Atlantic Council
CYBER STATECRAFT INITIATIVE

ai x ac

September 2025

Securing data in the AI supply chain

AI technologies involve a complex supply chain. Policymakers and technologists should better secure the AI supply chain's data components by taking a comprehensive view of the data in the AI supply chain and identifying which components can be more robustly secured and which risks require new mitigations.

Bottom lines up front

- AI technologies are underpinned by a complex supply chain—organizations, people, activities, information, and resources enabling research, development, deployment, and more—including human talent, compute, and institutional and individual stakeholders. Another core element is data.
- Policymakers can easily fall into the trap of focusing on one AI data component at one moment and another the next, risking a lopsided policy that fails to take account of all the AI data components important for AI R&D—and that need to be secured. These include training data, testing data, models (themselves), model architectures, model weights, Application Programming Interfaces (APIs), and Software Development Kits (SDKs).
- Developers, users, maintainers, governors, and securers of AI technologies should use this data component framework to identify relevant security mitigations. They should secure AI data components with existing data protections where sufficient (e.g., encryption for training data at rest) and leverage or develop new data protections for relatively AI-unique risks (e.g. the implanting of "neural backdoors" in AI models).

Executive summary

AI technologies are underpinned by a complex supply chain—organizations, people, activities, information, and resources enabling research, development, deployment, and more. The AI supply chain includes human talent, compute, and institutional and individual stakeholders. This report focuses on another element of the AI supply chain: data.

While there is a diversity of data types, structures, sources, and use cases in the AI supply chain, policymakers can easily fall into the trap of focusing on one AI data component at one moment (e.g., training data circa 2017), then switching focus to another AI data component next (e.g., model weights in current times), risking a lopsided policy that fails to take account of all the AI data components that are important for AI R&D. Put simply, a "one size fits all" approach to AI-related data runs the risk of creating a regulatory, technological, or governance framework that overfocuses on one element of the data in the AI supply chain while leaving other critical parts, and questions, unaddressed. This is insufficient for mitigating risks to AI data components, from errors to data leakage to intentional model exploitation and theft.

The data in the AI supply chain includes the data describing an AI model's properties and behavior as well as the data associated with building and using a model. It also includes AI models themselves and the different digital systems that facilitate the movement of data into and out of models. The report therefore spells out a framework to visualize the seven data components in the

Graphic 1: Components of AI Data Supply Chain

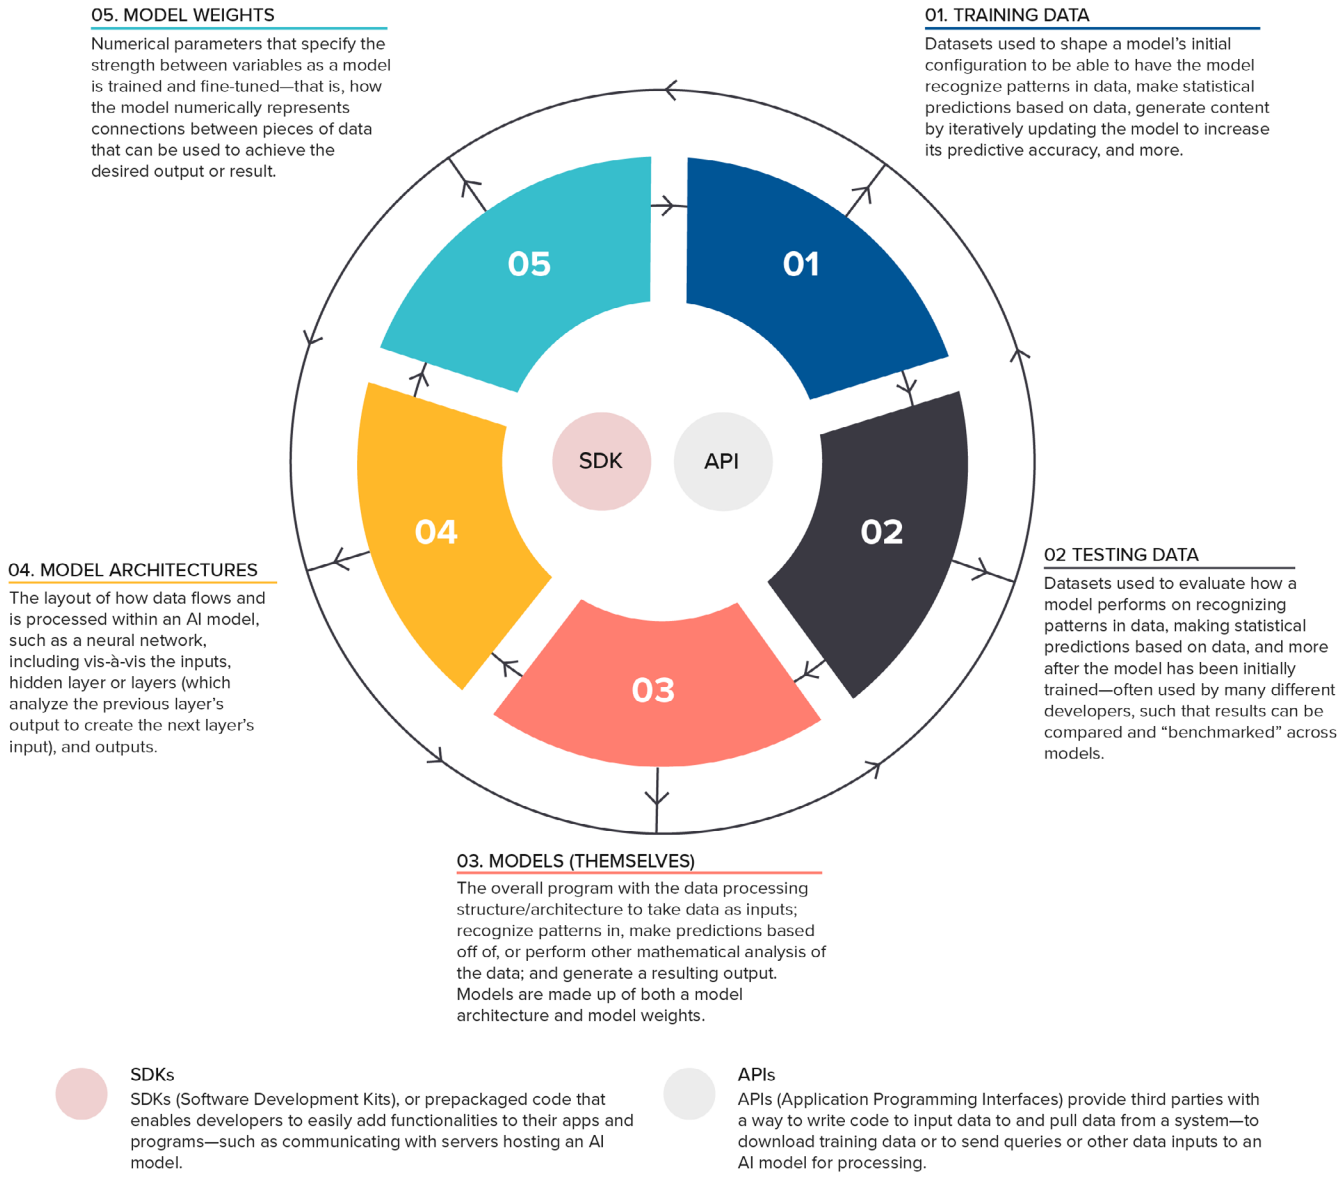

AI supply chain: training data, testing data, models (themselves), model architectures, model weights, Application Programming Interfaces (APIs), and Software Development Kits (SDKs).

It then uses the framework to map data components of the AI supply chain to three different ways that policymakers, technologists, and other stakeholders can potentially think about data risk: data at rest vs. in motion vs. in processing (focus on a data component within the supply chain and its current state); threat actor risk (focus on threat actors and risks to a data component within the supply chain); and supply chain due diligence and risk management (focus on a data component supplier or source within the supply chain and related actors).

In doing so, it finds that many risks to AI-related data are risks to data writ large that could be mitigated through existing best practices, such as NIST- and ISO-specified

Fig. 1: Three potential approaches to securing data in the AI supply chain

Step 1: : Identify Component and/or Source in Supply Chain	Step 2: Identify Security Approach	Step 3: Identify Mitigations from Existing Data and Supply Chain Security Best Practices
Approach one, two, and three: Training data, testing data, models (themselves), model architectures, model weights, APIs, or SDKs? Who is involved in sourcing the data, including the creation of the data, any subsequent processing of the data, and its dissemination?	**Approach one:** Is the data at rest, in motion, or in processing? (Focus on a data component within the supply chain and its current state.) **Approach two:** What are the threats, vulnerabilities, and consequences of security risks to specific components? (Focus on threat actors and risks to a data component within the supply chain.) **Approach three:** Who are the suppliers, and what are their security controls—or risks? (Focus on a data component supplier or source within the supply chain and related actors.)	**Approach one:** Identify relevant protections for data at rest, in motion, or in processing, such as NIST SP800-53 SC-28 ("Protection of Information at Rest"). Tailor to the use case and, if applicable, the orgavnization's supply chain role(s). **Approach two:** Identify relevant sources of information about the threat actor. Secure vulnerable systems (e.g., using NIST, ISO best practices) that match up against the threat actor's capabilities (and could be a way to access a specific data component in the AI supply chain). Tailor protections for the data component itself or components themselves (e.g. encryption, storage rules) based on the threat actor's targeting intent. **Approach three:** Conduct due diligence, to extent possible, on supply chain actors' ownership, maliciousness, and susceptibility to malicious influence—drawing on best practices from financial sector and other know-your-customer checklists Conduct due diligence, to extent possible, on the data security, cybersecurity, and supply chain security measures taken by those suppliers, to attempt to mitigate risks through the supply chain—drawing on best practices from GDPR compliance supply chain risk management, and so on.

data access controls, continuous monitoring systems, and robust encryption. Simultaneously, this report also finds that some security risks to AI data components do not map well to existing security best practices. At least two risk-mitigation pairings stand out: attempts to poison AI training data require data filtering mechanisms not well captured by existing measures, and which access controls or encryption would not appropriately mitigate; and emerging, malicious efforts to insert so-called neural backdoors into the behavior of neural networks require new security protections, too, beyond the realm of traditional IT data security. On top of implementing these two categories of mitigations themselves, this report emphasizes that organizations can leverage "know your supplier" best practices to ensure all other entities in their AI supply chains have security best practices

About this memo

Written by

Justin Sherman
Nonresident senior fellow
Atlantic Council

Justin Sherman is a nonresident senior fellow at the Atlantic Council's Cyber Statecraft Initiative. He is also the founder of Global Cyber Strategies, a Washington DC-based research and advisory firm, an adjunct professor at Georgetown University's School of Foreign Service, a contributing editor at Lawfare, and a columnist at Barron's.

About the center

The Cyber Statecraft Initiative works at the nexus of geopolitics, technology, and security to craft strategies to help shape the conduct of statecraft and to better inform and secure users.

The author thanks Trey Herr, Nitansha Bansal, Kemba Walden, Devin Lynch, Harriet Farlow, Ben Goldsmith, Kenton Thibaut, and other individuals who could not be thanked by name for their comments on earlier drafts of this report, as well as all the individuals who participated in the background and Chatham House Rule discussions about issues related to data, AI applications, and the concept of an AI supply chain.

for both non-AI-specific and AI-specific data risks.

Recommendations

This report concludes with three recommendations.

1. Developers, users, maintainers, governors, and securers of AI technologies should map the data components of the AI supply chain to existing cybersecurity best practices—and use that mapping to identify where existing best practices fall short for AI-specific risks to the data components of the AI supply chain.
2. Developers, users, maintainers, governors, and securers of AI technologies should "Know Your Supplier," using the supply chain-focused approach to mitigate both AI-specific and non-AI-specific risks to the data components of the AI supply chain.
3. Policymakers should widen their lens on AI data to encompass all data components of the AI supply chain. This includes assessing whether sufficient attention is given to the diversity of data use cases that need protection (e.g., not just training data for chatbots but for transportation safety or drug discovery) and whether they have mapped existing security best practices to non-AI-specific and AI-specific risks.