D&TA Recommendation: Framing "Consequential Decisions"

September 2024

Data & Trust Alliance

Executive summary	
• Artificial Intelligence ("AI") is a powerful technology that has the potential to revolutionize the way we live and work.	• When new regulation is necessary, it is useful to focus on whether and how AI is used to make consequential decisions that directly affect people's lives.
• Rather than attempt to regulate AI as a technology as *such*, policymakers should instead focus on potential harms the technology poses as *applied*, in specific use cases.	• Policymakers should also strive to harmonize regulation across jurisdictions
• When considering regulation, regulators should begin by reviewing existing regulation to determine whether it covers issues of concern. New regulation should only be introduced when necessary to address regulatory gaps.	• To assist policymakers in their efforts, this memorandum offers principles for how best to approach AI regulation to protect people while allowing innovation and creativity to thrive.

Artificial Intelligence (AI) can create unparalleled opportunities for economic advancement, creativity, and innovation. But like all technology, if misused, AI can create risks of harm. To protect people from those harms, policymakers may respond with unnecessary, confusing, or overly broad regulation—to the detriment of AI's great potential promise. A better approach is to prioritize balancing protections with maximally allowing people to reap AI's many benefits. This includes creating a policy landscape that promotes consistency and flexibility across jurisdictions to promote a healthy technology ecosystem.

Many companies, cross-industry groups, policymakers, academics, and members of civil society are diligently striving to advance this approach through proposed standards, frameworks, and regulatory initiatives. The approach below is designed to add to—not replace—those efforts. Our recommendation (see box on final page) is intended to preserve policymaker resources and drive regulatory balance, leaving room for existing technologies and beneficial, low-risk applications that do not raise novel issues.

New Regulation Should Address True Gaps

But first, before considering new regulation, policymakers should review the existing legal landscape. In many cases, existing laws and regulatory guidance sufficiently cover the relevant issues of concern, particularly if they are technology-agnostic. Relying on existing frameworks to the extent possible reduces the risk of creating overlapping or conflicting rules that lead to legal uncertainty and inconsistent protections. For example, civil rights, consumer protection, security, intellectual property, and importantly, data privacy rules often already address many of the most important risks associated with AI.

Only where gaps do exist should policymakers consider new regulations, and then they should tailor solutions in a way that preserves innovation. This work should be done at the national level with the goal of international alignment, where possible, and, in the U.S., state level efforts to fill gaps should be tailored and harmonized.

Focus Regulation on High-Risk AI Use Cases

Once a gap is identified, policymakers should then focus on scoping the potential AI-related use cases associated with potential harms at issue. Specifically, we suggest that policymakers focus on regulating high-risk use cases. Taking time to isolate and specify high-risk use cases prevents burdening unintended (and beneficial) ones and, when married with the landscape of regulations, is necessary to a balanced regulatory approach. Organizations should have the opportunity to demonstrate that their use cases may not actually pose high risks to individuals in practice, especially given a selected context or industry.

One effective emerging approach for addressing an important regulatory gap focuses on instances in which AI is used to make "consequential decisions" in specific and appropriately identified high-risk use cases. Strong implementations of this model are designed to trigger regulatory requirements only when three conditions are met: (1) AI is a "controlling factor" in (2) the provision or denial of goods or services, that (3) are related to specifically enumerated use cases.

We further recommend defining "controlling factor" as one that "makes or is dispositive in making" a relevant consequential decision. This establishes a clear standard that acknowledges the benefits of human oversight. It also avoids creating an overbroad definition that could include common business software tools that may be used to inform human decision-making without rising to the level of an AI consequential decision.

The "provision or denial of" clause importantly focuses the regulation on specific impacts and outcomes. This construct recognizes that AI is used in numerous and varied circumstances, many of which are unlikely to cause harm (e.g., IT optimization and other internal operations). In addition, pairing "consequential decision" with "controlling factor" recognizes that organizations leveraging AI to support (but not dictate) a consequential decision made by a person—for example to streamline a process or process large amounts of data—is not the appropriate target for additional requirements.

Colorado SB24-205 is an existent example of legislation leveraging the "consequential decision" model to regulate high-risk use cases. SB24-205 applies to developers and deployers of "high-risk artificial intelligence systems," which are defined as "any artificial intelligence system that, when deployed, makes, or is a substantial factor in making, a consequential decision." 6-1-17-01(9)(a). In that law, "consequential decision" is defined as "a decision that has a material legal or similarly significant effect on the **provision or denial** to any consumer of, or the cost or terms of" **education enrollment or opportunity, employment, financial services, essential government services, health-care services, housing, insurance, or legal services**[1]. While this framing aligns in many ways to the model we've outlined, Colorado uses "substantial factor" to describe the level of involvement necessary in decision making. This definition does not make a clear distinction between an AI-driven decision versus a decision that is ultimately made by a human. Therefore, future legislation should clarify the level of involvement in AI making the decision by changing the standard to "controlling factor," the definition of which appropriately accounts for human oversight.

[1] Colorado SB24-205, 6-1-17-01(3) (emphasis added)

Leverage Existing Regulations & Regulatory Bodies

No matter how they approach these issues, policymakers should always strive for consistency and harmonization in definitions and to reduce friction and enhance the innovative potential of AI systems. Any regulation should be developed or revised with the aim to promote regulatory harmonization and flexibility. The adoption of consistent nomenclature or flexible industry standards will be key to ensuring that organizations are able to continue to develop and adopt emerging technologies.

Likewise, when examining *how* to regulate AI, it is important to also consider *who* should do it. Existing regulatory bodies[2] are already well equipped to address most any new risks that AI may pose in their areas of responsibility. For example, in the U.S., financial regulators (e.g., OCC, FRB, and FDIC) have closely monitored and regulated the use of credit decisioning and quantitative models, including the use of AI in those contexts, for years. And the Department of Homeland Security and other federal agencies are proactively addressing critical infrastructure protection and other national security issues. These types of targeted approaches lead to greater consistency, more appropriately tailored rules, and help avoid a patchwork of conflicting requirements.

In contrast, regulation that presumes the creation of a new agency or similar bureaucracy is likely to lead to regulatory duplication and stifle investment in and development of AI technology without a corresponding benefit.

We should therefore leverage existing structures where possible, including by providing resources and support to ensure existing regulators have the knowledge and resources necessary to mitigate the risks of AI use in their domains. New regulatory bodies—like new regulation itself—should only be introduced when real gaps exist.

[2] https://www.ncsl.org/technology-and-communication/artificial-intelligence-2024-legislation

The Recommendation

To promote harmonization, enhance clarity and accountability, and reduce friction, regulators should start with existing legal frameworks and leverage the "consequential decision" approach when regulating to close specific gaps.

As policymakers work to limit harm, we recommend they seek to close gaps by regulating artificial intelligence systems focus on systems that are high risk:	01	Narrowly tailor regulation to address gaps in existing regulatory frameworks while encouraging and supporting existing bodies in meeting the evolving demands of modern technology.
	02	Pair "consequential decision" with at least one other trigger definition (such as "controlling factor") before introducing a due diligence requirement to reduce unnecessary overhead.
	03	Define "consequential decision" as a decision that has a **material legal or similarly significant effect** on the **provision or denial** to any consumer **of specifically enumerated appropriate high-risk use cases** to promote harmonization and reduce burden on companies.
We recommend governments seeking to expand existing guidance to include regulation of high-risk artificial intelligence systems:	01	Assess the potential gaps in existing guidance and authorities.
	02	Ensure existing regulators have the knowledge and resources necessary to mitigate the risks of AI in their domains including collaborating with the private sector to better understand the use and risks associated with relevant uses.
	03	Embed AI rules and guidelines into existing frameworks, as appropriate.

Authored and endorsed by a subset of Data & Trust Alliance organizations	Transcarent	Meta
	Salesforce	GM
	Responsible Innovation Labs	IBM
	Nielsen	AT&T
	NFL	

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. Our only KPI is adoption by practitioners.

Data & Trust Alliance released "Policy Recommendations for the Responsible Use of Artificial Intelligence" in June 2024. This paper builds upon that foundational work.

Learn more at dataandtrustalliance.org.