



---

**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](mailto: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)

# 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\*

*In this article, the authors explain how employers can avoid introducing artificial intelligence (AI) bias into their everyday decision-making by understanding what AI bias is, how it arises, and why employers should consider a bias audit if they use AI-driven decision-making tools.*

---

As employers increasingly use artificial intelligence (AI) to help screen job applicants, evaluate performance, forecast risk, and support everyday decision-making, it is critical you ensure your organization does not accidentally introduce “AI bias” into the mix. A slate of recent lawsuits alleging that the use of AI tools may disproportionately affect certain groups should serve as a warning, even if you believe your systems are facially neutral and used in good faith. Now, employers should make sure they understand what AI bias is, how it arises, and why to consider a bias audit if they use AI-driven decision-making tools.

## From Data to Decisions: Inputs, Features, and Weights

---

It is helpful to have a quick crash course in AI in order to understand how it works and how it can inadvertently lead to biased recommendations or decisions.

- AI systems are developed and trained with *data inputs*, also called “training data.” This could include structured data such as resumes, performance metrics, or credit history. It could also include unstructured data such as written text, video interviews, or audio recordings.
- From these inputs, AI developers select *features*, meaning the specific variables or characteristics the system will evaluate when making a decision or recommendation. For

example, features of an AI system evaluating resumes may designate education level or prior job titles as features to consider when evaluating resumes.

- The system then assigns *weights* to those features, which determine how much importance each feature has in producing the final output. For example, a system may give greater weight to education pedigree than years of experience, or prioritize uninterrupted work history over skills-based assessments.

## What Do We Mean by “Bias”?

---

Bias refers to a systematic tendency to favor certain outcomes, characteristics, or groups over others. Bias may be intentional or unintentional, explicit or implicit. Under the legal theory of *disparate impact*, liability can arise even in the absence of discriminatory intent when a facially neutral practice produces outcomes that disproportionately affect individuals based on protected characteristics such as race, sex, age, disability, or other protected statuses.

AI systems are shaped by the data on which they are trained. When that data reflects historical bias, structural inequities, design flaws, or incomplete information, the system may replicate or amplify those patterns. As a result, AI tools can unintentionally replicate existing inequities at scale.

Even when protected characteristics are excluded from the data inputs given to an AI system, certain features may function as *proxies* for protected traits, such as zip code correlating with race, or employment gaps correlating with disability or caregiving responsibilities. Outcomes are shaped by both the selection of features and the weight assigned to them, meaning these design choices can materially influence results and potentially create disparities at scale.

## Common Types of AI Tools Used by Employers and Businesses

---

AI appears in many familiar forms, often without being labeled as such.

- *Predictive Tools.* Predictive analytics uses historical data to identify patterns and forecast future outcomes, such as sales performance or creditworthiness. Since these tools rely on past decisions, they may reinforce existing disparities if historical practices were biased.
  - Example: An insurer uses predictive analytics to estimate the likelihood that a policyholder will file a future claim and to inform premium pricing based on prior claims data.
- *Machine Learning Systems.* Machine learning models learn patterns from large datasets without fixed decision rules. During training, the model continuously adjusts the weights assigned to different features based on prior outcomes. While this adaptability can improve accuracy, it can also make bias harder to identify, particularly as models evolve over time.
  - Example: A bank uses a machine learning model to assess loan applications by learning from historical lending data and adjusting the weight assigned to factors such as credit history and repayment behavior.
- *Scoring, Ranking, and Recommendation Tools.* Many AI systems generate scores, rankings, or recommendations that inform or influence human decisions, such as applicant rankings or performance scores. Even when humans remain involved, there is a risk of over-reliance on automated outputs, which can reduce meaningful oversight.
  - Example: A resume-screening tool ranks applicants by learning which candidates advance in the hiring process and adjusting its evaluation criteria based on prior hiring outcomes.
- *Language-Based and Generative Tools.* Some AI systems analyze or generate language, including resume-screening tools, chatbots, and performance-summary tools. Since these systems are trained on large volumes of text, they can replicate patterns and assumptions present in the training data.
  - Example: An AI system generates automated email or chat responses to customer inquiries based on patterns learned from prior communications.

## Where AI Bias Creates Legal Risk: Disparate Impact and Detection Challenges

---

One of the most significant challenges with AI bias is that it can be difficult to detect. Many AI systems operate as “black boxes,” making it hard to understand how inputs translate into outcomes. Without deliberate testing and documentation, biased results may go unnoticed until regulatory scrutiny or litigation arises.

### The 80/20 Rule

A potential measure for disparate impact when AI systems are involved is often evaluated using the 80/20 rule. Under this framework, a selection rate for a protected group that is less than 80 percent of the rate for the most favored group may indicate potential adverse impact. The 80/20 rule is not a definitive test of discrimination. Instead, it functions as a screening mechanism or warning indicator that may warrant closer review of a particular practice or decision-making process.

For example, if 60 percent of male applicants pass a screening assessment and only 45 percent of female applicants do, the resulting ratio of 75 percent falls below the 80 percent threshold. This outcome does not establish discrimination on its own, but it may signal potential bias and trigger further analysis.

## Why AI Bias Matters for Employers

---

AI bias matters because legal liability can arise even in the absence of discriminatory intent. Employers and businesses may be held responsible for practices that disproportionately impact protected groups when those outcomes are not job-related and consistent with business necessity.

As AI tools play an increasing role in employment decisions, courts are scrutinizing how these systems operate and whether they contribute to discriminatory results.

- One example is *Mobley v. Workday*, a class action pending in California federal court, in which a job applicant alleges that Workday’s AI-based screening tools systematically rejected him across more than 100 applications.

- Similarly, in *Harper v. Sirius XM*, pending in Michigan federal court, an applicant alleges that the employer relied on an AI-powered applicant tracking system that embedded historical bias by using data points functioning as proxies for race, resulting in his candidacy being downgraded and eliminated before advancing in the hiring process.

## What Employers Can Do Now: Consider an AI Bias Audit

---

As AI-driven tools become more embedded in employment decision-making, employers should take proactive steps to assess and mitigate bias using a defensible, structured approach. Employers can evaluate risk and implement practical safeguards by:

- Identifying where AI tools are used across the employment life cycle, including recruiting, hiring, onboarding, performance management, staffing and assignment decisions, employee relations, retention, and termination, to pinpoint where automated decision-making may create risk.
- Conducting bias audits and compliance reviews of third-party AI vendors, including reviewing vendor-provided bias audits or documentation, evaluating whether a tool qualifies as an automated decision tool under current and emerging laws, and advising on practical risk-mitigation strategies.
- Auditing internally developed or custom-built AI tools, including statistical testing for disparate impact across protected categories where data are available, explainable AI-based root-cause analysis, and recommendations to reduce or correct identified bias.
- Providing privileged legal assessments and regulator-ready documentation, including multi-jurisdictional compliance analyses, guidance on disclosure obligations, and summaries suitable for internal governance or external review.
- Establishing AI monitoring and governance frameworks, including regular bias audits, AI-assisted monitoring, regulatory updates, compliance workshops, and policy refresh guidance to address evolving legal requirements over time.

## Note

---

\* The authors, attorneys with Fisher Phillips, may be contacted at [ukahf@fisherphillips.com](mailto:ukahf@fisherphillips.com), [cviola@fisherphillips.com](mailto:cviola@fisherphillips.com), and [dwalton@fisherphillips.com](mailto:dwalton@fisherphillips.com), respectively.