Perspective Article *Open Access*

Journal of Pioneering Artificial Intelligence Research

ISSN: 3069-0846

DOI: doi.org/10.63721/26JPAIR0146

Regulating Innovation in an Era of Unparalleled Change Artificial Intelligence for Law: Regulating an Emerging Technology

Mara M. Burr, JD, LLM

Professor, Georgetown University Law Center, Washington, DC, USA

Citation: *Mara M Burr (2026) Regulating Innovation in an Era of Unparalleled Change Artificial Intelligence for Law: Regulating an Emerging Technology. J of Poin Artf Research 2(4), 1-4 WMJ-JPAIR-146*

***Corresponding author**: Mara M. Burr, JD, LLM, Professor, Georgetown University Law Center, Washington, DC, USA.

Submitted: 06.07.2026 **Accepted**: 10.07.2026 **Published**: 22.07.2026

Introduction

The Regulatory Challenge of Legal AI

Artificial intelligence is reshaping the legal profession at a pace that traditional regulatory systems were not designed to match. Large language models (LLMs) can summarize discovery, draft contracts, generate legal research memoranda, compare regulatory requirements, and assist self-represented litigants. These tools promise efficiency, lower costs, and broader access to legal information. Yet they also introduce serious risks: fabricated citations, confidentiality breaches, biased outputs, opaque decision-making, unauthorized practice of law concerns, and overreliance by attorneys or institutions that do not fully understand how the systems work.

The central problem for AI regulation in law is epistemic: regulators, judges, attorneys, clients, and even many developers often lack complete knowledge about how a model was trained, why it produced a particular answer, or how reliably it will perform in a new legal context. This knowledge gap makes it difficult to apply familiar legal concepts such as competence, negligence, causation, confidentiality, and professional supervision. At the same time, waiting for perfect understanding is not an option. AI is already embedded in legal workflows, and the public consequences of error can be severe. Decisions by several Bar Associations to sanction attorneys for providing AI-generated citations to cases that do not exist only underscore this point.

AI should be regulated through an adaptive, risk-based framework that protects core legal values while preserving room for innovation. The goal should not be to prohibit attorneys, courts, or companies from using AI. Instead, regulation should require accountable human judgment, transparent governance, careful validation, client confidentiality, and meaningful oversight for high-impact uses. In an era of unparalleled technological change, the law must move from reactive crisis management to durable governance.

The Promise of Artificial Intelligence for Law

AI can improve legal services in several important ways.

Reducing the amount of time required for routine legal tasks, including document review, contract comparison, transcript summarization, and first-draft research, allows for more critical thinking. It improves access to justice by helping non-attorneys understand forms, deadlines, and basic legal procedures. AI may assist organizations in identifying compliance risks earlier by monitoring large volumes of policies, communications, and transactions. Both courts and public agencies can improve administrative efficiency through streamlined procedures and processes.

These benefits matter because the legal system is often expensive, slow, and difficult to navigate. If properly governed, AI could help attorneys deliver faster work, help clients make better-informed decisions, and help institutions manage growing regulatory complexity. The case for regulation should therefore begin by recognizing that innovation itself serves public values. A regulatory approach that treats all legal AI as inherently suspect would risk freezing useful tools before their benefits can be realized.

The Epistemic Problem of AI Regulation

The difficulty is that AI systems are not ordinary tools. A calculator produces an answer through transparent rules. A word processor stores and displays what a user types. A large language model, by contrast, predicts language patterns from vast training data and can produce confident text without genuine legal understanding. In law, that distinction is critical. A legal answer must be accurate, jurisdiction-specific, current, ethically obtained, and tied to authoritative sources. A fluent but unsupported answer can be more dangerous than no answer at all because it creates the appearance of expertise.

The epistemic problem has at least three dimensions. First, there is a technical knowledge gap: users may not know what data was used to train the model, how the model weighs sources, or what limits exist in its reasoning. Second, there is an institutional knowledge gap: law firms, courts, and companies may adopt AI before they develop policies for validation, supervision, security, and disclosure. Third, there is a regulatory knowledge gap: lawmakers may regulate too early, locking in assumptions that become obsolete, or too late, allowing harmful practices to become normalized.

Risks Specific to Legal AI

Accuracy and Hallucination

The most visible legal AI risk is hallucination: the generation of false cases, statutes, quotations, or procedural claims. In ordinary business writing, a hallucination may be embarrassing. In litigation, it may mislead a court, harm a client, and trigger sanctions. Because legal authority depends on verifiable sources, any AI-assisted legal work must be checked against primary law and reliable secondary materials before it is used.

Confidentiality and Privilege

Legal practice depends on confidentiality. If attorneys place client information into public or poorly governed AI systems, they may expose privileged communications, trade secrets, personal information, or litigation strategy. Courts and regulators are increasingly attentive to whether a user had a reasonable expectation that information shared with an AI tool would remain confidential. For legal AI, privacy policies, data retention practices, training-use restrictions, and vendor controls are not technical details; they are professional responsibility issues.

Bias, Fairness, and Access to Justice

AI systems can reproduce or amplify bias in training data. In the legal context, that risk may affect bail assessments, sentencing recommendations, housing disputes, employment claims, immigration screening, or predictions about litigation outcomes. Bias is especially concerning when AI outputs appear neutral because they are produced by software. Regulation must therefore require testing for disparate impact, documentation of known limitations, and avenues for human challenge when AI affects rights, obligations, or access to legal remedies.

Unauthorized Practice and Overreliance

AI tools can make legal information appear more accessible, but they can also blur the line between general information and legal advice. Non-attorney platforms that generate individualized legal recommendations may raise unauthorized practice concerns. Attorneys face a related problem: overreliance. An attorney cannot delegate professional judgment to a machine. Competence requires understanding the tool's capabilities and limits, verifying outputs, and taking responsibility for final work product.

Existing Regulatory Approaches

The United States currently relies on a patchwork of professional responsibility rules, court orders, agency guidance, consumer protection law, privacy law, and emerging state legislation. This approach has the advantage of flexibility, but it can create uncertainty for attorneys and technology providers operating across jurisdictions. Bar guidance has generally emphasized that existing duties of competence, confidentiality, supervision, candor, and reasonable fees apply when attorneys use generative AI.

The European Union has taken a more comprehensive route through the AI Act, which uses a risk-based structure and imposes heightened obligations on certain high-risk AI systems. The EU model is important because it connects regulatory burden to potential harm rather than to the mere existence of AI. For legal systems, that insight is valuable: a grammar-assistance tool should not be treated the same as an AI system that predicts case outcomes, screens legal claims, or supports judicial administration.

Neither model is complete. A purely fragmented system may leave gaps and inconsistent expectations. A purely centralized system may become too rigid for fast-changing technology. The better path is a layered governance model: baseline duties for all legal AI uses, enhanced requirements for high-risk uses, and continuous updating through professional standards, court rules, audits, and regulatory sandboxes.

Proposed Framework for Regulating Legal AI

Risk Classification

Regulation should classify legal AI by use case. Low-risk tools might include spelling assistance, formatting, scheduling, or general document organization. Medium-risk tools might include contract summarization, due diligence support, or first-draft legal research. High-risk tools should include systems used for judicial decision support, legal eligibility determinations, litigation outcome prediction, automated client intake that affects access to representation, and tools that generate filings or legal advice without meaningful human review.

Human Accountability

Every legal AI framework should make one principle unmistakable: humans and institutions remain accountable. An attorney who submits an AI-assisted brief is responsible for its accuracy. A court that uses AI to manage cases must preserve judicial independence and due process. A company that deploys an AI compliance tool must be able to explain and monitor the system's performance. Accountability cannot disappear into the model.

Verification and Source Integrity

Legal AI systems should be designed and used in ways that support verification. Outputs that purport to state law should identify sources, distinguish binding from persuasive authority, and alert users when information may be incomplete or outdated. Attorneys and judges should be required to independently verify AI-generated legal authority before relying on it. For high-risk systems, validation should include documented testing against realistic legal tasks and periodic review after deployment.

Confidentiality, Security, and Vendor Governance

Law firms and legal departments should adopt vendor review standards before using AI with clients or confidential information. These standards should address data retention, training on user inputs, encryption, access controls, breach notification, audit rights, and deletion practices. Attorneys should not place confidential information into tools unless the confidentiality protections are understood and appropriate for the matter.

Disclosure and Consent

Disclosure rules should be targeted rather than symbolic. An attorney should not be required to disclose every use of AI for grammar or formatting. Disclosure may be appropriate when AI materially contributes to legal analysis, when a court order requires it, when client confidential information will be processed by an external system, or when the client's objectives reasonably depend on knowing how work is being performed. Client consent should be informed, practical, and tied to meaningful risk.

Regulatory Sandboxes and Continuous Learning

Because AI changes quickly, regulators should create mechanisms for experimentation under supervision. Regulatory sandboxes can allow courts, legal aid organizations, law firms, and vendors to test AI tools with safeguards, reporting obligations, and independent

evaluation. This approach encourages innovation while generating evidence about what works, what fails, and what protections are necessary.

Applying the Framework: Practical Recommendations

- For attorneys: Adopt written AI policies, verify legal authorities, protect confidential information, document human review, and train all users on the limits of generative AI.
- For law firms and legal departments: Inventory AI tools, classify use cases by risk, review vendor terms, maintain audit trails for high-risk uses, and establish approval processes before deploying new systems.
- For courts: Require candor when AI materially affects filings, discourage blanket bans that block useful innovation, and preserve due process whenever AI assists administrative or adjudicative functions.
- For regulators and bar associations: Harmonize guidance where possible, update professional responsibility rules through comments and opinions, and focus enforcement on uses that threaten clients, courts, privacy, or fundamental rights.
- For developers: Build legal AI tools that support source verification, access controls, logging, explainability, and clear warnings about limitations.

Conclusion

Artificial intelligence is not merely another legal technology or tool. It is a general-purpose system that can influence how law is researched, practiced, taught, administered, and accessed. Its promise is real, but so are the risks. The central regulatory challenge is to avoid two failures at once: allowing ungoverned systems to undermine legal accuracy, confidentiality, fairness, and accountability, and imposing rigid rules that prevent beneficial innovation.

The best approach is adaptive regulation grounded in our shared legal values. Legal AI should be governed by risk classification, human accountability, verification, confidentiality, meaningful disclosure, and continuous learning. In an era of unparalleled change, the law's task is not to stop innovation. It is to discipline innovation so that speed does not outrun judgment, and technological power remains answerable to the rule of law.