

The Intersection of Agentic AI and Emerging Legal Frameworks

19 December 2024

The evolution of artificial intelligence (AI) has introduced systems capable of making autonomous decisions, known as agentic AI. While generative AI essentially “creates” – providing content such as text, images, etc. – agentic AI “does” – performing tasks such as searching for and ordering products online. These systems are beginning to emerge in public-facing applications, including Salesforce’s Agentforce and Google’s Gemini 2.0.

As agentic AI continues to proliferate, legal systems must adapt to address the risks and harness the benefits of AI systems that are able to think more logically and take action rather than merely guide or create. Initiatives such as the California Consumer Privacy Act (CCPA) and its proposed modifications for automated decision-making technologies (ADMT) highlight the ways in which regulators are working to ensure privacy and accountability in the AI-driven era.

Practical Uses of Agentic AI

Agentic systems pursue complex goals using sophisticated reasoning with limited human supervision. Unlike traditional generative AI systems that respond to prompts, agentic AI can execute tasks using third parties as a user’s “agent.” For example, when prompted to book a flight an agentic AI will access flight databases, search for available flights based on the user’s preferences and budget, evaluate trade-offs in price and travel time, and finally book the flight by interacting with the airline’s booking system, inputting all necessary information for the passenger.

Further applications in various industries include health care, where it can assist in disease diagnosis by analyzing patient data and medical imaging, and finance, by enabling fraud detection and credit risk assessment through advanced data analytics. Retail companies can also utilize agentic AI to personalize shopping experiences, recommending products based on user behavior.

On a consumer level, portable agentic AI gadgets such as Rabbit R1 have introduced consumers to the early stages of autonomous decision-making. The device demonstrated how agentic AI can navigate third-party apps and perform tasks such as ordering food or booking rides via voice commands. On a small scale such as this, misunderstood prompts have minor consequences to users, perhaps leading to a wrong delivery order or sending a rideshare driver to the wrong location. However, when applied to a more complex use-case scenario, the ramifications of a misunderstood prompt are magnified.

Understanding Agentic AI and its Governance

While businesses will likely be able to streamline workflows and save resources, the legal and ethical implications of deploying such systems in consumer-facing roles demand careful consideration. For example, the integration of agentic AI into legal contract review, modification, and (someday soon)

negotiation raises a number of important implications that expose businesses and consumers to the greater risks of automating legally binding documents without human supervision and nuanced judgment.

ADMT and Emerging Regulatory Frameworks

The California Privacy Protection Agency (CPPA) has proposed a set of national standards for regulating Automated-Decision Making Technology (ADMT) to address growing concerns in agentic AI. Defined under the California Consumer Privacy Act (CCPA), ADMT includes “any technology that processes personal information and uses a computation to execute a decision, replace human decisionmaking or substantially facilitate human decisionmaking.”^[1] Importantly, the definition of “substantially facilitate human decision-making” specifies that ADMT includes instances where its output serves as a key factor in a human’s decision-making process.

However, the CCPA excludes certain technologies that do not independently execute decisions or significantly influence human decision-making from this definition. Examples include basic tools like spellchecks or calculators, which organize or compute data without making autonomous decisions.^[2] These distinctions establish the regulatory focus on agentic AI systems (coined ADMT) capable of independent or heavily influencing decision-making.

Cybersecurity Audits

The proposed rule mandates regular cybersecurity audits for businesses processing personal information that presents significant risks to consumers’ security. “Significant risks” apply to businesses that meet specific thresholds, namely deriving more than 50% of their annual revenues from selling consumer information or processing sensitive information of consumers or households. These businesses are required to complete audits annually without gaps by a qualified, objective, and independent professional covering all aspects of the business’s cybersecurity program to identify gaps, document findings, and outline plans to address any weaknesses.^[3]

For AI and ADMT use cases, Article 9 of the proposed rule compels businesses using ADMT technologies to adopt a proactive approach to security and risk management. For example, banks that use ADMT technologies in [loan origination automation](#) process sensitive consumer information. A breach could expose consumers’ credit histories or Social Security numbers, leading to identity theft or financial fraud. The rule’s proposal for robust cybersecurity audits would identify vulnerabilities in how data is stored and processed, mitigating unauthorized access and data breaches.

Risk Assessments

Building on the cybersecurity requirements of the proposed rule, Article 10 shifts the focus to management of the risks to consumer privacy for businesses using ADMT and AI systems. These assessments are required for high-stakes tasks that require extensive profiling, such as decisions on creditworthiness, health care eligibility, admission into academic programs, and hiring.^[4] Businesses must employ compliance members to conduct risk assessments to determine “whether the risks to consumers’ privacy from the processing of personal information outweigh the benefits to the consumer, the business, other stakeholders, and the public.”^[5]

This balancing test is especially impactful for AI and ADMT use cases in areas like hiring or customer profiling, where the potential for bias or harm must be weighed against operational efficiency. For example, businesses using ADMT to conduct emotional assessments to determine who to hire must conduct a risk assessment because it uses ADMT “for a significant decision concerning a consumer.”^[6] An ADMT-driven hiring tool that disproportionately affects certain demographics would need to demonstrate that its benefits outweigh these risks or face suspension under Article 10.

Risk assessments must be updated every three years or whenever significant changes occur in the technology or data processing activities.^[7] This ensures that AI models remain compliant as they evolve. These provisions mandate balancing privacy protection and innovation to ensure that ADMT systems are deployed responsibly and ethically.

Transparency and Accountability

Enhancing the cybersecurity focus of Article 9 and the forward-looking risk management framework of Article 10, Article 11 establishes rules for businesses leveraging ADMT to mandate transparency, fairness, and consumer control.

Article 11 requires businesses to provide consumers with a pre-use notice detailing the purpose of ADMT usage, its potential outcomes, and consumers’ rights in opting out or accessing the system’s logic and outputs.^[8] This aims to provide transparency in how personal data is processed and used by businesses that allow consumers to retain rights to their data.

In addition to transparency, Article 11 imposes requirements on businesses to evaluate the performance of ADMT systems, ensuring they operate as intended and do not result in unlawful discrimination.^[9] This includes testing systems for bias and verifying the quality of their outputs. By establishing these requirements, Article 11 aims to reinforce accountability in the deployment of ADMT while safeguarding consumer rights.^[10]

Risks Associated with Agentic AI

The use of agentic AI poses several risks, such as bias in decision-making, which can lead to unfair outcomes, especially in hiring or lending scenarios. Additionally, over-reliance on autonomous systems can result in operational disruptions if systems fail or produce erroneous outputs. Data security remains a critical concern, with potential exposure of sensitive information to breaches or misuse.

Mitigation and Management Strategies

To address these risks, businesses should implement rigorous testing and validation processes to detect and correct biases. Employing robust cybersecurity measures, such as encryption and regular audits, can mitigate data security threats. Human oversight mechanisms should be integrated to validate critical decisions made by agentic AI, ensuring accountability and reliability. Furthermore, businesses should invest in continuous education and training for staff to understand and manage AI systems effectively.

The CCPA’s proposed regulations provide a regulatory framework to govern the adoption and integration of agentic AI and ADMT. OpenAI’s [whitepaper](#) from December 2023 highlights that with proper governance, agentic AI can enhance productivity while maintaining safety and reliability.

A streamlined process using agentic tools to reduce time spent on consumer engagement development and routine legal tasks promises significant efficiency gains for businesses. The cybersecurity requirements reduce the risk of consumer harm by protecting personal data from breaches and enhancing the reliability of AI-driven processes, especially in legal matters that involve confidential client information. Article 10's emphasis on risk assessments builds on this foundation, requiring businesses to weigh risks against rewards.

The regulation ensures that agentic AI innovations remain equitable and aligned with ethical standards. Additionally, the transparency provisions aim for greater oversight of agentic AI systems to ensure that they are not only effective but also comprehensible to end users. As agentic AI systems develop for more complex decision-making scenarios, transparency will provide regulators and private parties with a level of control over systems.

However, with transformative potential comes new risks that the CCPA's proposed regulations may not cover. As is common in new technologies, the law lags behind innovation. When companies rush to adopt agentic AI, a lack of oversight creates risk in diminished output reliability, increased vulnerabilities to complex decision outputs, and labor displacements. Cybersecurity threats remain a pressing concern, and the compliance costs and operational complexities may hinder widespread adoption and invite legal loopholes. The balancing test under § 7154(a) – weighing privacy risks against operational benefits – can be subjective, leaving room for legal disputes. Similarly, Article 11's transparency mandates compel businesses to disclose the logic behind automated decisions, which can introduce conflicts between consumer rights and intellectual property protections.

Conclusion

Agentic AI offers transformative potential but introduces significant legal and ethical challenges. The CCPA proposed regulations provide a foundation for addressing these issues by emphasizing cybersecurity, risk assessment, and transparency. As legal frameworks evolve to keep pace with innovation, balancing accountability with progress will be critical to ensuring agentic AI is deployed responsibly and equitably.

Implementing robust governance frameworks is essential to navigate the complexities of agentic AI. Such frameworks guide the development and deployment of AI models, promoting the use of robust, unbiased, and high-quality data to derive outcomes. They also reduce compliance risks by establishing clear guidelines and standards aligning AI systems with legal and regulatory requirements, ensuring AI solutions operate ethically and legally. Moreover, integrating agentic AI into various sectors necessitates a paradigm shift towards inclusive design and democratic innovation. This involves moving beyond simply “adding” marginalized groups to AI discussions and ensuring that diverse perspectives are integral to AI development. By involving stakeholders from various disciplines, we can develop new theories, evaluation frameworks, and methods to navigate the complex nature of AI ethics, steering AI development in a direction that is beneficial and sustainable.

[\[1\]](#) See § 7001 (f).

[\[2\]](#) See § 7001. Definitions (f)(4).

[\[3\]](#) See § 7120; § 7121; § 7122; § 7123.

[\[4\]](#) See § 7150(a)(3)(A).

[\[5\]](#) See § 7152(a).

[\[6\]](#) See § 7150(c)(1).

[\[7\]](#) See § 7155(a)(2).

[\[8\]](#) See § 7220(a)/(c); § 7222.

[\[9\]](#) See § 7201(a)(1).

[\[10\]](#) See § 7201(a)(2).

Author(s)

Chanley T. Howell
Partner

**Alexander J.
Liederman**
Associate