herein.

Jurisdiction-specific analysis is required for any actual engagement.