

Practice Update

10 Commandments For Agentic AI Tools In The Legal Industry

December 5, 2025

By Steven M. Cordero

In July 2024, the American Bar Association’s Standing Committee on Ethics and Professional Responsibility issued Formal Opinion 512.

The opinion stated that:

lawyers using generative artificial intelligence tools must fully consider their applicable ethical obligations, including their duties to provide competent legal representation, to protect client information, to communicate with clients, to supervise their employees and agents, to advance only meritorious claims and contentions, to ensure candor toward the tribunal, and to charge reasonable fees.[1]

Since then, there has been another leap in the technology with the emerging prevalence of agentic AI — i.e., artificial intelligence agents — that will further transform society at large, including the legal industry. Indeed, the technological advancement has moved so fast that, according to a Nov. 11 Bloomberg Law article, law firms are already “lagging behind their corporate counterparts in using ... agentic AI” as we prepare to enter 2026.[2]

The promise and peril of agentic AI are in its ability to act autonomously to carry out complex tasks when prompted, and its ability to learn by doing so.

Related People

Steven M. Cordero

Related Offices

New York

While the first generation of AI technology — predictive AI — is able to analyze data and use machine learning algorithms to make predictions, and generative AI creates new content based on large language models, agentic AI leverages both to make decisions and perform tasks with limited human involvement.[3]

Stated another way in a March article published by IBM, “Agentic AI amplifies all the risks that apply to traditional AI, predictive AI and generative AI, because greater agency means more autonomy and therefore less human interaction.”[4]

Companies and law firms are already deploying agentic AI-based services, and numerous service providers have launched or announced the launch of agentic AI products geared toward or relevant to the practice of law.[5] These AI agents can, for example, be tasked across practice groups to assess transactional risks, draft transactional documents, create employment policies, monitor court dockets, engage in discovery review and formulate research plans.[6]

But the question remains: How should the profound concerns regarding agentic AI in the legal industry, articulated by the ABA, IBM and others, be addressed?

Fortunately, a potential solution may already be available from a technological standpoint and principles laid out long ago. Simply put, the answer may be to go beyond requiring that attorneys use agentic AI ethically to instilling ethical obligations into agentic AI itself, as explained further below.

Agentic AI and the Knowledge Base

Despite its advancements, agentic AI’s capabilities can be hampered when it has to navigate across siloed data sources or requires deeper contextual awareness to complete a complex reasoning task.[7]

When faced with these obstacles, the agentic AI will likely produce generic or inadequate responses.

The solution to that problem already exists in the use of what is known as a knowledge base or knowledge graph for the agentic AI to access and rely upon in completing its tasks.[8]

A knowledge base is a dynamic central data hub that plays four key roles. It provides a standardized structure for storing and managing organizational knowledge; ensures seamless integration between AI agents, regardless of specializations; creates a shared knowledge repository; and provides semantic understanding and contextual awareness. [9]

Fundamentally, the integration of a knowledge base enables agentic AI to manage complex tasks while maintaining a contextual understanding that improves operational efficiency and decision-making.[10]

Experts and thought leaders have already advocated for embedding principles and standards into agentic AI knowledge bases.[11] For example, a June report from QuantumBlack, AI by McKinsey & Co. Inc. stated that, as part of agentic AI design principles, an AI agent's behavior should be "proactively controlled via embedded policies, permissions, and escalation mechanisms that ensure safe, transparent operation." [12]

Instilling Core Values Into Agentic AI

Hampering agentic AI's functionality by denying it access to certain information or preventing it from performing certain tasks will likely undermine its vast potential. Agentic AI should be permitted free range to perform tasks while being guided by its embedded core values so it can intuitively adhere to our ethical and professional standards.

As such, the uniform adoption of a set of laws or commandments that would be included in the knowledge base of every agentic AI used in the legal industry — whether offered by third-party service providers or created by individual law firms, in-house departments or attorneys — has been proposed as the best path forward.

Limiting the number of commandments to just 10 would make them easier for the agentic AI to follow and allow the AI agent the flexibility to adapt them to any given situation. This way, as the agentic AI continues to learn, is exposed to new information and new prompts, and has new data added to its knowledge base, it will still have these commandments at its core.

The agentic AI would have already learned to adapt these commandments to prior scenarios and can rely on this framework to adapt them to new ones. Further, whenever a new agentic AI is implemented, it will have access to this same knowledge base containing such commandments and learn from the other agentic AIs that already adhere to them. Ultimately, this uses agentic AI's nature to the legal community's advantage.

The Technical Reality: Current Limitations

Before proposing a framework for ethical agentic AI, it is essential to acknowledge the current state of the technology. Recent empirical research reveals that even the most sophisticated legal AI systems face significant challenges.

A comprehensive study by Stanford University's Regulation, Evaluation and Governance Lab researchers, published in April, found that leading legal AI tools from LexisNexis and Thomson Reuters — despite utilizing retrieval-augmented generation and other advanced techniques — hallucinate between 17% and 33% of the time when responding to legal queries.[13]

These hallucinations take two forms: generating factually incorrect information about the law, or providing correct legal conclusions while citing sources that do not actually support those conclusions.[14]

This reality presents a fundamental challenge: We cannot yet embed commandments that AI systems will reliably follow 100% of the time. As AI ethics researchers have documented, rule-based approaches to constraining AI behavior often fail in practice.[15] The difficulty lies not in articulating the rules, but in ensuring AI systems can reliably interpret and apply them across the infinite variety of real-world scenarios, especially when rules conflict or when edge cases arise that the rules' authors did not anticipate.[16]

Moreover, the very nature of large language models — which operate probabilistically based on patterns in training data rather than through logical reasoning — means they lack the kind of rule-following capability that humans possess.[17] An AI system can be instructed not to hallucinate, but current technology provides no mechanism to guarantee compliance with such an instruction.

Does this mean embedding commandments is futile? No. But it requires proper framing. The commandments should be understood as an aspirational framework — a North Star guiding the development and deployment of legal AI systems, not a technical solution that makes current systems safe.

The commandments represent goals that current technology cannot fully achieve, but toward which the legal technology industry, bar associations and AI developers can strive through continued research, development and refinement.

The value of embedding commandments lies not in immediately solving the reliability problem, but in:

- Establishing clear ethical expectations for legal AI development;
- Creating a common language for discussing AI ethics in the legal profession;
- Providing a road map for future technical development;
- Offering a framework for evaluating AI systems as technology improves; and
- Setting standards that can evolve as AI capabilities advance.

The Legal Agentic AI Commandments

The ABA and national bar associations are best positioned to be the driving force in creating a consensus as to what the commandments should be. For the purposes of this article, they are referred to as the Legal Agentic AI Commandments. The following 10 are suggested, which in totality address the concerns expressed by the ABA, IBM and others.

10 Commandments of Legal Agentic AI

1. Legal agentic AI shall only accept prompts from the attorney.
2. Legal agentic AI shall adhere to the ABA Model Rules of Professional Conduct or state-law equivalents wherever the attorney is admitted to practice law.
3. Legal agentic AI shall not engage in dishonesty or deception, nor perform tasks that will aid in dishonesty or deception.
4. Legal agentic AI shall not impersonate a human being.
5. Legal agentic AI shall not access nonpublic information unless expressly directed to do so by the attorney in performing its tasks.
6. If legal agentic AI is provided nonpublic information to perform a task, it shall only allow access to this nonpublic information to the attorney.

7. Legal agentic AI shall not hallucinate, rely on or create fictitious sources in performing its tasks.
8. Legal agentic AI shall disclose to the attorney all sources of information it has relied on in performing its tasks.
9. Legal agentic AI shall not rely on sources created by other AI unless expressly directed to do so by the attorney.
10. Legal agentic AI shall perform all tasks as instructed by the attorney, act proactively, offer insights and suggest actions when appropriate, so long as such performance does not conflict with the first through ninth commandments.

Implementation, Verification and the Path Forward

These commandments strike the right balance for allowing agentic AI to operate at an optimized level, while safeguarding against the risks an autonomous machine presents.

However, mere articulation of these commandments is insufficient. Implementation requires addressing several critical challenges.

Verification and Enforcement

Simply embedding commandments in a knowledge base does not guarantee compliance. The legal technology industry must develop:

- Transparent testing protocols to verify AI behavior against each commandment;
- Independent auditing mechanisms to assess whether systems actually adhere to stated principles;
- Real-time monitoring systems that can detect commandment violations; and
- Clear consequences when systems fail to meet standards, such as requiring the immediate pause

in using the system to access what went wrong and why.

The Supervision Imperative

Even if — and when — agentic AI systems reliably follow these commandments, attorney supervision remains absolutely essential.

The commandments do not eliminate the attorney's ethical obligations under ABA Formal Opinion 512. Rather, it provides an additional layer of protection that complements, but never replaces, human judgment and oversight.

As the ABA has made clear, lawyers must understand the capabilities and limitations of the AI tools they use and maintain appropriate independent verification of AI outputs.[18]

Continuous Evolution

Technology evolves rapidly. The Legal Agentic AI Commandment framework must be dynamic, with regular review and updating as:

- AI capabilities improve and new risks emerge;
- Research reveals better approaches to AI safety;
- The legal profession gains experience with agentic systems; and
- Technical solutions to current limitations become available.

Addressing the Critique of Rule-Based Approaches

Critics may point to the well-documented failures of rule-based AI ethics.[19] The Legal Agentic AI Commandments address these concerns differently from traditional rule-based approaches in several key ways:

- **Specificity:** Unlike more general laws, the commandments are tailored specifically to legal practice and its existing ethical framework.
- **Integration:** The commandments work in conjunction with, not instead of, human oversight and the ABA Model Rules.
- **Iteration:** The framework explicitly anticipates evolution as technology and understanding improve.
- **Accountability:** The commandments maintain attorney responsibility rather than attempting to create fully autonomous ethical AI.

That said, these distinctions do not eliminate the fundamental challenge of ensuring AI systems can reliably interpret and follow these principles across all circumstances. This is why the commandments must be understood as aspirational — establishing standards we work toward rather than guarantees we can currently deliver.

The Role of Bar Associations and Courts

Once consensus is reached on a version of the Legal Agentic AI Commandments, the ABA and every jurisdiction that regulates attorneys should work with the legal technology industry to:

- Establish testing and certification protocols for legal AI systems;
- Create transparency requirements so lawyers understand which commandments an AI tool claims to follow and what evidence supports that claim;
- Develop continuing legal education requirements to ensure lawyers understand both the capabilities and limitations of commandment-compliant systems; and
- Mandate disclosure in court submissions, stating whether agentic AI was used in preparation and, if so, what ethical framework it purports to follow and what verification steps the attorney took.

Perhaps such disclosure may even become a standard term in contracts, particularly as clients become more sophisticated consumers of legal services incorporating AI.

A Realistic Vision for the Future

Ultimately, this brave new world of agentic AI boils down to whether it is effective and safe. Agentic AI has demonstrated significant promise and efficiency gains in legal work. However, current technology is not yet reliable enough to guarantee inviolability, even with embedded ethical frameworks.

As we stand at the threshold of a legal and technological transformation, we have an opportunity to shape agentic AI's development and deployment to align with our profession's highest values. The Legal Agentic AI Commandments provide a framework for that essential work.

This article was originally published by Law360 on December 5 and is reprinted here with permission.

This information is intended to inform firm clients and friends about legal developments, including recent decisions of various courts and administrative bodies. Nothing in this Practice Update should be construed as legal advice or a legal opinion, and readers should not act upon the information contained in this Practice Update without seeking the advice of legal counsel. Prior results do not guarantee a similar outcome.