AI Governance:

Policies, Procurement, and Product Risk

Frankfurt Kurnit Tech Law Summit

May 12, 2026

Daniel Goldberg and Andrew Folks

Where Companies Struggle with AI Governance

1. Misunderstanding What AI Governance Actually Is
2. Not Knowing What's Running in Your Org
3. Vendor Contracts and Training Data
4. High-Risk AI Systems
5. Conversational AI
6. AI Training Transparency

Issue 1: Misunderstanding What AI Governance Actually Is

AI Policy is Not AI Governance

- Many companies treat AI governance as a policy question: draft acceptable use rules, publish them, done. That is not governance.

- AI governance is a framework of policies, processes, controls, and oversight used to ensure AI systems are deployed responsibly, safely, and in compliance with legal and business requirements.

- AI governance covers:
  - AI inventory and use-case mapping
  - Risk and impact assessments
  - Bias, fairness, and discrimination testing
  - Model testing, monitoring, and validation
  - Vendor and third-party oversight
  - Transparency, disclosures, and consumer rights
  - Incident response and remediation

Using NIST AI Risk Management Framework as the Organizing Structure

- A practical, risk-based framework that is jurisdiction-agnostic and increasingly treated by regulators as a de facto benchmark.

- Four Core Functions:
  - Govern – accountability and oversight
  - Map – systems, data flows, and impacts
  - Measure – risk, bias, reliability
  - Manage – mitigation and continuous improvement

Issue 2:
Not Knowing What's Running in Your Org

Shadow AI: The Inventory Problem

- Employees are using AI tools that legal never approved, never reviewed, and may not even know exist. The data flowing through them includes confidential, personal, and proprietary business information.

- How It Happens:
  - Prohibition without a workable alternative leads to circumvention
  - Team-level tool adoption without procurement involvement
  - SaaS agreements signed outside of legal review
  - Pilots that quietly become production deployments

- Dangers:
  - You cannot assess risk against systems you don't know exist
  - Retroactive remediation (data deletion, contract renegotiation, breach notification) is expensive, disruptive, and often incomplete

The Fix: Intake and Inventory

- For any AI tool under consideration, ask:
  - What data does this tool process?
  - Does the vendor train on customer or employee data?
  - Does it make significant decisions about people?
  - Has it been reviewed against governance criteria and applicable regulatory obligations?

- Maintain a detailed inventory: what tools are approved, who approved them, when, and under what conditions.

- Prohibition without a workable alternative drives Shadow AI underground, not away. Make disclosure easy, approval achievable, and consider an amnesty window for tools already in use.

- This is not a one-time exercise. Intake and inventory should be embedded in your existing risk assessment and vendor management processes.

Issue 3:
Vendor Contracts and Training Data

Service Provider or Third Party?

- When an AI vendor trains on personal information, are they a service provider under CCPA or have they become a third party?

- The answer determines:
  - Whether a sale or sharing has occurred
  - What opt-out and disclosure obligations apply
  - Who bears liability when something goes wrong

- The boundaries of the service provider exception are unsettled, particularly where vendors use data to improve models that serve other customers

- The CCPA regulations permit a service provider to use personal information:
  - "to build or improve the quality of the services it is providing to the business . . . provided that the service provider . . . does not use the personal information to perform services on behalf of another person."

Deidentification: Promise and Pitfall

- Deidentification is both a compliance strategy and a source of risk. Companies use it to unlock data for AI development and potentially take that data outside the scope of certain privacy laws, but only if strict legal and technical standards are actually met.
- The problem: most vendor claims that data is "anonymous" or "deidentified" do not explain the actual methodology, safeguards, or testing used.
- LLMs present unique risks: models can retain, reproduce, or allow inference of information from training data even when that data was never meant to persist.
- A key question: Was the data actually deidentified before ingestion into the LLM or did deidentification happen after the fact at all?
- Under AB 1008 (effective January 1, 2025), CCPA personal information includes "artificial intelligence systems that are capable of outputting personal information."

Your Data May Be Training Competitor Models

- Training data can reveal confidential business information, workflows, pricing strategies, customer insights, and product direction, even when it contains no personal information.
- When your data trains a vendor's model, the resulting benefits may accrue to the vendor or even your competitors.
- Aggregated or transformed data can still expose competitive intelligence and internal decision-making patterns.
- AI vendor diligence should evaluate not only privacy risk, but confidentiality obligations, IP leakage, and long-term competitive impact.

What Your AI Vendor Contract Should Say

- **Training restrictions:** Prohibit training on customer data except as permitted for a service provider under applicable privacy law, with CCPA-compliant limitations expressly referenced.
- **Deidentification standards:** Define deidentification standards with specificity, not left to vendor interpretation.
- **Output ownership:** Establish clear ownership of AI-generated outputs, derivatives, embeddings, and improvements.
- **Data use restrictions:** Prohibit use of customer data for product improvement, benchmarking, or shared or multi-tenant model development.
- **Deletion and retention:** Define data deletion and retention obligations triggered at termination, not just upon request.
- **Audit rights:** Secure access to technical documentation regarding AI training, data retention, and security practices, with meaningful audit rights attached.

Issue 4:
High-Risk AI Systems

A Growing Patchwork Regulating High-Risk AI

Law	Key Requirement	Status
California CCPA ADMT Regs	Risk assessments, notices, access and opt-out rights	Effective Jan. 1, 2027
Colorado AI Act	Annual impact assessments, consumer rights, antidiscrimination	Effective Jan. 1, 2027 (will be amended)
Connecticut SB5	Disclosures, safety protocols, subscription transparency	Awaiting Gov. signature
New York LL 144 Illinois HB 3773	Bias audits, notice to job candidates	In effect & enforceable
EU AI Act	Risk-tiered obligations	Phased implementation through 2028

EU AI Act – Brief Update

- Risk-based framework (unacceptable / high / limited / minimal) with layered obligations for providers and deployers.

- Recent developments – Provisional Agreement
  - High-risk obligations delayed until Dec. 2, 2027, with some to 2028.
  - Digital watermarking obligation delayed until Dec. 2, 2026.
  - Enforcement for General Purpose AI models centralized to AI Office.
  - Personal data processing permitted where strictly necessary to detect and correct biases in AI systems, subject to safeguards.

- What's NOT changing:
  - Transparency obligations still on for Aug. 2, 2026.
  - AI literacy requirement unchanged.

When Does AI Become High-Risk?

- Under CCPA, ADMT is a technology that uses computation to replace or substantially replace human decision-making for a "significant decision."

- Significant Decisions:
  - Employment or independent contracting
  - Financial or lending services
  - Housing
  - Healthcare
  - Education

- Advertising — Not included. But not a blanket exception.

Obligations Once You're In Scope

- For Covered Systems:
  - Pre-use notices to consumers
  - Opt-out mechanisms (unless the decision can be appealed to a human reviewer)
  - Consumer access rights to logic and output
  - Risk assessments (requiring executive attestation signed ***under penalty of perjury***)

- Compliance Timeline
  - Risk assessments: already required as of January 1, 2026
  - Consumer-facing notices, and access and opt-out rights: required January 1, 2027

- Employment Is the Highest-Impact Area:
  - Connecticut SB5: advance disclosures and written notice to applicants and employees
  - NYC LL 144: annual bias audit required before deployment
  - California: employment-specific CCPA regulations expected

Issue 5: Conversational AI

The Chatbot Regulatory Landscape

- In Q1 2026 alone, 36 states introduced over 70 bills targeting AI chatbots. This is one of the most active areas of AI legislation right now.

- At least ten states have laws in effect or imminent: New York, California, Washington, Oregon, Nebraska, Idaho, Iowa, Utah, Colorado, and Tennessee, with more expected before year end.

- In the absence of federal legislation, this patchwork is accelerating. Companies deploying conversational AI across multiple states face overlapping and sometimes conflicting obligations with no single compliance standard.

- Regulatory Focus: Categories under Scrutiny
  - Companion and social AI
  - Mental health AI
  - Minor-facing AI
  - Licensed professional
  - General purpose and customer service

- Litigation Reality: Wrongful death suits have already been filed. California creates a private right of action at $1,000 per violation plus attorneys' fees.

Obligations Once You're In Scope

- Disclosure:
  - Affirmative disclosure of AI status — baseline requirement across all jurisdictions

- Mental Health and Crisis Protocols:
  - Detect expressions of suicidal ideation or self-harm
  - Refer users to crisis service providers
  - Prohibit chatbots from impersonating licensed mental health professionals

- Minor-Specific Protections:
  - Recurring reminder every 3 hours that the chatbot is not human and to take a break
  - Block sexually explicit content
  - Prohibit fostering emotional dependency or simulating romantic relationships
  - Annual reporting obligations (California)

- Exemptions (Only Some States)
  - Customer service, education, financial services, healthcare support, business productivity, or technical assistance.

- Don't forget UCL, privacy laws, consumer protection laws, etc.!

Chat Data

- Many companies are using AI to review chat conversations for purposes of moderation, personalization, etc.
- Takeaways
  - Secondary uses and purpose limitation principle may require consent
  - Wiretapping risk
  - Third-party vendor oversight

Issue 6: Training Data Transparency

California AB 2013 — Effective January 1, 2026

- Who Is Covered:
  - Developers of generative AI systems made publicly available to Californians
  - "Developer" defined broadly: anyone who designs, produces, or *substantially modifies* an AI system (including via retraining, fine-tuning, or potentially adding a wrapper to a model)

- Developers of publicly available generative AI systems must disclose a high-level summary of:
  - Sources or owners of the datasets
  - Whether datasets were purchased or licensed, or obtained from public domain
  - Whether datasets were processed or modified by developer
  - Descriptions of how datasets further purpose of the AI system/service
  - Number and description of data points
  - Whether datasets include personal information, aggregate consumer information, or data protected by IP rights.

- Applies retrospectively – must disclose historical training data practices

- Enforceable via California's Unfair Competition Law – with its private right of action.

Beyond California

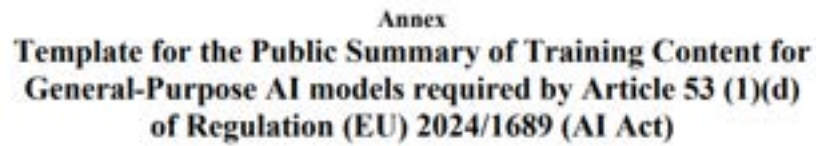

Annex

Template for the Public Summary of Training Content for General-Purpose AI models required by Article 53 (1)(d) of Regulation (EU) 2024/1689 (AI Act)

- EU AI Act — General Purpose AI Models:
  - Must publish summaries of training data
  - Must implement policies complying with EU copyright law
  - For content owners: training data disclosures are now part of your litigation intelligence function

- Connecticut SB5 — Provenance Data:
  - "Covered providers" must embed tamper-resistant provenance data in AI-generated or materially altered content
  - Does not require disclosure of personal information, trade secrets, or confidential information

- Takeaways

Questions?

Daniel Goldberg

Chair, Data Strategy Group
Frankfurt Kurnit Klein + Selz
dgoldberg@fkks.com

Andrew Folks

Associate, Data Strategy Group
Frankfurt Kurnit Klein + Selz
afolks@fkks.com