



---

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

---

Editor's Note: AI-Generated Arbitration Awards  
Victoria Prussen Spears

**The Silicon Arbitrator: AI-Generated Arbitration Awards and the Federal Arbitration Act—Part I**

David L. Evans

Executive Preemption and Artificial Intelligence: President Trump's Bet on  
Deregulatory Federalism  
Bepi Pezzulli

AI at the Arsenal Gate: How the Pentagon's New AI Strategy Rewrites Defence  
Supply Chain Rules for Contractors  
Nicolas Groffman

Courts Grapple with Privilege Implications of AI  
Angela L. Dunning, Jonathan S. Kolodner, David A. Last, Rahul Mukhi,  
Jennifer Kennedy Park, Arminda B. Bepko, and Nathaniel P. Reynolds

Instacart Agrees to Settlement in Federal Trade Commission Lawsuit over Alleged  
Deceptive Marketing Tactics and AI-Driven Pricing  
Bruce D. Sokler, Alexander Hecht, Erek L. Barron, Christian Tamotsu Fjeld, and  
Aliza Inam

Start-Up Corner: Venture Basics: Understanding Anti-Dilution  
Jim Ryan

- 235 Editor’s Note: AI-Generated Arbitration Awards**  
Victoria Prussen Spears
- 239 The Silicon Arbiter: AI-Generated Arbitration Awards and the Federal Arbitration Act—Part I**  
David L. Evans
- 257 Executive Preemption and Artificial Intelligence: President Trump’s Bet on Deregulatory Federalism**  
Bepi Pezzulli
- 279 AI at the Arsenal Gate: How the Pentagon’s New AI Strategy Rewrites Defence Supply Chain Rules for Contractors**  
Nicolas Groffman
- 295 Courts Grapple with Privilege Implications of AI**  
Angela L. Dunning, Jonathan S. Kolodner, David A. Last, Rahul Mukhi, Jennifer Kennedy Park, Arminda B. Bepko, and Nathaniel P. Reynolds
- 303 Instacart Agrees to Settlement in Federal Trade Commission Lawsuit over Alleged Deceptive Marketing Tactics and AI-Driven Pricing**  
Bruce D. Sokler, Alexander Hecht, Erek L. Barron, Christian Tamotsu Fjeld, and Aliza Inam
- 309 Start-Up Corner: Venture Basics: Understanding Anti-Dilution**  
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)

# The Silicon Arbiter: AI-Generated Arbitration Awards and the Federal Arbitration Act—Part I

David L. Evans\*

*The emergence of artificial intelligence (AI) systems has made urgent a question that would have seemed absurd to the drafters of the Federal Arbitration Act (FAA): whether algorithms might generate binding arbitral awards without human involvement. This two-part article concludes that, under the FAA as currently written, such awards cannot be enforced. Yet this article does not end with a negative prognosis for AI's place in arbitration. Instead, it makes suggestions for legislative reform addressing informed consent, explainability, bias mitigation, accountability, and meaningful judicial review.*

*After a brief introduction, this first part of this two-part article frames the enforceability question precisely, distinguishing pure AI arbitration from AI-assisted human arbitration and explaining why this article focuses on judicial enforcement rather than initial contractual validity; establishes the textualist interpretive framework, explaining why the Supreme Court's recent decision in *Loper Bright Enterprises v. Raimondo* controls and what that methodology requires; applies the major questions doctrine from *West Virginia v. EPA*, demonstrating that recognizing AI-generated awards would constitute precisely the kind of major question requiring explicit congressional authorization; and applies textualist analysis to the FAA, demonstrating through dictionary definitions, statutory structure, and Section 10's vacatur grounds that the statute presupposes human arbitrators.*

*The conclusion of this article, to be published in the next issue of *The Journal of Robotics, Artificial Intelligence & Law*, will address constitutional concerns under both the Due Process Clause and Article III; respond to counterarguments based on party autonomy, engaging particularly with Professor Szalai's recent purposivist arguments; distinguish AI-assisted arbitration from AI-generated awards and provide guidance on where the line falls; propose a legislative framework for those who would chart a different course; and offer a conclusion.*

---

## Introduction

---

On November 3, 2025, the American Arbitration Association-  
International Centre for Dispute Resolution (together, AAA)

launched what it called an “AI Arbitrator” for documents-only construction disputes.<sup>1</sup> The system, trained on over 1,500 construction arbitration awards, generates draft decisions that human arbitrators review before issuing final awards. Significantly, the AAA is using a “human-in-the-loop” framework, perhaps cognizant that questions persist about whether a purely artificial intelligence (AI)-generated award could be enforced. Start-ups have begun marketing AI dispute resolution products that purport to render binding decisions without any meaningful human oversight.<sup>2</sup> These developments raise a fundamental legal question that deserves rigorous analysis: Would awards issued solely by AI arbitrators be confirmable under the Federal Arbitration Act?

The Federal Arbitration Act (FAA)<sup>3</sup> was enacted in 1925, when the very notion of AI resided solely in the realm of science fiction. Congress could not have contemplated, nor did it purport to address, the possibility that parties might one day submit their disputes to algorithmic resolution rather than human judgment. The question now presented is: When parties agree that an AI system shall resolve their dispute and render a binding decision, does the resulting determination constitute an “award” that courts may confirm under the FAA?

This article posits that the answer is no. Under the FAA as currently written, AI-generated arbitral awards cannot be enforced, regardless of party consent. This conclusion follows from principles of statutory interpretation that the Supreme Court has recently reinforced. The FAA’s use of “arbitrator” unambiguously contemplates human decision-makers. Party autonomy, however foundational to arbitration, operates within statutory bounds and cannot expand “arbitrator” to include entities Congress never contemplated.

The textual obstacles are not merely formal. The FAA’s vacatur provisions create what might be called a “vacatur gap”: a jurisprudential void where algorithmic errors, hallucinations, and systematic biases would be immune from the statutory checks Congress designed to ensure arbitral integrity. Section 10 permits courts to vacate awards procured by “corruption” or “fraud,” rendered by arbitrators exhibiting “evident partiality,” or reflecting “misconduct.” Each ground presupposes a human decision-maker capable of moral failure. When the arbitrator is a machine, these safeguards become nullities.

Alternative dispute resolution (ADR) practitioners hold different perspectives on the use of AI to decide cases. Some see

enormous benefits in the new technology. Others express profound concerns about replacing human judgment with algorithmic decision-making. But regardless of one's view on the desirability of AI arbitrators, the legal question of enforceability must be answered through statutory interpretation, not technological enthusiasm or anxiety.

The enforceability question has immediate practical significance. Without clear legal authorization, parties who agree to AI arbitration face uncertainty about whether resulting "awards" can be confirmed in court under the FAA. This uncertainty undermines arbitration's core advantage: providing economical, efficient and enforceable resolution of disputes. This article addresses that uncertainty by providing the statutory analysis the question demands.

But this article's negative conclusion is not fixed. The technological capabilities that make AI arbitration conceivable will continue to advance. Commercial pressures toward faster, cheaper dispute resolution will intensify. If society determines that AI-generated awards should be enforceable, the appropriate path is legislative action rather than judicial reinterpretation. Accordingly, this article proposes a framework for what such legislation should address.

This two-part article proceeds as follows. Here, the next section frames the enforceability question precisely, distinguishing pure AI arbitration from AI-assisted human arbitration and explaining why this article focuses on judicial enforcement rather than initial contractual validity. The third section establishes the textualist interpretive framework, explaining why the Supreme Court's recent decision in *Loper Bright Enterprises v. Raimondo* controls and what that methodology requires. The fourth section applies the major questions doctrine from *West Virginia v. EPA*, demonstrating that recognizing AI-generated awards would constitute precisely the kind of major question requiring explicit congressional authorization. The fifth section applies textualist analysis to the FAA, demonstrating through dictionary definitions, statutory structure, and Section 10's vacatur grounds that the statute presupposes human arbitrators.

The conclusion of this article will contain the sixth section, which addresses constitutional concerns under both the Due Process Clause and Article III; the seventh section, which responds to counterarguments based on party autonomy, engaging particularly with Professor Szalai's recent purposivist arguments; the eighth section, which distinguishes AI-assisted arbitration from

AI-generated awards and provides guidance on where the line falls; the ninth section, which proposes a legislative framework for those who would chart a different course; and the tenth section, which contains the conclusion.<sup>4</sup>

## Defining the Problem

---

Can parties to a dispute agree to have an AI system decide the dispute for them? In a trivial sense, of course the answer is yes. If A and B agree to flip a coin to resolve their dispute, and both honor that agreement, the dispute is resolved. No legal enforcement mechanism is necessary.

If A and B agree to have an AI system decide the dispute for them, and A is dissatisfied with the outcome or otherwise does not voluntarily abide by the AI system's decision, can B obtain a court order enforcing that decision? That is the question this article addresses. It does not examine whether parties can validly agree to AI arbitration as a matter of contract formation or whether such agreements are unconscionable or violate public policy as a threshold matter. It is assumed that parties have validly agreed to AI arbitration and the focus is solely on whether the resulting award can be judicially enforced under the FAA.

Before examining enforceability, it is essential to clarify what is meant by "AI arbitration," as current practice reveals a spectrum of AI involvement. At one end, human arbitrators routinely use AI tools for document review, legal research, or drafting assistance, much as they might use legal databases or word processing software. These tools enhance efficiency, but the human arbitrator retains complete decisional control. The AAA institutional AI guidelines contemplate this model and stress that AI "shall not replace the arbitrator's analysis of the facts, law, and evidence."

In the middle of the spectrum lie "hybrid" or "human-in-the-loop" models where AI systems draft proposed awards that human arbitrators then review, revise, and formally issue. The AAA's recently announced AI Arbitrator operates precisely this way: the AI analyzes submissions and generates a draft award, but a trained human arbitrator reviews, corrects if necessary, and issues the final award. AAA materials emphasize that "no case is decided solely by AI" and the award "remains the human arbitrator's award, not the AI's."

At the far end sits pure AI arbitration, in which parties designate an AI system as the decision-maker without meaningful human adjudicatory oversight. The system receives submissions and autonomously generates a binding award.

This article's enforceability analysis focuses primarily on the third model—pure AI arbitration—and argues that such awards cannot be confirmed under the FAA. Hybrid models present closer questions that are addressed in the eighth section, where we provide guidance on distinguishing permissible AI assistance from impermissible AI replacement. The distinction matters because statutory interpretation depends on whether a human decision-maker with legal responsibility stands behind the award.

This spectrum defines the central question with precision: If parties validly agree to pure AI arbitration through clear contractual language, and the AI system issues an award, can that award be confirmed under FAA Section 9? This article focuses on federal law under the FAA; the New York Convention on the Recognition and Enforcement of Foreign Arbitral Awards presents different questions that merit separate analysis.

## The Framework for Statutory Interpretation

---

In *Loper Bright Enterprises v. Raimondo*, the Supreme Court overruled *Chevron U.S.A. Inc. v. Natural Resources Defense Council, Inc.*, ending decades of judicial deference to agency interpretations of ambiguous statutes.<sup>5</sup> The Court held that courts must exercise independent judgment in determining statutory meaning, using traditional tools of statutory construction. “Courts must exercise their independent judgment in deciding whether an agency has acted within its statutory authority,” the Court declared. While *Loper Bright* addressed administrative law, its textualist methodology applies equally to interpreting the FAA.

As Chief Justice John Roberts wrote for the majority:

[Ambiguous] statutes, no matter how impenetrable, do—in fact, must—have a single, best meaning. That is the whole point of having written statutes; every statute's meaning is fixed at the time of enactment.<sup>6</sup>

The implications extend beyond administrative law. *Loper Bright* reflects a commitment to textualism and to the judiciary's

responsibility to “say what the law is.”<sup>7</sup> If courts may not defer to agencies on statutory meaning, they certainly may not defer to private parties’ contractual arrangements that conflict with statutory requirements. The FAA’s terms must be given their ordinary meaning as courts determine it, not as stretched by litigants seeking novel procedures.

This textualist approach is particularly significant for our analysis. The FAA was enacted in 1925, long before anyone conceived of artificial intelligence as we understand it today. A purposivist approach might suggest that courts should adapt the statute to accommodate technological progress. But *Loper Bright* forecloses that move. The statute’s meaning was “fixed at the time of enactment,” and courts cannot update it to reflect modern innovations, no matter how beneficial. If the FAA does not authorize confirmation of AI-generated awards based on its 1925 meaning, only Congress can change that through amendment.

The Supreme Court has consistently emphasized that determining ordinary meaning requires attention to context. In *Taniguchi v. Kan Pacific Saipan, Ltd.*, the Court examined whether the term “interpreter” in a fee-shifting statute included document translators.<sup>8</sup> The Court consulted dictionaries contemporary to the statute’s enactment, examined how the term was used in nearby statutory provisions, and considered the statute’s structure. This contextual approach to ordinary meaning provides the template for analyzing “arbitrator” in the FAA.

Similarly, in *Dubin v. United States*, the Court examined whether “means of identification” in the identity theft statute included the means itself or only the underlying identity.<sup>9</sup> The Court emphasized that statutory terms must be read in context, with attention to how Congress used related terms throughout the statute. This structural analysis applies with equal force to the FAA’s repeated use of “arbitrator” and related terms throughout its provisions.

The Supreme Court’s emphasis on context extends beyond nearby words to the statute’s overall structure. In *FDA v. Brown & Williamson Tobacco Corp.*, the Court held that statutory interpretation requires consideration of “the specific context in which that language is used, and the broader context of the statute as a whole.”<sup>10</sup> When Congress uses terms that describe human characteristics or actions, courts presume Congress meant to refer to humans unless the statute clearly indicates otherwise.

The Court's arbitration jurisprudence has consistently attended to statutory text. In *Epic Systems Corp. v. Lewis*, Justice Neil Gorsuch's majority opinion emphasized that the FAA "requires courts to enforce arbitration agreements according to their terms."<sup>11</sup> But this principle presupposes that the underlying procedure qualifies as arbitration within the statutory meaning. An agreement to resolve disputes through coin flips or trial by combat would not be enforceable under the FAA, however clearly expressed. The question is whether AI-generated decision-making falls within the statutory definition, and *Loper Bright* instructs that courts must answer this question through independent interpretation of the text.

## The Major Questions Doctrine

---

The major questions doctrine, articulated most fully in *West Virginia v. EPA*, provides additional grounds for concluding that the FAA does not authorize enforcement of AI-generated awards.<sup>12</sup> The Court held that courts should not presume Congress intended to delegate decisions of "vast economic and political significance" absent clear authorization.<sup>13</sup> "[E]xtraordinary grants of regulatory authority" require "clear congressional authorization."<sup>14</sup> This reflects a presumption against hidden elephants in statutory mouseholes and insistence that Congress speak clearly when effecting major policy changes.<sup>15</sup> While the major questions doctrine traditionally constrains administrative agencies, its underlying logic—that Congress speaks clearly when assigning powers of vast economic and political significance—applies with equal force to judicial interpretation of the FAA. The Court's skepticism toward implicit delegations of authority is not limited to the executive branch; it reflects a broader insistence that major policy shifts require legislative enactment.

Recognizing AI-generated arbitral awards would constitute precisely the kind of major question requiring explicit congressional authorization. It would transform private dispute resolution, affecting billions of dollars in commercial transactions. It would raise profound questions about due process, accountability, and the meaning of adjudication. It would require courts to develop entirely new frameworks for reviewing algorithmic decisions. All this would rest on statutory language that, on its face, contemplates human arbitrators exercising human judgment.

The FAA's silence on AI-generated awards cannot be read as implicit authorization for such a fundamental departure from the statute's design. Congress in 1925 could not have contemplated algorithmic adjudication, and its use of "arbitrator" must be understood against the background of human decision-making that was the only form of arbitration then known. To read the statute as authorizing AI-generated awards would locate in congressional silence a delegation of vast significance, precisely the interpretive move that *West Virginia v. EPA* forbids. Such a transformation requires explicit legislative authorization, not judicial inference from statutory silence.

## The Statutory Presumption of Human Arbitrators

---

The FAA's confirmation provision appears deceptively simple. Section 9 authorizes courts to confirm "the award made pursuant to the arbitration" when parties have agreed that judgment shall be entered upon the award. But this simple language raises fundamental questions when applied to AI-generated decisions: What constitutes "the award"? What does "arbitration" mean? And crucially, what is an "arbitrator"?

These questions require examining the FAA's text, structure, and context to determine what Congress understood these terms to mean in 1925. The analysis proceeds in four steps: examining dictionary definitions contemporary to the statute's enactment; analyzing the statute's structural use of arbitrator-related terms; confronting Section 10's vacatur provisions, which reveal most decisively that the FAA presupposes human arbitrators; and examining the nature of arbitral judgment itself.

## The Plain Meaning of "Arbitrator" in 1925

When Congress used the term "arbitrator" in 1925, it adopted a term with settled legal meaning. *Black's Law Dictionary* (2d ed. 1910) defined "arbitrator" as "a private, disinterested person, chosen by the parties to a disputed question, for the purpose of hearing their contention, and giving judgment between them; to whose decision (award) the litigants submit themselves either voluntarily, or, in some cases, compulsorily, by order of a court."<sup>16</sup>

This definition reveals several assumptions about arbitrators that current AI systems cannot satisfy. First, the arbitrator must be “a person.” Second, the arbitrator must be “disinterested,” capable of having interests that might conflict with the parties’ and of making ethical judgments about those conflicts. Third, the arbitrator engages in distinctly cognitive processes: “hearing their contention” and “giving judgment,” verbs that connote human understanding, deliberation, and reasoned decision-making.

*Webster’s New International Dictionary* (1913) similarly defined “arbitrator” as “a person, or one of two or more persons, chosen by parties in controversy, to determine their differences.”<sup>17</sup> Again, the definition explicitly requires a “person” and contemplates the exercise of determination, an inherently mental process involving weighing competing arguments, evaluating evidence, and reaching reasoned conclusions.

*Bouvier’s Law Dictionary* (1914) described the arbitrator’s role in terms that unmistakably contemplate human judgment: “A private extraordinary judge chosen by the parties who have a matter in dispute, invested with power to decide the same. Arbitrators are so called because they have generally an arbitrary power; for if their authority is restrained . . . they are not, in strictness of language, arbitrators, but referees.”<sup>18</sup> This definition’s reference to “arbitrary power” assumes a decision-maker capable of exercising discretion, a quintessentially human capacity. The word derives from the Latin *arbiter*, meaning one who goes to see, a witness, a judge.

Unsurprisingly, none of these definitions remotely contemplates an algorithmic decision-maker. To the contrary, each emphasizes precisely human capacities: discretion, judgment, and the ability to hear and determine disputes through reasoned analysis. When Congress used the term “arbitrator” in 1925, it adopted these settled meanings. The statute’s text must be interpreted accordingly.

## Structural Evidence: The FAA’s References to Human Arbitrators

The FAA’s structure reinforces that arbitrators must be human. Section 5 authorizes courts to “designate and appoint an arbitrator or arbitrators” when parties fail to agree on selection procedures.<sup>19</sup> This appointment power assumes arbitrators who can be identified, contacted, and formally designated. Section 7 provides that

“arbitrators . . . may summon in writing any person to attend before them or any of them as a witness.”<sup>20</sup> Advocates and opponents of AI decision-making in arbitration can disagree about whether an AI system can understand testimony, evaluate credibility, or exercise sound judgment. But as a textual matter, the statutory requirement that witnesses be summoned to “attend before” the arbitrators points clearly in the direction of a human arbitrator. A witness cannot really “attend before” a computer program. The legal authority to issue summonses contemplates arbitrators who understand testimonial relevance, can evaluate credibility, and exercise judgment about what evidence to seek.

Section 9’s requirement that awards be “in writing” and “signed by the arbitrators” similarly presupposes human decision-makers.<sup>21</sup> While AI systems can generate text, the concept of a signature traditionally involves personal authentication of a document’s contents by someone who takes responsibility for them. An AI system cannot “sign” in this meaningful sense because it cannot take moral or legal responsibility for its outputs.

The Dictionary Act’s definition of “person” does not resolve this question in favor of AI arbitrators. While the Act defines “person” to include “corporations, companies, associations, firms, partnerships, societies, and joint stock companies, as well as individuals,”<sup>22</sup> the Supreme Court has explained that this definition applies only when “the context indicates” such an expansive reading is appropriate.<sup>23</sup> Context here indicates otherwise.<sup>24</sup>

One might argue that if corporations qualify as “persons,” AI systems should too. But this analogy fails. Every entity the Dictionary Act includes is a legal person capable of acting through human agents.<sup>25</sup> Corporations serve as arbitrators only by designating human representatives who conduct proceedings and render decisions. The corporation provides institutional structure; humans provide judgment. An AI system, by contrast, would render the decision without human agency.<sup>26</sup> The Dictionary Act’s inclusion of corporations thus does not support extending “person” to include AI systems that operate without human cognition.

While the Court does not generally look to comparative law when construing federal statutes, the conclusion that arbitrators must be human finds support in comparative law. Several civil law jurisdictions have codified what the FAA’s text implies. Portugal’s Voluntary Arbitration Act mandates that “[a]n arbitrator must be a natural person” with full legal capacity.<sup>27</sup> The Netherlands

Arbitration Act contains the same requirement: arbitrators “must have legal capacity and be natural persons.”<sup>28</sup> France requires natural person arbitrators in domestic arbitration, though this restriction does not apply to international arbitration.<sup>29</sup> A recent comparative survey of arbitration statutes across 11 jurisdictions concluded that none “clearly and unequivocally provide for a robot to be used as an arbitrator.”<sup>30</sup> While the FAA does not contain equally explicit language, its use of human-referent pronouns in Section 5 and its vacatur grounds presupposing moral agency place it functionally in the same category. The international consensus reinforces a reading of the FAA that aligns with the global understanding of what arbitration requires.

### **The Section 10 Problem: Human-Centric Vacatur Grounds**

Section 10 provides decisive textual evidence that the FAA presupposes human arbitrators. The statute identifies four grounds for vacating arbitral awards, each premised on distinctly human characteristics or capacities.<sup>31</sup> When examined closely, these provisions reveal that Congress assumed arbitrators would be persons capable of corruption, fraud, partiality, misconduct, and exceeding delegated authority, all concepts presupposing human agency, judgment, and moral capacity.

#### *Section 10(a)(1): Corruption, Fraud, or Undue Means*

Section 10(a)(1) permits vacatur “where the award was procured by corruption, fraud, or undue means.”<sup>32</sup> Each term in this provision presupposes human moral agency. “Corruption” requires intentional wrongdoing by a person capable of understanding ethical obligations and choosing to violate them. “Fraud” involves deliberate deception, an act requiring the intent to mislead. Federal courts have interpreted “undue means” to require conduct that is “immoral if not illegal” and “tantamount to corruption.”<sup>33</sup>

These terms cannot coherently apply to AI systems. An algorithm cannot be “corrupt” because it cannot understand or violate ethical duties. It cannot commit “fraud” because it cannot form the intent to deceive. It cannot employ “undue means” because it cannot engage in the kind of intentional misconduct the phrase contemplates.<sup>34</sup> When an AI system produces a flawed output, this reflects programming errors or data problems, not moral turpitude. The

statute's language assumes a decision-maker capable of choosing between right and wrong conduct, and current AI systems possess no such capacity.

### *Section 10(a)(2): Evident Partiality or Corruption*

Section 10(a)(2) permits vacatur “where there was evident partiality or corruption in the arbitrators, or either of them.” “Partiality” means favoring one party over another based on relationships, interests, or prejudices. An arbitrator exhibits partiality when personal relationships, financial interests, or preexisting opinions cause them to favor one party unfairly. While AI systems can produce skewed results reflecting biased training data, this is not “partiality” in any ordinary sense. The statute contemplates vacatur based on the arbitrator's mental state, a concept inapplicable to algorithms.

The Supreme Court established the modern framework for evident partiality in *Commonwealth Coatings Corp. v. Continental Casualty Co.*, holding that arbitrators must disclose “any dealings that might create an impression of possible bias.”<sup>35</sup> The Court emphasized that arbitrators are held to higher standards than judges because parties have limited recourse against arbitrators. This standard presupposes arbitrators capable of having interests, relationships, and preferences. An AI system has no financial interests to disclose, no professional relationships creating obligations, and no personal preferences favoring particular outcomes.

Federal courts addressing partiality claims examine whether arbitrators deliberately tilted proceedings or decisions to favor particular parties.<sup>36</sup> The circuit courts apply tests asking whether “a reasonable person would have to conclude that an arbitrator was partial to one party.”<sup>37</sup> These standards assume a decision-maker capable of partiality in the psychological sense. An algorithm cannot be partial; it can only produce outputs reflecting its training data and architecture.

The conceptual mismatch between Section 10(a)(2)'s language and AI systems creates an enforcement gap. The provision exists to protect parties from biased decision-makers. Yet when applied to AI systems, it provides no meaningful protection because algorithmic bias falls outside the statute's framework. A party might prove that an AI arbitrator's training data systematically disadvantages certain claimants, but this showing would not establish “evident partiality” as courts have interpreted that term.

### *Section 10(a)(3): Misconduct and Misbehavior*

Section 10(a)(3) permits vacatur where “the arbitrators were guilty of misconduct in refusing to postpone the hearing, upon sufficient cause shown, or in refusing to hear evidence pertinent and material to the controversy; or of any other misbehavior by which the rights of any party have been prejudiced.” The phrase “guilty of misconduct” is inherently moral. Guilt implies culpability, a conscious choice to act improperly despite awareness of correct conduct. This language makes sense only for human actors capable of ethical judgments.

An AI system cannot engage in “misconduct” because it cannot comprehend behavioral norms or choose to violate them. It operates according to its programming; when it produces unintended results, this reflects programming errors or data problems, not wrongful conduct. The specific examples in Section 10(a)(3)—refusing to postpone a hearing or refusing to hear evidence—further reveal the statute’s presumption of human arbitrators who exercise judgment about procedural matters.

### *Section 10(a)(4): Exceeding Powers*

Section 10(a)(4) permits vacatur where “the arbitrators exceeded their powers, or so imperfectly executed them that a mutual, final, and definite award upon the subject matter submitted was not made.” This provision assumes arbitrators capable of understanding the scope of their delegated authority and either respecting or exceeding those limits. The concept of “exceeding powers” presupposes an agent who understands what they are authorized to do and chooses to do something beyond that authorization.

Applied to human arbitrators, this ground makes perfect sense. An arbitrator exceeds their powers when they decide issues the parties did not submit, apply law the parties contractually excluded, or award relief the parties’ agreement does not permit. These actions involve understanding the parties’ intentions, recognizing boundaries, and making decisions that transgress those boundaries. AI systems cannot “exceed their powers” in this meaningful sense because they have no understanding of authority, boundaries, or delegation. They execute algorithms. When an AI produces an output beyond what parties intended to authorize, this reflects defective programming or ambiguous inputs, not deliberate overreach.

Taken together, these textual requirements create what might be called a “vacatur gap”: a jurisprudential void where algorithmic

errors, hallucinations, and biases are immune from the statutory checks Congress designed to ensure arbitral integrity. Section 10 exists to police arbitral misconduct, but its provisions presuppose a decision-maker capable of moral failure. An algorithm cannot be bribed, cannot harbor hidden loyalties, cannot willfully ignore evidence. Yet an algorithm can produce outcomes infected by training data bias, fabricated citations, or systematic errors that disadvantage particular classes of litigants. Under the current statutory framework, such outcomes would be unreviewable. The very safeguards Congress enacted to protect parties from arbitral abuse become nullities when the arbitrator is a machine. This gap is not a minor technical problem; it represents a fundamental breakdown in the enforcement architecture that makes arbitration a legitimate alternative to judicial resolution.

### **The Nature of Arbitral Judgment**

Beyond Section 10's specific terms, the FAA's structure presupposes that arbitration involves the exercise of human judgment. Arbitrators weigh evidence, assess credibility, interpret contracts, and exercise discretion in fashioning remedies. Leading scholarship on AI has demonstrated that current AI models, despite their sophistication, operate through statistical pattern recognition fundamentally different from legal reasoning.<sup>38</sup> AI systems identify patterns in training data and generate outputs statistically likely to match those patterns. They do not engage in the purposive interpretation of legal texts or the moral evaluation of human conduct that arbitration requires.<sup>39</sup>

Consider witness credibility assessment, central to many arbitral proceedings. A human arbitrator observes demeanor, detects hesitation, notes inconsistencies in tone, and integrates these observations with documentary evidence and common sense. This involves what philosophers call "practical wisdom": the capacity to perceive morally salient features of a situation and respond appropriately. While AI systems can process transcripts and identify verbal patterns, they lack the capacity for genuine comprehension that characterizes human judgment. As leading AI researchers have acknowledged, current systems operate through sophisticated correlation rather than understanding; they identify statistical patterns without grasping meaning.<sup>40</sup>

Contract interpretation provides another illustration. When the Supreme Court instructs arbitrators to construe contracts by asking what a reasonable person in the parties' position would have understood a term to mean, it presupposes an interpreter capable of imaginatively inhabiting human perspectives.<sup>41</sup> An AI system simulates this inquiry by generating text statistically similar to outputs in its training data. But simulation is not the thing itself. The reasonable person standard assumes a mind capable of understanding intention, context, and commercial purpose.

\* \* \*

Editor's note: This article will conclude in the next issue of *The Journal of Robotics, Artificial Intelligence & Law*.

## Notes

---

\* The author, counsel in the Boston office of Murphy & King, may be contacted at [devans@murphyking.com](mailto:devans@murphyking.com).

1. Press Release, AAA-ICDR Launches AI Arbitrator for Construction Disputes (Nov. 3, 2025), <https://www.adr.org/news/aaa-icdr-launches-ai-arbitrator>.

2. See, e.g., <https://www.arbitrus.ai/>.

3. 9 U.S.C. §§ 1-16 (2024).

4. This article was written in December 2025 and reflects AI technology and the law as of that date.

5. *Loper Bright Enterprises v. Raimondo*, 603 U.S. 369 (2024) (overruling *Chevron U.S.A. Inc. v. Nat. Res. Def. Council, Inc.*, 467 U.S. 837 (1984)).

6. *Loper Bright Enters. v. Raimondo*, 603 U.S. 369, 400 (2024) (quoting *Wisconsin Central Ltd. v. United States*, 585 U.S. 274, 284 (2018)).

7. *Loper Bright*, 603 U.S. at 385 (quoting *Marbury v. Madison*, 5 U.S. (1 Cranch) 137, 177 (1803)).

8. *Taniguchi v. Kan Pac. Saipan, Ltd.*, 566 U.S. 560, 566-68 (2012).

9. *Dubin v. United States*, 599 U.S. 110, 118-20 (2023).

10. *FDA v. Brown & Williamson Tobacco Corp.*, 529 U.S. 120, 132-33 (2000).

11. *Epic Sys. Corp. v. Lewis*, 584 U.S. 497, 506 (2018).

12. *West Virginia v. EPA*, 597 U.S. 697 (2022).

13. *Id.* at 721.

14. *Id.* at 723.

15. *Whitman v. Am. Trucking Ass'ns*, 531 U.S. 457, 468 (2001).

16. *Black's Law Dictionary* 83 (2d Ed. 1910).

17. Webster's New International Dictionary 147 (1913).
18. Bouvier's Law Dictionary 228 (1914).
19. 9 U.S.C. § 5 (2024).
20. 9 U.S.C. § 7 (2024).
21. 9 U.S.C. § 9 (2024).
22. 1 U.S.C. § 1 (2024).
23. *Rowland v. Cal. Men's Colony*, 506 U.S. 194, 200 (1993) (explaining that Dictionary Act definitions apply "unless the context indicates otherwise").
24. *Return Mail, Inc. v. U.S. Postal Serv.*, 587 U.S. 605, 616 (2019) (explaining that Dictionary Act definitions "are not absolute" and yield when "the context indicates otherwise").
25. The AAA's own operations illustrate this distinction. The AAA administers arbitrations and promulgates rules, but it does not itself decide disputes. Human arbitrators on AAA panels render awards. When the AAA describes its new "AI Arbitrator," the system generates draft decisions that human arbitrators then review, revise as needed, and issue. The structure thus maintains the human decision-maker requirement even as it incorporates AI tools.
26. See *Burwell v. Hobby Lobby Stores, Inc.*, 573 U.S. 682, 706-07 (2014) (recognizing that corporations can exercise rights only through human actions and decisions).
27. Lei n.º 63/2011, de 14 de Dezembro, art. 9(1) (Port.) ("O árbitro deve ser uma pessoa singular e plenamente capaz.").
28. *Wetboek van Burgerlijke Rechtsvordering [Rv] art. 1023* (Neth.).
29. *Code de procédure civile [C.P.C.] art. 1450* (Fr.) ("Seule une personne physique ayant le plein exercice de ses droits peut être arbitre."). This requirement applies only to domestic arbitration; French law permits parties to appoint legal entities as arbitrators in international arbitration. See C.P.C. arts. 1506, 1508.
30. Robert Walters, *Robots Replacing Human Arbitrators: The Legal Dilemma*, 34 *Info. & Comm'ns Tech. L.* 129 (2025) (surveying Australia, Cayman Islands, European Union, India, Indonesia, Japan, Singapore, Qatar, Philippines, United Kingdom, and United States).
31. *Hall St. Assocs., L.L.C. v. Mattel, Inc.*, 552 U.S. 576, 584 (2008) (holding that FAA's vacatur grounds are exclusive and cannot be supplemented by contract).
32. 9 U.S.C. § 10(a)(1) (2018).
33. *Bonar v. Dean Witter Reynolds, Inc.*, 835 F.2d 1378, 1383 (11th Cir. 1988) (explaining that "undue means" requires conduct that is "immoral if not illegal" and "tantamount to corruption").
34. *A.G. Edwards & Sons, Inc. v. McCollough*, 967 F.2d 1401, 1403 (9th Cir. 1992) (holding that "undue means" requires intentional misconduct beyond mere error).
35. *Commonwealth Coatings Corp. v. Cont'l Cas. Co.*, 393 U.S. 145, 150 (1968).

36. *Scandinavian Reinsurance Co. v. Saint Paul Fire & Marine Ins. Co.*, 668 F.3d 60, 73 (2d Cir. 2012) (examining whether arbitrator's conduct reflected deliberate favoritism).

37. *Morelite Constr. Corp. v. N.Y.C. Dist. Council Carpenters Benefit Funds*, 748 F.2d 79, 84 (2d Cir. 1984).

38. Maxi Scherer, *Artificial Intelligence and Legal Decision-Making: The Wide Open?*, 36 *J. Int'l Arb.* 539, 554-57 (2019).

39. *Id.* at 564.

40. See Stuart Russell & Peter Norvig, *Artificial Intelligence: A Modern Approach* 1-35 (4th ed. 2020) (explaining that current AI systems operate through statistical pattern recognition rather than semantic understanding); see also Gary Marcus & Ernest Davis, *Rebooting AI: Building Artificial Intelligence We Can Trust* 15-45 (2019) (arguing that deep learning systems lack genuine comprehension and common-sense reasoning).

41. *Stolt-Nielsen S.A. v. AnimalFeeds Int'l Corp.*, 559 U.S. 662, 682-84 (2010).