



MAYER|BROWN

GOVERNING AUTONOMOUS SYSTEMS: REGULATORY COMPLIANCE AND ENTERPRISE OVERSIGHT

Arsen Kourinian, Lei Shen, Oliver Yaros

AGENDA

1. Overview of AI Regulatory Landscape
2. Agentic AI Governance in Contracts
3. Broader AI Governance Considerations

01

OVERVIEW OF AI REGULATORY LANDSCAPE

AGENTIC AI IS A SUBSET OF AI SYSTEMS REGULATED UNDER AI LAWS

EU AI Act

"AI system" means a machine-based system that is designed to operate with ***varying levels of autonomy*** and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments.

European Commission Guidance

Central to the concept of autonomy is "human involvement" and "human intervention" and thus human-machine interaction. At one extreme of possible human-machine interaction are systems which are designed to perform all tasks though manually operated functions. At the other extreme are systems that are capable to operate ***without any human involvement or intervention, i.e. fully autonomously.***

California AI Transparency Act

"Artificial intelligence" or "AI" means an engineered or machine-based system that ***varies in its level of autonomy*** and that can, for explicit or implicit objectives, infer from the input it receives how to generate outputs that can influence physical or virtual environments.

South Korea AI Basic Act

The term "AI system" means an artificial intelligence-based system that infers outputs such as predictions, recommendations, and decisions that affect real and virtual environments for given goals ***with various levels of autonomy and adaptability.***

DOMESTIC AND GLOBAL AI LAWS

Comprehensive AI Laws

- EU AI Act
- Colorado ADMT Act
- South Korea's Artificial Intelligence Development and Establishment of a Foundation for Trustworthiness
- Vietnam Law on AI (No. 134/2025/QH15)

Narrower AI Laws in the US

- Texas Responsible AI Governance Act (prohibited practices)
- Frontier Model Laws (California, New York and Illinois)
- Utah Generative AI Law (transparency)
- Illinois AI law amending Human Rights Act (transparency and avoiding discrimination)
- Illinois AI Video Interview Act (transparency and consent)
- New York City Local Law 144 (transparency and bias audit)
- California Generative AI: Training Data Transparency (AB2013) (transparency regarding training data)
- California AI Transparency Act (SB942) (AI detection tool and manifest and latent disclosures on AI-generated content)
- Chatbot Laws (transparency and safety protocols)

Interaction with Existing Laws

- Data privacy law considerations (state Attorneys' General and California Privacy Protection Agency)
 - **CCPA: New regulations on "Automated Decisionmaking Technology"**
- Employment laws (Equal Employment Opportunity Commission)
- Unfair and deceptive practices (Federal Trade Commission)
- Explaining credit decisions following adverse action (Consumer Financial Protection Bureau)

PROHIBITED AI SYSTEMS

EU AI ACT	TEXAS RESPONSIBLE AI GOVERNANCE ACT
Use subliminal, manipulative or deceptive techniques to distort a behavior/decision	It is prohibited to develop or deploy an AI system with the intent to incite or encourage a person to : (a) commit physical self-harm, including suicide; (b) harm another person; or (c) engage in criminal activity.
Exploit vulnerabilities like age, disability, or social or economic situation to distort behavior	It is prohibited to develop or deploy an AI system with the sole intent for the AI system to infringe, restrict, or otherwise impair an individual's rights guaranteed under the United States Constitution.
Classify individuals based on social behavior or personal characteristics (social scoring) leading to detrimental or unfavorable treatment that is out of context or disproportionate	It is prohibited to develop or deploy an AI system with the intent to unlawfully discriminate against a protected class in violation of state or federal law. The law notes, however, that a disparate impact is not sufficient by itself to demonstrate an intent to discriminate. Insurance entities and federally insured financial institutions are exempt from this prohibition.
Make risk assessments of individuals to assess or predict the risk of them committing a crime (predictive policing)	It is prohibited to develop or distribute an AI system with the sole intent of producing, assisting or aiding in producing, or distributing unlawful visual material or deepfake videos or images . It is also prohibited to develop or distribute an AI system that engages in text-based conversations that simulate or describe sexual conduct , while impersonating or imitating a child under 18.
Untargeted scraping of the internet or CCTV for facial images to build or expand facial recognition databases	
Inferring emotions in the workplace and educational institutions, unless for medical or safety reasons	
Biometric categorization systems based on biometric data (like facial images or fingerprints) to deduce or infer individuals' race, political opinions, trade union membership, religious or philosophical beliefs, or sexual orientation	
Real-time biometric identification system in publicly accessible spaces by law enforcement, subject to narrow exceptions	

HIGH-RISK AI SYSTEMS / AUTOMATED DECISION-MAKING TECHNOLOGY

EU AI ACT	COLORADO ADMT ACT
Education and vocational training	Education enrollment or an education opportunity
Employment, worker's management, and recruitment	Employment or an employment opportunity that creates or may create an employer-employee relationship
Essential public and private goods, services, and benefits • Evaluate the creditworthiness of natural persons or establish their credit score • Life and health insurance	Essential government services and public benefits, including eligibility and renewal determinations
Law enforcement, administration of justice and democratic processes, and immigration and border control	A financial or lending service
Biometrics • Remote biometric identification system (except ID verification) • Biometric categorization based on sensitive/protected attributes • Emotion recognition	Insurance, including underwriting, pricing, coverage, claims adjudication, or other determinations that materially affect access to benefits
Safety components in critical infrastructure (e.g., digital infrastructure, road traffic, or in the supply of water, gas, heating, or electricity)	Healthcare services
Safety component in regulated products	The lease or purchase of residential real estate in Colorado

HIGH-RISK OBLIGATIONS UNDER THE EU AI ACT

PROVIDER	DEPLOYER
Risk management system	Follow instructions of use
Data governance	Human oversight
Maintain technical documentation and logs	Relevant and representative input data
Transparency and instructions of use to deployer	Monitoring and notice of serious incidents
Ensure accuracy, robustness, cybersecurity, and accessibility	Maintain autogenerated logs
Quality management system and declaration of conformity	Impact assessments
Human oversight	Transparency
Name, trademark, CE marking, and registration	Honor data subject rights, including automated decisions
Post-market monitoring, corrective measures, and reporting of serious incidents	

HIGH-RISK OBLIGATIONS UNDER THE COLORADO ADMT ACT

DEVELOPER	DEPLOYER
Transparency: Provide information to the deployer regarding the Covered ADMT, including the intended uses, harmful or inappropriate uses, training data used, known limitations and risks, instructions for use, monitoring, and meaningful review, and other information necessary for the deployer's compliance.	Transparency: Provide a pre-use notice to consumers that the deployer is using a Covered ADMT to materially influence a consequential decision and how a consumer may obtain additional information under the law. This requirement may be satisfied by posting a public notice that is reasonably accessible at points of consumer interaction, including through a link or posting reasonably proximate to the interaction or transaction.
Change notice: Provide a notice to the deployer regarding material updates, intentional and substantial modifications, and changes to the intended use of, limitation for, or risk-mitigation for the Covered ADMT within a reasonable amount of time.	Adverse outcome notice and rights: Provide a notice to consumers within 30 days of the Covered ADMT materially influencing a consequential decision that results in an adverse outcome. The post-adverse outcome notice must include a description of the decision and Covered ADMT's role and instructions on how to request additional information, including regarding the consumer's rights to access and correct factually or materially inaccurate personal data and an opportunity for meaningful human review and reconsideration of the consequential decision, to the extent commercially reasonable.
Records: Maintain records demonstrating compliance for not less than three years after creation, including system version identifiers, changelogs, and documentation and notices of material updates provided to deployers.	Records: Maintain records demonstrating compliance for not less than three years or longer if required by other law.

EU AI ACT – TRANSPARENCY OBLIGATIONS

PROVIDER	DEPLOYER
Transparency about user interaction with an AI system	Transparency about using emotion recognition or biometric categorization systems
Marking synthetic audio, image, video, or text as such	Transparency about deepfake image, audio, or video content
	Transparency about artificially generated or manipulated text published with the purpose of informing the public on matters of public interest

CHATBOT LAWS IN THE UNITED STATES

CHATBOTS	COMPANION CHATBOTS
Basic transparency and chatbot laws in California, New Jersey, Utah, and Maine require making it clear to users that they are interacting with AI, and not a human being.	Some states have companion chatbot laws (California, New York, Georgia, Idaho, Iowa, Nebraska, Oregon, Washington, and Connecticut)
	Maintain and publish protocols for suicide / self-harm prevention
	Notifications with certain frequency (e.g., every hour or every three hours) that the user is interacting with AI, and not a human being
	Prevent chatbot from discussing sexual content with minors

CALIFORNIA GENERATIVE AI: TRAINING DATA TRANSPARENCY

- Effective January 1, 2026
- Developers of AI systems (or those that substantially modify it, e.g., finetuning or retraining) must provide a high-level summary on their website about the training data used to train the generative AI system.
- This law emphasizes the importance of maintaining a data provenance record.

CALIFORNIA AI TRANSPARENCY ACT

- Effective August 2, 2026
- Developers of generative AI systems that are made publicly accessible and have more than one million monthly (*subject to change SB1000*) visitors and users need to:
 - (1) Make available an AI detection tool at no cost to the user;
 - (2) Offer the user the option to include a manifest disclosure in image, video, or audio content clearly identifying the content as AI-generated; and
 - (3) Include a latent disclosure in AI-generated image, video, or audio content.

CONNECTICUT SB5& WASHINGTON HB1170

- Effective October 1, 2026 (CT) and February 1, 2027 (WA)
- Covered providers that create, code, or otherwise produce a generative AI system that has more than one million monthly visitors and is publicly accessible to consumers for personal use are required to:
 - (1) Include provenance data in audio, image, or video content; and
 - (2) Make provenance data difficult to tamper with, remove, or disassociate.
- Doesn't apply to:
 - B2B use, sale, licensing, or distribution of generative AI;
 - Video game or interactive experiences; or
 - Systems used solely for upscaling, noise reduction, or compression.

02

AGENTIC AI GOVERNANCE IN CONTRACTS

UNDERSTAND YOUR USE CASE

- The customer and agentic AI provider should discuss whether the AI system is intended to engage in high-risk or consequential decision-making.
- If not, terms can be added to clarify that the agentic AI will not be deployed in these regulated areas.
- If the agentic AI will be involved in high-risk or consequential decisions, the parties may consider addressing party-role obligations based on applicable regulations.
- Parties may also clarify that neither party will engage in prohibited practices under AI regulations.



DATA GOVERNANCE

- The customer should evaluate which databases it wants to give the agentic AI access to.
- Consider restricting access to databases that contain sensitive data or trade secrets.
- Evaluate what external sources the agentic AI can access to execute actions.




TRANSPARENCY

- If the agentic AI will communicate with external parties, communicate disclosure that individuals are interacting with AI and not a human being.
- If the agentic AI will collect personal information from individuals, discuss method of communicating a privacy notice to data subjects at or before the point of collection.
- Discuss flowing down your terms of use and acceptable use policies to the agentic AI providers.
- Identify contact person in company that can intervene if agentic AI deviates and external parties raise concerns.

HUMAN OVERSIGHT

- Evaluate circumstances where human oversight and intervention necessary with agentic AI.
- The contract may specify at what juncture a human needs to be alerted for AI decision-making and what type of decisions and actions do not require intervention.



An abstract graphic on the left side of the slide, consisting of several overlapping, translucent blue rectangular planes that create a sense of depth and movement, resembling a staircase or a series of steps ascending towards the top right.

TECHNICAL CONTROLS

- Contract terms may consider applying technical controls to limit the agentic AI's actions, conduct continuous monitoring of the agentic AI, and maintenance of logs.



LATENT DISCLOSURES

- The contract may contain a prohibition on the licensee tampering with the agentic AI's ability to include latent disclosures in AI-generated content, to address California AI Transparency Act requirements.

CONTRACTUAL AND PROCUREMENT ISSUES WHEN PROCURING AGENTIC AI

Due Diligence and Procurement:

- Assess the vendor's AI governance maturity, security posture and incident history
- Identify party roles at the outset (provider, deployer etc) as obligations follow role
- Diligence any embedded or downstream LLM the vendor relies on, including its online terms
- Confirm whether the vendor trains on customer data by default and whether that can be switched off
- Map the intended use case and confirm it is not high-risk or prohibited before terms are drafted
- Decide the contracting vehicle: bespoke AI addendum, or an AI policy incorporated by reference

Regulatory Compliance:

- EU AI Act: role allocation, transparency and high-risk obligations where triggered
- UK/EU GDPR: DPIA, lawful basis, data subject rights and transfer controls
- Sector specific reg: e.g. UK FCA operational resilience, outsourcing and critical-third-party expectations for regulated clients
- Sector and consumer-facing conduct requirements where the agent interacts with customers
- Cooperation and information obligations to support the client's own regulatory reporting

Provisions to Prioritize:

- Warranties tied to documented specifications and instructions of use, not perfection
- Liability: carve-outs from the cap for data breach, confidentiality and IP infringement
- IP: input and output ownership, express assignment of output rights, and an output infringement indemnity
- Data governance: access scoping, no-training commitments, data processing addendum
- Oversight: human-in-the-loop triggers, technical action limits, monitoring and logging
- Audit and information rights extending to subcontractors and the underlying model

Exit and Transition:

- Portability of data, logs, audit trails and configuration in usable form
- Deletion or return of client data, with certification
- Restrictions on retaining client data, derived embeddings or model weights for retraining
- Termination for cause on material breach, serious incident or loss of rights to the underlying model
- Transition assistance, step-in rights and a run-off period

SCENARIO 1 – AUTONOMOUS CUSTOMER-FACING SERVICE AGENT

Scenario

You license an agentic AI assistant that interacts directly with customers, answers queries and takes autonomous actions
Vendor's product is built on a third-party LLM; agent collects personal data in the course of interaction

Key Warranties

Performance against documented specifications (not perfection)
Fitness for the client's disclosed purpose; compliance with applicable AI and data protection law
No use of client inputs or outputs to train the vendor's models
Tool retains latent and manifest disclosure capabilities

Critical Contractual Terms

Transparency: customers told they are interacting with AI; privacy notice at point of collection
Flow-down of AUP; named intervention contact for deviations
Technical controls, continuous monitoring and log retention
Carve-outs from liability cap for data breach, confidentiality and IP infringement
Audit and information rights

Risk Allocation

Resist terms shifting all AI-decision responsibility onto you
Specify human-alert triggers; define actions not requiring intervention
Allocate liability for erroneous commitments the agent makes to customers

Regulatory & Compliance

Confirm tool not deployed for high-risk decision-making unless expressly scoped
DPIA and data-subject-rights obligations under UK/EU GDPR
Sector specific operational resilience expectations where client is regulated

Exit & Transition

Export of logs, customer records and configuration
Deletion of client data; no retention for model training
Termination for cause on material breach, serious incident or loss of LLM rights

SCENARIO 2 — AGENTIC AI EMBEDDED IN THE CLIENT'S PRODUCT

Scenario

You source an AI component or foundation model from a vendor and embeds it into your own product, licensed to customers under the client's name and brand

The product is used by customers in the EU

Key Warranties

Vendor has all rights and licenses in training data; model appropriately sourced

Outputs do not infringe third-party IP; no open-source or copyleft contamination

GPAI transparency documentation is current

Conformity with vendor's declared specifications

Critical Contractual Terms

IP ownership: training data to provider; input and output data to client; express assignment of output rights

IP infringement indemnity for outputs (vendors offer only limited, conditional indemnities)

Escrow or continuity of model and documentation

Reasonable change-management timeframes

Risk Allocation

Role allocation is the central issue: white-labelling or substantial modification transforms you into the provider

Watch vendor indemnities shifting liability for fine-tuned or modified models onto you

Resist warranting downstream deployer obligations you cannot control

Regulatory & Compliance

EU AI Act extraterritorial reach where AI system is on the EU market or output used in the EU

Identify whether provider obligations are triggered and plan compliance accordingly

Exit & Transition

Continued license for tail of existing customer contracts

Restrictions on vendor retraining on client data or derivative weights

Termination for cause on IP challenge, insolvency or withdrawal of the model

SCENARIO 3 – AUTONOMOUS SOFTWARE ENGINEERING AGENT

Scenario

You deploy an agentic coding tool with access to source repositories that generates code, opens pull requests and can merge changes within defined limits

Tool draws on a third-party foundation model and public code

Key Warranties

Training data appropriately sourced and licensed

Generated code does not infringe third-party IP and is free of open-source or copyleft contamination

Provenance information available for generated code

Security warranties covering the tool's own vulnerabilities and repository access

No use of proprietary code to train or improve the vendor's models

Critical Contractual Terms

Ownership of generated code vests in you, with express assignment of any vendor rights

IP infringement indemnity for outputs, with the vendor's conditions negotiated and disclosed

Audit rights and vulnerability disclosure and remediation commitments

Risk Allocation

Resist terms shifting all AI-decision responsibility onto you

Specify human-alert triggers; define actions not requiring intervention depending on deployment context

Allocate liability for erroneous code changes the agent makes that impact high risk systems

Regulatory & Compliance

Confirm tool not deployed for decision-making affecting high risk systems unless expressly scoped

Sector specific operational resilience expectations where client is regulated

Exit & Transition

Export of logs, repository changes records and configuration

Deletion of code designated as client data; no retention for model training

Usual termination for cause on material breach, serious incident or loss of LLM rights

03

BROADER AI GOVERNANCE CONSIDERATIONS

GLOBAL AI GOVERNANCE

ACCOUNTABILITY

- Leadership Policy
- Developer Policy
- Deployer Policy
- Employee AI Use Policy


AI GOVERNANCE TEAM

- Cross-disciplinary group
- HR, IP, Privacy, Cyber, Engineer, Data Scientist, etc.
- Training in regular cadence
- Understanding company's AI use cases/inventory

MITIGATION MEASURES

- Transparency & Explainability
- Fair and unbiased
- Accuracy
- Robustness
- Safe and secure
- Privacy-enhanced
- Human oversight

- Technical documentation and logs
- Continuous monitoring
- Contestability
- **Contract and oversight of AI vendors and model providers**
- Decommissioning AI system



DATA GOVERNANCE

- Input data for use of AI
- Training data for AI development
- Maintaining data provenance record

LEGAL COMPLIANCE

RISK MANAGEMENT

- Identify/rank risks
- Document risk assessment via AI impact assessment

Vendor management just one subcomponent

MAYER | BROWN

This Mayer Brown publication provides information and comments on legal issues and developments of interest to our clients and friends. These materials are provided by Mayer Brown and reflect information as of the date of presentation. The foregoing is not a comprehensive treatment of the subject matter covered and is not intended to provide legal advice. The contents are intended to provide a general guide to the subject matter only and should not be treated as a substitute for specific advice concerning individual situations.

Readers should seek legal advice before taking any action with respect to the matters discussed herein.

You may not copy or modify the materials or use them for any purpose without our express prior written permission.

Mayer Brown is a global legal services provider comprising associated legal practices that are separate entities, including Mayer Brown LLP (Illinois, USA), Mayer Brown International LLP (England & Wales), Mayer Brown Hong Kong LLP (a Hong Kong limited liability partnership) and Tauil & Chequer Advogados (a Brazilian law partnership) (collectively, the "Mayer Brown Practices"). The Mayer Brown Practices are established in various jurisdictions and may be a legal person or a partnership. PK Wong & Nair LLC ("PKWN") is the constituent Singapore law practice of our licensed joint law venture in Singapore, Mayer Brown PK Wong & Nair Pte. Ltd. More information about the individual Mayer Brown Practices and PKWN can be found in the Legal Notices section of our website.

"Mayer Brown" and the Mayer Brown logo are the trademarks of Mayer Brown. © 2025 Mayer Brown. All rights reserved.