AI Agents, Privacy, and the Importance of Context in Data Regulation

August 2025

Data & Trust Alliance

Executive Summary

Most current AI agents—used in customer service, logistics, employee support, fitness, and coding—operate well within existing U.S. data privacy and sector-specific regulations. As AI agent technology advances further in the future, policymakers should focus on the **context** in which the AI agent is used, *not* the technology itself. By evaluating factors like data sensitivity, purpose, user relationships, and potential harm, policymakers can craft smarter, risk-based laws that protect privacy while enabling innovation. Context isn't just helpful, it's essential.

Note: While this paper focuses on U.S. data privacy, context is also important for other jurisdictions.

Introduction

Published 08.06.2025

Authored by the Policy Committee of the Data & Trust Alliance

AI agents are the next evolution in AI and are transforming how individuals and organizations interact with technology—automating tasks and entire workflows, streamlining services, and delivering highly personalized experiences. AI agents act with varying degrees of autonomy to perform tasks on behalf of users, based on goals, preferences, and environmental inputs, and some AI agents use those interactions to continuously improve and self learn.

Just like with any other AI system or use of a consumer's personal data, the risks associated with AI agents depend heavily on what data is being used or accessed, who is using the data, for what purpose, and under what conditions. In fact, when built and deployed responsibly, AI agents can offer more privacy protections and serve as adaptive and secure alternatives to older, less adaptive systems. Policymakers should approach AI agents with that in mind, beginning with existing legal and regulatory frameworks and addressing gaps or unique issues as they emerge. As lawmakers consider how to govern these new use cases responsibly, it is essential to recognize that *context matters*.

What is an AI Agent?

AI agents are autonomous software tools that perform tasks, make decisions, and interact with their environment intelligently and rationally. They use artificial intelligence to learn, adapt, and take action based on real-time feedback and changing conditions. AI agents can work independently or as part of a bigger system, learning and changing based on the data they process.

Regulating the Present: How U.S. Frameworks Already Address Common AI Agents[1]

AI agents introduce new and powerful ways to augment humans by automating tasks. Below are five examples of how today's AI agents are used across sectors in ways that are within the bounds of U.S. data privacy laws such as state privacy laws, sector-specific rules (like Health Insurance Portability and Accountability Act "HIPAA" where applicable), and long-standing oversight by agencies such as the Federal Trade Commission (FTC). Here are some examples showing how AI agent use cases fit within established legal frameworks.

1 Renner, Matt, and Chaban, Matt A.V. "601 Real-World Gen AI Use Cases from the World's Leading Organizations." Google Cloud Blog. Last modified April 9, 2025. https://cloud.google.com/transform/101-real-world-generative-ai-use-cases-from-industry-leaders.

 © 2025 The Data & Trust Alliance, all rights reserved

Inventory Management Agents in Retail Supply Chains

A major U.S. retailer is deploying AI agents to optimize inventory forecasting and restocking decisions by analyzing structured internal datasets such as historical sales, weather forecasts, and promotional calendars.

Because AI agents process enterprise operational data that does not include personal identifiers or customer data, they do not implicate privacy regulation or present privacy risks. Internal data handling is covered by standard enterprise data governance practices and vendor contracts, ensuring sufficient oversight under existing commercial data regulations.

Customer Service Agents in the Appliance Industry

A nationwide appliance manufacturer is implementing AI agents to assist in customer support. These agents reference internal enterprise documentation to walk users through common product issues and, when necessary, escalate inquiries to human agents.

These customer AI agents are designed, trained, and operate on structured, non-sensitive company data—such as product manuals and repair instructions. These manuals and instructions are often readily publicly available online. There may also be a need to use customer information such as service history, warranty status, and contact details. In these cases, the AI agent data use is no different than when a human agent completes the same task. Since the customer data remains purpose-bound and the AI agent does not introduce automated profiling or decision-making beyond support triage, existing privacy frameworks are adequate.

Fitness Coaching Agents Handling Health-Adjacent Data

A fitness company serving tens of thousands of members across the U.S. has introduced an AI agent in its mobile app to help members find classes, check amenities, and receive personalized fitness tips.

While the AI agent uses preference data and activity history, it does not access or infer protected health information (PHI), and therefore, is not subject to HIPAA. State privacy laws do apply, requiring disclosures about how data is collected and used, as well as opt-out options for data sharing or sale. Because the AI agent operates in a narrow, consent-based context, serving only known users through a private app interface, existing privacy frameworks help ensure consumer information is safeguarded.

Customer-Facing Agents in Customer Support Operations

A large digital platform company has adopted AI agents to support its frontline employees. These AI agents summarize past customer communications, highlight prior touchpoints, and help staff respond efficiently to user requests.

Here, the AI agent accesses and processes previously collected customer communication data that is already governed under customer data usage policies and applicable privacy regulations.

This function is governed by existing data privacy laws, including FTC guidelines on fair data practices and state laws like the California Consumer Privacy Act (CCPA), which requires notice and purpose limitations. Provided that role-based access controls are in place and the AI agent's data use is consistent with the information the company provided to customers at the time of collection and use, current regulations are sufficient.

Code Agents for Enterprise Software Development

An electric vehicle and software company have deployed AI coding agents to assist developers with debugging, code generation, and refactoring. The AI agents are customized to the company's internal code base and engineering standards.

Because these AI agents interact exclusively with internal code repositories and software documentation—not personal data—they do not invoke data privacy obligations. Their use falls under internal IT controls and secure development lifecycle protocols. To the extent any telemetry or developer interactions are logged, such data would fall under existing employee privacy policies and relevant workplace data use laws.

© 2025 The Data & Trust Alliance, all rights reserved

Implications for Regulation: A Contextual Framework for AI Agents

As the above real-world examples demonstrate, many of today's uses of AI agents are consistent with existing privacy laws and frameworks. As AI agents become more integrated into business operations, public services, and daily life, policymakers will inevitably face decisions about how and when to regulate them. The real-world use cases discussed demonstrate that there is nothing inherently riskier about AI agents from a privacy perspective. Indeed, many current AI agent deployments are simply tools to help individuals and businesses use the data they already have, albeit in a more automated way.

Early generations of any new technology are often just applications of the technology to existing business processes, while subsequent generations often usher in new uses, innovations, and business models. The same may eventually happen for emerging or future AI agents, particularly those that handle more sensitive data, operate with higher degrees of autonomy, or function in high-risk contexts. As these novel AI agents emerge, gaps in existing frameworks may emerge. When that time comes, **context must guide the design, scope, and application of those regulations.** This approach enables more precise protections without unnecessarily constraining innovation.

If and when gaps emerge in how existing privacy frameworks apply to AI agents, policymakers should apply the following principles:

Adopt a Risk-Based Approach
Regulation should be proportionate to the level of risk posed by the data and the task. AI agents analyzing operational supply chain data do not warrant the same scrutiny and oversight as those recommending changes to a patient's medical treatment plan. A risk-based approach ensures that protections are strongest where the consequences are most significant, such as in credit risk, education, healthcare, or law enforcement contexts. And existing sector-specific regulations (e.g., GLBA, HIPAA) already provide additional safeguards in such areas.

Move Beyond Broad Definitions
Instead of relying solely on expansive or vague definitions of "AI" or "personal data," regulatory approaches should focus on the *specific data being used*, *who is using it*, *why it is being used*, and *how it might affect individuals*. Many AI agents are performing functions that are uncontroversial when humans or well-established software and productivity tools perform the same task with the same data.

Recognize and Center the User
Many of today's AI agents function either directly with or on behalf of a clearly identified user—such as an AI assistant helping gym members book classes in the gym's mobile app. These contexts inherently include transparency, consent, and predictable data flows.

© 2025 The Data & Trust Alliance, all rights reserved

Promote Outcome-Based, Technology-Agnostic Rules
Rather than regulating based on the underlying technology, legislation should target *outcomes* focusing on whether data use is explainable and transparent. This allows the regulatory framework to remain relevant even as AI capabilities evolve, ensuring it applies equally to generative models, rule-based agents, or emerging brain-computer interfaces.

Ensure Transparency and User Understanding
The level of transparency required will vary depending on the context. In some circumstances, it will be obvious to consumers that they are using an AI agent (e.g., using an AI agent to schedule a fitness class) and additional transparency will not be needed. In other cases, additional transparency may not be needed even where the use of an AI agent is neither apparent to the consumer nor significant (e.g., where an AI agent manages inventory for an online purchase the consumer makes). But where AI agents make decisions that materially affect individuals, users should be given notice and additional insight into the rationale behind those outcomes, including the key factors that influenced the final decision. *Note: The Data & Trust Alliance has done work on transparency between AI developers and deployers including data provenance, and there are other mechanisms like model cards and system cards to promote transparency.*

Incorporate Ethical and Sector-Specific Standards
Context also means acknowledging that sectors like healthcare, education, and finance are already subject to ethical codes, industry standards, and statutory obligations. Rather than creating duplicative or conflicting AI-specific rules, new regulation should align with and reinforce these existing safeguards.

Why Context Must Guide Regulation

Without a contextual lens, regulators risk applying a one-size-fits-all approach that regulates AI technology uniformly when not all AI presents the same risks. Policymakers risk overregulating safe, low-risk applications that can impede already existing business processes, while leaving other more sensitive or high-risk applications inadequately addressed. Context transforms abstract data points into meaningful regulatory signals by asking:

- **What type of data is being used and in what use case?** A class-scheduling AI agent using gym preferences isn't equivalent to a diagnostic AI agent analyzing genomic sequences.
- **Why is the data being collected?** Is it to streamline a service the user expects, or to monetize the data in opaque ways?
- **What role does the user play in the interaction?** An AI agent acting directly with or on behalf of a user is different from an AI agent performing background tasks.
- **What harm could result if things go wrong?** Could misuse lead to fraud, discrimination, misdiagnosis, or reputational damage—or is the worst case simply a late shipment or bad recommendation?

© 2025 The Data & Trust Alliance, all rights reserved

- **What are the expected benefits?**
  Does the AI agent provide benefits such as improved access to education, health outcomes, or quality of life?
- **What do users reasonably expect?**
  Were they clearly informed, and do they understand their rights and obligations?

The most widely deployed AI agents, like those managing inventory, resolving customer issues, or assisting software developers, are often low-risk and privacy-aligned when implemented with care. But tomorrow's AI agents may operate in less predictable environments, across highly sensitive domains, with new forms of autonomy or insight into human behavior.

As use cases emerge, contextual regulation—one that aligns risk, purpose, sensitivity, and transparency—offers a path forward. It is not only more precise and resilient but better reflects how people and organizations experience AI in the real world.

By grounding regulation in use case analysis, we can design rules that are protective where needed, permissive where appropriate, and capable of evolving as AI itself continues to evolve.

We stand ready to share insights and information with lawmakers to ensure that the future of AI is both safe and transformative. As emerging applications test the boundaries of existing frameworks and demand even greater clarity around context, intent, and proportionality, it will be essential to understand context to effectively regulate.

These and other novel use cases will continue to challenge conventional regulatory categories and demand new ways of thinking about context, purpose, and risk. The work ahead is significant. Together, we must remain vigilant and adaptive, ensuring that our governance frameworks evolve in step with and support innovation, and that context remains our compass in navigating the path forward. ●

© 2025 The Data & Trust Alliance, all rights reserved

Authored and endorsed by a subset of Data & Trust Alliance organizations

AT&T	Nielsen
IBM	Regions Bank
Mastercard	Salesforce
Meta	Transcarent

About D&TA

The Data & Trust Alliance is a group of industry-leading enterprises committed to a future powered by the responsible use of data and AI.

We leverage the collective expertise and influence of our members— among them the leading deployers of data and AI in business—to create and adopt practices that enhance trust in data, in AI models, and in the people and process through which they are deployed.

Learn more at dataandtrustalliance.org.

Please note that the views expressed in these materials do not necessarily reflect the views of any D&TA member organization.

© 2025 The Data & Trust Alliance, all rights reserved