



**The Journal of Robotics,
Artificial Intelligence & Law**

Editor's Note: It's Still About the Data

Victoria Prussen Spears

Federal Energy Regulatory Commission Issues Order Providing Guidance for
"Co-Locating" Power Plants with Data Centers in PJM Interconnection Footprint

David A. Applebaum, Emil Barth, Diane Gremillion Evans, Mary Franco,
Ryan C. Norfolk, Jay T. Ryan, and Michael A. Yuffee

Artificial Intelligence in Clinical Decision-Making: Regulatory Roadmap and
Reimbursement Strategies

Shawn Maree Bishop, Nathan A. Brown, Kelly M. Cleary, Virgil A. Miller,
Hans Christopher Rickhoff, and Mario Luis Ramirez

Why You Need to Care About AI Bias and How a Bias Audit Can Help You Avoid Danger

Usama Kahf, Chelsea Viola, and David J. Walton

California Appellate Courts Remind Practitioners to Avoid Citing AI Hallucinations
in Legal Briefs

William M. Hensley

AI Platform Risk Assessments: Time for Action

Amy S. Mushahwar, P. Kai Knight, and Tricia Y. Wagner

Federal Communication Commission Adds All Foreign-Made UAS and UAS Critical
Components to Covered List

Jennifer L. Richter, Steven A. Rowings, Sean T. Conway, Virginia Hiner Antypas,
Halie B. Peacher, Sharanya Sriram, and Alexandra M. Van Cleef

A National AI Platform Takes Shape: What Corporate Innovators Need to Know About
the Genesis Mission

Gregory Szewczyk and Harlan Mechling

New AI Regulations Come Into Play with the Texas Responsible Artificial Intelligence
Governance Act

Katherine Franco

Start-Up Corner: Venture Basics: Understanding Protective Provisions

Jim Ryan

- 167 Editor’s Note: It’s Still About the Data**
Victoria Prussen Spears
- 171 Federal Energy Regulatory Commission Issues Order Providing Guidance for “Co-Locating” Power Plants with Data Centers in PJM Interconnection Footprint**
David A. Applebaum, Emil Barth, Diane Gremillion Evans, Mary Franco, Ryan C. Norfolk, Jay T. Ryan, and Michael A. Yuffee
- 181 Artificial Intelligence in Clinical Decision-Making: Regulatory Roadmap and Reimbursement Strategies**
Shawn Maree Bishop, Nathan A. Brown, Kelly M. Cleary, Virgil A. Miller, Hans Christopher Rickhoff, and Mario Luis Ramirez
- 189 Why You Need to Care About AI Bias and How a Bias Audit Can Help You Avoid Danger**
Usama Kahf, Chelsea Viola, and David J. Walton
- 195 California Appellate Courts Remind Practitioners to Avoid Citing AI Hallucinations in Legal Briefs**
William M. Hensley
- 201 AI Platform Risk Assessments: Time for Action**
Amy S. Mushahwar, P. Kai Knight, and Tricia Y. Wagner
- 209 Federal Communication Commission Adds All Foreign-Made UAS and UAS Critical Components to Covered List**
Jennifer L. Richter, Steven A. Rowings, Sean T. Conway, Virginia Hiner Antypas, Halie B. Peacher, Sharanya Sriram, and Alexandra M. Van Cleef
- 215 A National AI Platform Takes Shape: What Corporate Innovators Need to Know About the Genesis Mission**
Gregory Szewczyk and Harlan Mechling
- 221 New AI Regulations Come Into Play with the Texas Responsible Artificial Intelligence Governance Act**
Katherine Franco
- 225 Start-Up Corner: Venture Basics: Understanding Protective Provisions**
Jim Ryan

EDITOR-IN-CHIEF

Steven A. Meyerowitz

President, Meyerowitz Communications Inc.

EDITOR

Victoria Prussen Spears

Senior Vice President, Meyerowitz Communications Inc.

BOARD OF EDITORS

Jennifer A. Johnson

Partner, Covington & Burling LLP

Paul B. Keller

Partner, Allen & Overy LLP

Garry G. Mathiason

Shareholder, Littler Mendelson P.C.

James A. Sherer

Partner, Baker & Hostetler LLP

Elaine D. Solomon

Partner, Blank Rome LLP

Edward J. Walters

Chief Strategy Officer, vLex

John Frank Weaver

Director, McLane Middleton, Professional Association

START-UP COLUMNIST

Jim Ryan

Partner, Morrison & Foerster LLP

THE JOURNAL OF ROBOTICS, ARTIFICIAL INTELLIGENCE & LAW (ISSN 2575-5633 (print) /ISSN 2575-5617 (online) at \$495.00 annually is published six times per year by Full Court Press, a Fastcase, Inc., imprint. Copyright 2026 Fastcase, Inc. No part of this journal may be reproduced in any form—by microfilm, xerography, or otherwise—or incorporated into any information retrieval system without the written permission of the copyright owner. For customer support, please contact Fastcase, Inc., 729 15th Street, NW, Suite 500, Washington, D.C. 20005, 202.999.4777 (phone), or email customer service at support@fastcase.com.

Publishing Staff

Publisher: David Nayer

Production Editor: Sharon D. Ray

Cover Art Design: Juan Bustamante

Cite this publication as:

The Journal of Robotics, Artificial Intelligence & Law (Fastcase)

This publication is sold with the understanding that the publisher is not engaged in rendering legal, accounting, or other professional services. If legal advice or other expert assistance is required, the services of a competent professional should be sought.

Copyright © 2026 Full Court Press, an imprint of Fastcase, Inc.

All Rights Reserved.

A Full Court Press, Fastcase, Inc., Publication

Editorial Office

729 15th Street, NW, Suite 500, Washington, D.C. 20005

<https://www.fastcase.com/>

POSTMASTER: Send address changes to THE JOURNAL OF ROBOTICS, ARTIFICIAL INTELLIGENCE & LAW, 729 15th Street, NW, Suite 500, Washington, D.C. 20005.

Articles and Submissions

Direct editorial inquiries and send material for publication to:

Steven A. Meyerowitz, Editor-in-Chief, Meyerowitz Communications Inc.,
26910 Grand Central Parkway, #18R, Floral Park, NY 11005, smeyerowitz@
meyerowitzcommunications.com, 631.291.5541.

Material for publication is welcomed—articles, decisions, or other items of interest to attorneys and law firms, in-house counsel, corporate compliance officers, government agencies and their counsel, senior business executives, scientists, engineers, and anyone interested in the law governing artificial intelligence and robotics. This publication is designed to be accurate and authoritative, but neither the publisher nor the authors are rendering legal, accounting, or other professional services in this publication. If legal or other expert advice is desired, retain the services of an appropriate professional. The articles and columns reflect only the present considerations and views of the authors and do not necessarily reflect those of the firms or organizations with which they are affiliated, any of the former or present clients of the authors or their firms or organizations, or the editors or publisher.

QUESTIONS ABOUT THIS PUBLICATION?

For questions about the Editorial Content appearing in these volumes or reprint permission, please contact:

David Nayer, Publisher, Full Court Press at david.nayer@clio.com or at
202.999.4777

For questions or Sales and Customer Service:

Customer Service
Available 8 a.m.–8 p.m. Eastern Time
866.773.2782 (phone)
support@fastcase.com (email)

Sales
202.999.4777 (phone)
sales@fastcase.com (email)

ISSN 2575-5633 (print)
ISSN 2575-5617 (online)

AI Platform Risk Assessments: Time for Action

Amy S. Mushahwar, P. Kai Knight, and Tricia Y. Wagner*

In this article, the authors explain why companies that adopt artificial intelligence should perform risk assessments now.

Companies spent 2025 racing to adopt artificial intelligence (AI). The data shows that AI did not just create new risks; it also acted as a high-speed searchlight, exposing the infrastructure gaps many organizations have carried since the late 1990s. Now, we are closing an era of deferred maintenance.

Why Act Now?

Many of you are facing risk now. Your board is asking about AI risk. Your engineers are deploying models faster than Legal can review them. Your vendor contracts do not address who owns training data. And regulators are watching. The recent executive order establishing a national AI policy framework signals that heightened regulatory and enforcement may heat up, even if a preemption battle ensues.

Stakeholders, regulators, and boards now expect visible, defensible action. Building a robust governance framework takes time, so organizations that begin now will be better positioned to meet future requirements. Notably, under California's new mandatory risk framework, AI risk assessment is a required component of enterprise risk assessment, with a compliance deadline of December 31, 2027.

Mitigate Risk and Use the NIST AI RMF as Your Operating Spine

The National Institute of Standards and Technology (NIST) Artificial Intelligence Risk Management Framework (AI RMF)¹ provides a sector-agnostic, defensible structure for AI governance,

and is quickly becoming the industry standard. It offers practical tools, including an implementation playbook and crosswalks to other governance frameworks. These tools enable organizations to align legal, risk, and engineering teams while maintaining traceability from policy to practice.²

Table 1		
Function	What to Implement Now	Evidence to Retain
Govern	Charter an AI risk committee; define roles and responsibilities, thresholds, and escalation	Charter; RACI (Responsible, Accountable, Consulted, and Informed matrixes), meeting minutes, risk appetite statements
Map	Inventory AI use cases and systems architecture, data flows, stakeholders, and potential harms	Inventory, system ownership (individual and joint), data lineage diagrams, impact assessments, risk register prioritized, legal, operational, and engineering alignment
Measure	Define metrics for performance, robustness, bias, privacy, and security	Test plans (pre- and post-deployment), datasets, results, acceptance criteria, signoffs
Manage	Monitor in-production; implement rollback, retraining, and incident handling	Monitoring dashboards, drift alerts, incident logs, change approvals

This mapping enables organizations to speak a common language across legal, risk, and engineering teams and to demonstrate continuous improvement even as regulatory requirements evolve.

Make It Actionable: Three Critical Foundations

Infrastructure Reality: Define Accountability at Every Handoff in Policies, Procedures, and Data Maps

Who owns AI outputs when customer data is trained by engineering, deployed by product, and used for decisions Legal is liable

for? Map data flows, model lineage, and system ownership. Identify who owns data at each stage, from training and fine-tuning through deployment and decision-making, and who has authority to pause or override systems when risks emerge. Policies that cannot be executed in production are not governance; they create risk without a roadmap for execution.

Legal-Engineering Alignment: Test Policies Against System Reality

Can you honor all your data subject access requests in a trained model? Ensure that privacy, deletion, access, and transparency commitments are technically feasible. Can you explain decisions your algorithm makes? Legal teams must understand how systems operate in practice; engineering teams must understand the legal consequences of design choices.

Board-Ready Oversight: Ground Reporting in Infrastructure Reality

Document AI risk appetite, unacceptable uses, and testing standards. Provide quarterly dashboards on high-risk systems, incidents, and regulatory milestones. Board reporting should reflect system operations and risk reality, not just compliance status.

Then, Operationalize Across Your Organization

- *Incident Response.* Update playbooks for AI-specific issues—bias, drift, adversarial events. Define escalation paths and document decisions.
- *Contracts and Third Parties.* Update templates for training data rights, safety/bias/privacy warranties, model change disclosures, and audit rights.
- *State Law Compliance.* Maintain a register of obligations by jurisdiction. Adopt the strictest common denominator for enterprise standards.
- *Tabletop Exercises.* Conduct realistic scenarios that mimic real incidents (e.g., technical partial information, time pressure, and competing priorities). Include Legal, Engineering, Product, Communications, and the leadership team.

Pull actual logging interfaces in the tabletop so you are aware of what logging is available for your most critical AI platforms.

- *Regulatory Monitoring.* Assign responsibility for tracking Department of Justice, Commerce Department, agency rulemaking, and state updates.

Long-Term Planning: Phased Approach with Time Frames

Effective AI governance requires a phased approach. Organizations should begin by mapping AI use and establishing governance and documentation, then implement testing, contractual, and technical controls, and ultimately focus on ongoing monitoring, reporting, and transparency to ensure responsible, sustainable oversight.

Table 2		
Phase	Time Frame	Focus
Phase 1	0-3 months	Mapping, governance, ownership, documentation
Phase 2	3-9 months	Testing, contracts, technical controls
Phase 3	9-18 months+	Monitoring, reporting, transparency

Phase 1: Governance and Documentation (0-3 months)

- Map AI usage;
- Assign accountable owners for AI risk and compliance;
- Form cross-functional and diverse review groups (Legal, risk, information technology, business);
- Create a system of record for all AI systems in use;
- Update incident response plans for AI-specific risks; and
- Assess and update policies.

Rationale. These foundational steps establish oversight and visibility. They can be launched immediately and should be completed quickly to demonstrate good faith to regulators and stakeholders.

Phase 2: Strengthen Testing and Controls (3-9 months)

- Broaden testing protocols (e.g., for subgroup fairness, privacy, security);
- Revise contracts and agreements for AI-specific obligations (training data rights, audit rights, model change disclosures);
- Implement technical controls for monitoring, rollback, and retraining; and
- Schedule the first tabletop for AI response.

Rationale. This phase builds on the governance foundation. It requires coordination across teams and may involve vendor negotiations and technical upgrades. Regulators increasingly expect demonstrable progress within the first year, with an improving compliance narrative over time.

Phase 3: Continuous Monitoring and Reporting (9-18 months and ongoing)

- Shift to ongoing monitoring (alerts, dashboards, drift detection);
- Implement quarterly reporting to boards and leadership team; and
- Prepare public summaries or model cards as needed for transparency.

Rationale. Continuous monitoring is an ongoing commitment. Initial systems should be in place within 12 to 18 months, with regular updates and improvements as the regulatory landscape evolves.

What Regulators Will Ask for by Priority

Regulators do not expect perfection; they expect visible progress and a credible improvement narrative. Here is what to have ready on a prioritized basis because full compliance is not feasible immediately.

Have Now (Foundation)

- AI system inventory with risk tiers and ownership;
- Updated incident response plans for AI-specific risks; and
- Charter for AI governance committee.

Build in Year 1 (Demonstrating Progress)

- AI policy and standards; iterate—it will not be comprehensive initially;
- Written protocols for testing and validation, but this may need to be done sooner rather than later, particularly where systems affect employment, housing, or vulnerable populations such as children or seniors;
- Vendor diligence questionnaires and updated contracts;
- Impact sector-specific assessment templates for the Health Insurance Portability and Accountability Act, the Gramm-Leach-Bliley Act, Family Educational Rights and Privacy Act, Customer Proprietary Network Information, NIST/Cybersecurity Maturity Model Certification; and
- Minutes from risk governance committees and training records.

Maintain Ongoing (Operational Maturity)

- Model cards and data sheets;³
- Change logs and approval records;
- Compliance mapping for state and federal laws (living document); and
- Monitoring dashboards and drift alerts.

Key Takeaway

Regulatory uncertainty is real, but defensible steps exist. Use the NIST AI RMF as your foundation, stay compliant with state laws, monitor federal updates, and implement ongoing oversight. Acting now reduces enforcement risk, demonstrates leadership in responsible AI practices, and enables prepared, measured judgment if—and when—an AI incident occurs.

We want your team to avoid a scenario such as:

discovering your customer service AI was making eligibility decisions it wasn't authorized to make. Legal thought they'd prohibited automated decisioning. Engineering thought the model was advisory-only. Product thought they'd disclosed it. Nobody had mapped who owned the output or who could stop the model. As a result, you fumble around for hours trying to find out who has access to shut down the model.

AI is moving quickly, and operational documentation will ensure you have sufficient knowledge to act when necessary.

Organizations building AI governance programs in 2026 should begin with infrastructure mapping and governance chartering. Early action positions you ahead of evolving requirements and ensures your AI tools are reliable and compliant.

In Summary

- *Act Now.* Boards and regulators expect visible progress on AI governance.
- *California Deadline.* Risk Assessments including certain high-risk AI usage due by December 31, 2027.
- *Use NIST AI RMF.* Adopt a defensible, sector-agnostic framework.
- *Start with Infrastructure Mapping.* Define ownership and accountability early.

Notes

* The authors, attorneys with Lowenstein Sandler LLP, may be contacted at amushahwar@lowenstein.com, kknights@lowenstein.com, and twagner@lowenstein.com, respectively.

1. <https://airc.nist.gov/airmf-resources/airmf/>.

2. The playbook is available at <https://airc.nist.gov/airmf-resources/playbook/>; crosswalks at <https://airc.nist.gov/airmf-resources/crosswalks/>.

3. Examples of model cards in action from Hugging Face, at <https://huggingface.co/docs/hub/model-cards>; see a helpful article on model card standardization at Cornell University, Model Cards for Model Reporting, available at <https://arxiv.org/abs/1810.03993>.