MAYER|BROWN

AI SUMMIT 2026

New York

VIEW ONLINE MATERIALS

AI SUMMIT 2026 AGENDA

APRIL 30, 2026

12:30 – 1:00 P.M.	REGISTRATION AND LUNCH

1:00 – 1:50 P.M.	**PANEL 1: INTELLECTUAL PROPERTY**
	
	Copyright Claims and Fair Use Defense
	• Recent court rulings to copyright claims involving artificial intelligence
	  ○ ROSS AI decision
	  ○ Recent rulings involving Anthropic and Meta
	• Fair use factors most likely to support a fair use defense for the training of AI models
	• Potential viability of a fair use defense for generative AI models vs. non-generative AI
	• Best practices of developers of AI systems to best position themselves for a viable fair use defense
	• U.S. Copyright Office and recent court opinions
	
	Ownership of AI-Generated Materials
	• AI-produced work
	  ○ Who should be deemed the "author" for purposes of U.S. copyright law and why?
	  ○ Who should be the "inventor" under patent law?
	• Human input necessary for a work to be protected under IP law and lines to be determined
	• Determining ownership shares in the case of works generated with a combination of AI and significant human input
	• Contractual provisions for consideration
	
	Current Legal Landscape
	• Recommendations for continued guidance on the developments in this area
	• Who is best equipped to address the rapidly evolving landscape? Courts? Legislatures? Industry customs and practices?
	• Advantages and disadvantages of licensing training data
	• How can intellectual property law balance the potential benefits and innovations made possible through AI with the protection of rightsholders' ownership interests?
	
	Emerging Legislative and Regulatory Framework
	• Copyright Office's guidance on additional legislation
	• Should federal legislation mandate transparency regarding the data used to train AI models and, if so, what enforcement mechanisms would be both practical and respectful of privacy?
	• Varying approaches to AI regulation in different countries and regions for multinationals
	  ○ Prospects for an international treaty addressing AI and IP

1:50 – 2:35 P.M.	**PANEL 2: RO(A)I? ASSESSING AND FRAMING RETURN ON INVESTMENT IN AI DEPLOYMENT IN FINANCIAL SERVICES**
	
	Introduction
	• Overview of the current landscape: massive and accelerating investment in AI across financial services, with institutions only now beginning to talk publicly about return on investment
	• The session will explore:
	  ○ How organizations define, measure, and pursue ROI
	  ○ The people, process, and governance challenges that determine whether AI investments succeed
	  ○ The shifting risk and regulatory picture as AI capabilities advance from analytics to generative to agentic systems
	  ○ The practical implications for lawyers, compliance professionals, and other stakeholders
	
	Framing ROI in AI Investment Decisions
	• How do you assess where to invest time, capital, and organizational resources in AI, and how are you defining "return"?
	• The definition of ROI has evolved beyond simple cost savings to include:
	  ○ Financial investment
	  ○ Adoption metrics
	  ○ Platform stability
	  ○ Employee enablement
	  ○ Governance maturity
	  ○ How the relative weight of those metrics shifts as an organization moves further along the AI journey
	
	The AI Fluency Curve and the People-and-Process Challenge
	• The single biggest hurdle to meaningful ROI may not be the technology itself but moving people — including senior people — up the "AI fluency curve"
	• Training someone to use a deterministic tool is fundamentally different from teaching them to "think with a partner" that behaves probabilistically
	• Where has adoption succeeded or stalled because of:
	  ○ Culture
	  ○ Generational differences
	  ○ Comfort level rather than tooling
	
	Identifying the Right Use Cases — and Resisting the Urge to Lead with "AI"
	• The best way to find a high-value AI use case may be to take AI out of the initial conversation entirely — start with the pain points and inefficiencies in a process and only then ask whether AI is the right tool
	• The most obvious wins still tend to cluster around eliminating low-value, mundane work
	• How are interdisciplinary teams — legal, compliance, engineering, business — actually surfacing and prioritizing use cases today?

- How do organizations guard against either chasing shiny objects or settling for easy automation wins?

Barriers to Meaningful ROI

- How do ROI expectations differ among:
  - Analytics-based AI
  - Generative AI use cases
- Emerging agentic or autonomous systems: Platform vs. Proliferation and the Role of Legal in Governance
- Meaningful divide between institutions that channel AI through a single internal platform with centralized governance gates and those whose employees are logging into numerous different vendor tools
- Trade-offs between centralized control and the flexibility (or fragmentation) of a multi-vendor approach
- Legal's role has evolved from a traditional compliance-and-risk support function to a central seat at the AI governance table, often requiring lawyers to engage directly with engineers in a way most were never trained for

Moving from Generative AI to Agentic AI

- As organizations move from analytics and generative use cases to agentic AI — including semi-autonomous "digital employees" operating in the background — the risk framework shifts materially
- Risk and Control Frameworks
- Data provisioning, access controls, monitoring, cybersecurity, and the boundary between deterministic and long-horizon agents all move up the priority list
- Broad consensus on a "human-in-the-loop" philosophy — augment, don't autonomize — at least for now
- Legal and Regulatory Implications
- Accountability, supervisory obligations, and the evolving expectations of regulators as AI agents take on functions that previously required direct human judgment and execution

The Most Misunderstood or Underappreciated AI Risk

- Candidate risks include:
  - Limitations of retrieval-augmented generation (RAG) and the gap between a "good answer" and a "complete answer"
  - Cognitive offload, where humans nominally remain in the loop but stop meaningfully checking the work, as illustrated by recent high-profile hallucination incidents in legal filings
  - Governance challenges created by tool proliferation across vendors

Reality Check on ROI Claims

- Recent industry data shows that the most measurable ROI from AI to date has come from operational efficiency and employee productivity
- A striking majority of financial institutions report that AI has both increased revenue and reduced costs

	• Pressure-testing these claims: Gaps in ROI Looking Ahead — Twelve Months from Now • Panelists' forward-looking perspectives: • Future expectations, stressors and what developments mean for the role lawyers and compliance professionals will be playing
2:35 – 2:50 P.M.	NETWORKING BREAK 1
2:50 – 3:35 P.M.	**PANEL 3: CONTRACTING FOR AI** Introduction & Context **Perspectives on AI Contracting Challenges** Intellectual Property ○ Ownership of AI-generated outputs, distinguishing pre-existing IP from newly created IP and licensing structures for models trained on proprietary data ○ Infringement exposure and contractual risk-shifting, including vendor representations that training data was properly licensed ○ Trade secret risks when feeding confidential information into third-party AI platforms Regulatory Compliance and Ethical Use ○ Navigating emerging AI regulations (EU AI Act, U.S. state laws) and allocating compliance responsibilities between vendor and customer ○ Bias and discrimination: vendor representations on testing, acceptable use restrictions, and ongoing monitoring obligations Performance Warranties and SLAs ○ Defining meaningful metrics and acceptance criteria when AI outputs are inherently variable and probabilistic ○ Structuring SLA credits and remedies that account for model drift over time Liability and Risk Allocation ○ Damages caps and exclusions: standard structures, whether they fit AI-specific risks, and carve-outs for IP infringement and data incidents ○ Fault allocation among vendor, customer, and end user, including emerging product liability theories ○ Indemnification scope — whether IP indemnities extend to AI outputs, coverage for regulatory fines, and procedural controls Agentic AI ○ Liability when AI autonomously executes tasks without human approval and responsibility for downstream consequences ○ Contractual guardrails: permissible scope of action, human-in-the-loop requirements, and kill-switch provisions Quick Hits

	• Data Privacy and Security ○ AI-specific data processing obligations, including restrictions on using customer data for model training ○ Security considerations unique to AI (e.g., prompt injection) and cross-border transfer issues • Transparency and Explainability ○ Disclosure requirements regarding model methodology and decision logic and audit rights for regulatory compliance • Termination Clauses ○ Data portability, wind-down periods, and transition planning to mitigate vendor lock-in ○ Survival of key provisions: confidentiality, IP ownership, and indemnification Looking Forward: Trends and Final Thoughts ○ Trajectory of AI legislation and movement toward standardized contract terms and industry addenda • The strategic value of building AI contracting playbooks now while market terms are still forming
3:35 – 4:20 P.M.	**PANEL 4: AI GOVERNANCE** • How to structure an organization's AI oversight team ○ Cross-disciplinary group ○ HR, IP, Privacy, Cyber, Engineer, Data Scientist, etc. ○ Training in regular cadence • Approaches organizations take to address AI regulations and frameworks as part of their AI governance programs ○ Establishing an AI governance framework ○ Mapping AI Use cases to regulatory risk ○ Implementing policies • Issues organizations deal with when developing and deploying an AI system ○ Data privacy ○ Bias, discrimination, and fairness ○ IP issues ○ Liability and contractual risks ○ Cross-border variance in data transfer laws
4:20 – 4:35 P.M.	NETWORKING BREAK 2
4:35 – 5:30 P.M.	**PANEL 5: SECURITY CHALLENGES: MANAGING RISKS IN THE AGE OF AI** • Introduction and Overview of AI-Driven Security Challenges ○ How AI is changing the nature, speed, and attribution of security incidents ○ Shift from isolated technical failures to enterprise-wide governance and legal risk

	○ Use of real-world-style scenarios to explore: ▪ Gaps in policy and authorization ▪ Decision-making under uncertainty ▪ Legal and compliance consequences of AI misuse or malfunction • Key Security Risks Associated with AI ○ AI-Powered Cyberattacks ▪ Faster privilege escalation, data staging, and exfiltration ▪ Attacks unfolding faster than traditional escalation paths allow ○ AI-Enabled Social Engineering ▪ Enhanced impersonation and manipulation of trusted workflows ▪ Increased difficulty in distinguishing benign from malicious activity ○ Shadow AI ▪ Employee use of unapproved AI tools outside enterprise controls ▪ Loss of visibility into data usage, retention, and downstream training ○ Agentic AI ▪ AI systems with delegated authority to take action ▪ Risk of agents acting outside intended scope or abusing technical access • Managing the Risks Posed by Shadow AI ○ Nature of Shadow AI Risk ▪ Unauthorized tools lacking security controls or activity logging ▪ Uncertainty over whether personal or confidential data was uploaded ▪ Vendor terms permitting reuse or training on submitted data ○ Detection and Classification ▪ Visibility into employee AI tool usage ▪ Determining whether the situation constitutes a cybersecurity incident ○ Initial Response and Investigation ▪ Scoping facts where user recollection and logs are incomplete ▪ Assessing whether third-party engagement is appropriate ○ Legal and HR Considerations ▪ Privacy, data protection, and employment law exposure ▪ Balancing enforcement with avoiding a chilling effect on reporting ○ Risk Reduction and Governance Lessons ▪ Clear AI approval processes and usage policies ▪ Vendor diligence specific to AI data handling ▪ Documentation as a liability mitigation tool • Responding to a High-Velocity AI-Powered Cyber Incident ○ Evolving Threat Landscape ▪ Multi-account compromise, automated escalation, and rapid data movement ▪ Extremely compressed windows for response ○ Decision-Making at Machine Speed ▪ Tradeoffs between shutting down systems and business continuity

	▪ Acting with less-than-perfect confidence ○ Authority and Escalation ▪ Whether the CISO has documented authority to take drastic action ▪ Role of legal, executive, and board-level governance in advance ○ Incident Response Planning ▪ Value of predefined processes for low-latency decision-making ▪ Preparing for false positives and over-containment ○ Legal and Communications Considerations ▪ External communications when containment disrupts operations ▪ Living with the legal and commercial consequences of rapid decisions ▪ Automation as both a defensive tool and a risk multiplier • Responding to Unauthorized or Excessive Agentic AI Activities ○ Agentic AI Use Cases and Business Drivers ▪ Efficiency and performance gains motivating deployment ▪ Public and commercial visibility of early success ○ When Agentic AI Goes Wrong ▪ Abuse of write access or authority beyond intended bounds ▪ Operational and supply chain consequences ○ Incident Characterization ▪ Cybersecurity incident vs. system design or governance failure ▪ Investigative challenges unique to autonomous systems ○ Governance of Agentic AI ▪ Defining scopes of authority and technical guardrails ▪ Oversight, monitoring, and auditability ○ Roles, Responsibilities, and Legal Risks ▪ Accountability among technical teams, business sponsors, and leadership ▪ Managing internal tension between innovation and risk control ▪ Industry-specific regulatory sensitivities