CIPL Response to the Office of Science and Technology Policy's Request for Information on the Development of an Artificial Intelligence (AI) Action Plan

Submitted March 14, 2025

The Centre for Information Policy Leadership (CIPL)[1] welcomes the opportunity to respond to the U.S. Office of Science and Technology Policy (OSTP)'s request for information on the Development of an Artificial Intelligence (AI) Action Plan.[2] CIPL commends the OSTP's efforts to define priority policy actions to sustain and enhance U.S. AI leadership and innovation. As a data- and privacy policy think tank, CIPL's comments focus on concrete policy actions relating to data governance practices that support innovation and ensure effective and resilient AI development and deployment.

AI is transformational to industries and society alike, driving growth, societal benefits, and opportunities for all. From healthcare to financial services, access to vast amounts of data is central to AI's success. However, this reliance on data – specifically, personal information and sensitive information – raises important ethical, legal, and societal considerations. Judicious and proportionate rules, accompanied by accountable self-governance structures within organizations, can foster development and deployment of AI in a manner that engenders customers' trust and favors more rapid AI innovation and adoption. An effective AI Action Plan would ensure that existing laws, such as those covering data privacy, do not act as a barrier to the use of data where it is needed for AI training, development, and deployment. It would also mitigate and potentially resolve the patchwork of regulatory approaches created by U.S. state laws on data privacy and AI that erode the competitive advantage of operating in the U.S. Furthermore, an effective AI development plan should incentivize and reward organizations who invest in and develop trusted and robust AI technologies.

At the global level, U.S. leadership and cooperation with allies is needed to ensure that AI governance frameworks are interoperable and that mechanisms for trusted cross-border data

---

[1] **The Centre for Information Policy Leadership (CIPL)** is a global privacy and data policy think tank within the Hunton law firm that is financially supported by the firm, 85+ member companies that are leaders in key sectors of the global economy, and other private and public sector stakeholders through consulting and advisory projects. CIPL's mission is to engage in thought leadership and develop best practices for the responsible and beneficial use of data in the modern information age. CIPL's work facilitates constructive engagement between business leaders, data governance and security professionals, regulators, and policymakers around the world. For more information, please see CIPL's website at www.informationpolicycentre.com. Nothing in this document should be construed as representing the views of any individual CIPL member company or Hunton. This document is not designed to be and should not be taken as legal advice.

[2] National Science Foundation, "Request for Information on the Development of an Artificial Intelligence (AI) Action Plan," 90 FR 9088 (Feb. 6, 2025), available at https://www.federalregister.gov/documents/2025/02/06/2025-02305/request-for-information-on-the-development-of-an-artificial-intelligence-ai-action-plan.

flows remain strong, so as to ensure access to and the sharing of the rich datasets required for AI development.

I. RELEVANT WORK AND EXPERTISE

For more than 20 years, CIPL has been a thought leader on organizational accountability and risk-based solutions as the foundation for smart regulation, responsible data governance, and critical innovation, including the development and deployment of AI. Our benchmarking report, "Building Accountable AI Programs: Mapping Emerging Best Practices to the CIPL Accountability Framework," outlines best practices and case studies on how 20 leading organizations are responsibly developing and deploying AI through the lens of CIPL's Accountability Framework.[3] This Framework can and is being used by organizations to develop their own internal data and AI governance programs both in the absence of applicable laws or in conjunction with such laws.

In an environment where AI regulations are on the rise globally and in U.S. states, CIPL issued a report in 2023, outlining an approach to AI regulation that avoids overly prescriptive rules and emphasizes a flexible, risk- and outcomes-based approach that relies on a range of proven organizational accountability measures. Our report also considered laws that already regulate AI-related risks to help inform self-regulatory and other potential frameworks for AI governance. Specifically, our paper providing "Ten Recommendations for Global AI Regulation," which proposes a three-tiered approach that would minimize potential risks of harm to both individuals and society, while enabling the responsible development and deployment of AI.[4]

CIPL's most recent discussion paper, "Applying Data Protection Principles to Generative AI: Practical Approaches for Organization and Regulators,"[5] considers key privacy and data protection concepts and explores how they can be effectively applied to the development and deployment of AI models and systems.

---

[3] CIPL, "Building Accountable AI Programs: Mapping Emerging Best Practices to the CIPL Accountability Framework", February 23, 2024, https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl_building_accountable_ai_programs_23_feb_2024.pdf.

[4] CIPL, "Ten Recommendations for Global AI Regulation", October 2023, https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl_ten_recommendations_global_ai_regulation_oct2023.pdf.

[5] CIPL, "Applying Data Protection Principles to Generative AI: Practical Approaches for Organizations and Regulators", December 2024, https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl_applying_data_protection_principles_genai_dec24.pdf.

Given CIPL's expertise in AI within the data policy sphere, our response[6] will focus specifically on the following themes:

- Organizational accountability as a driver of trust, investment, and adoption
- Smart regulation for efficiency and security
- Transitioning to the workplace of the future
- Resolving the tensions between data privacy and AI Innovation
- Key components of a data and privacy governance framework for AI
- Ensuring global leadership

II. Organizational Accountability as a Driver of Trust, Investment, and Adoption

Accountability requires organizations to put in place demonstrable internal governance mechanisms that ensure responsible AI development and use, as well as compliance with applicable laws. The implementation of formal internal governance processes and procedures designed to foster accountability within an organization can enable the laws regulating AI to be less prescriptive and more flexible. For these reasons, CIPL encourages policies that inspire and support the adoption of organizational accountability frameworks for the development and deployment of AI technologies. This approach helps to advance the benefits of AI while also identifying and addressing potential harms.

Organizational accountability is recognized as a key building block for effective organizational governance and a foundation for smart regulation.[7] By enabling responsible AI development and use, accountability-based governance frameworks promote trust in the AI ecosystem – a vital component for promoting investment in, and fostering the adoption and use of AI by individuals and organizations. Building such trust is a strategic imperative for the U.S. as well as a commercial necessity for businesses,[8] but recent research suggests that public trust in AI in the U.S. lags behind other countries. Unless this "trust gap" is closed, other countries could see faster adoption of AI. [9]

A U.S. AI Action Plan should encourage and incentivize organizational data governance centered on accountability. As noted, the CIPL Accountability Framework has been used by organizations as a blueprint for AI and data governance programs. The core elements in CIPL's

---

[6] This document is approved for public dissemination. The document contains no business-proprietary or confidential information. Document contents may be reused by the government in developing the AI Action Plan and associated documents without attribution.

[7] CIPL, "Organizational Accountability – Past, Present, and Future," October 30, 2019, cipl_accountability_past_present_future_oct19.pdf.

[8] KPMG, *Trust in Artificial Intelligence,* 2023, https://kpmg.com/xx/en/our-insights/ai-and-technology/trust-in-artificial-intelligence.html#accordion-797cc1e444-item-70e3e0a5d2.

[9] Ina Fried, "Exclusive: Trust in AI is Much Higher in China than in the U.S.," *Axios*, February 13, 2025, https://www.axios.com/2025/02/13/trust-ai-china-us?utm_term=emshare.

Accountability Framework are captured in the diagram below and are documented in more detail in the Appendix to this submission.

Source: CIPL

While a core set of accountability practices should be required for organizations developing and deploying AI, policymakers should encourage and incentivize the adoption of a range of practices, frameworks, tools, and technologies. The goal should be to create an environment where organizations view the adoption of thoughtful accountability frameworks—including self-regulatory or co-regulatory codes of conduct and certifications—as differentiators for creating value and deepening trust, in addition to fulfilling existing legal obligations established, for example, in consumer protection, intellectual property, and civil rights law, as well as applicable sectoral state laws pertaining to personal information and AI.

III. SMART REGULATION FOR EFFICIENCY AND SECURITY

AI is generating wide and growing societal benefits, but regulatory guidance may be needed to address a range of potential risks and harms associated with AI that could undermine these benefits. CIPL recommends a flexible, risk-based approach that builds on existing laws and standards and on accountability practices of organizations. This approach should be backed by innovative oversight and co-regulatory instruments. Any legislative or regulatory approach to AI should follow these overarching recommendations:[10]

- ***Create a flexible and adaptable framework that defines the outcomes to be achieved, rather than prescribing details of how to achieve them.*** To be effective, an AI Action Plan must be able to remain relevant as technology and use-cases continue to advance. Any rules should be technology neutral: a framework that is overly prescriptive and specific to individual technologies or current business models and practices can become quickly outdated by inhibiting beneficial innovations.

---

[10] CIPL, " Ten Recommendations for Global AI Regulation," October 10, 2023, https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl_ten_recommendations_global_ai_regulation_oct2023.pdf.

- ***Adopt a risk-based approach that considers risks and benefits holistically***. A risk-based regulatory or governance framework for AI would provide non-exhaustive criteria to assist organizations to determine the likelihood and severity of potential harms associated with their AI products or services and the measures available to mitigate them while also preserving the intended benefits and desired functionality of the AI applications. Assessing and understanding potentially harmful impacts in the development or deployment of AI applications would allow organizations to tailor their mitigations to the actual risks and avoid the implementation of unnecessary measures that could stifle innovation or beneficial deployment.

- ***Build on existing legal foundations***. An AI action plan should build on existing legal frameworks, including federal and state regulations and legislation. Many sectors where AI finds application are already highly regulated (*e.g.*, healthcare and finance), and existing laws and regulations (including consumer protection, privacy, civil rights, and intellectual property laws) already provide requirements, compliance structures, and remedies that apply to the use of AI. Relying on existing legal frameworks reduces the risk of creating overlapping, conflicting, or excessive rules that could lead to burdensome, uncertain, and inconsistent legal obligations. Also, where appropriate, existing rules can be supplemented by industry standards and co-regulatory tools developed in partnership with stakeholders, such as codes of conduct, certifications, and assurance models.

- ***Empower individuals through transparency, explainability, and mechanisms for redress.***

    - *Transparency.* Developers and deployers of AI should provide context-appropriate and meaningful transparency about the operations of AI systems, while preserving privacy and data protection, security, safety, and trade secrets.

    - *Explainability* is an aspect of transparency and a means of boosting accountability and trust. It requires developers and deployers to explain meaningfully how AI systems affect automated decisions and outcomes that impact individuals, while bearing in mind other concerns, such as security, safety, and accuracy.

    - *User Feedback and Redress* also boost accountability and trust. For externally facing tools, where individuals do not understand an AI-enabled automated decision, or believe they have been harmed by AI, there should be clear options for user feedback, inquiries, complaints, and further transparency, where appropriate. Other factors may include the right to contest the decision, a requirement for human review, and, ultimately, redress by enforcement authorities, where appropriate.

- ***Create mechanisms for coordination and cooperation across agencies and enforcement bodies***. AI is used across sectors governed by different agencies and

enforcement bodies. While each agency should maintain competence over its own remit for purposes of legal certainty, an AI action plan can set high-level AI policies and goals applicable across all sectors and industries, and require alignment, coordination, and joint action between different regulatory bodies where appropriate. This coordination is more critical than ever in an environment of fragmenting AI regulatory approaches. The growing privacy and AI patchwork inhibits U.S. innovation and competitiveness in the long run.

IV. Transitioning to the Workplace of the Future

AI tools represent a tremendous opportunity and require a critical new skill set for employees in the future. Businesses should prepare their workforce for the impact of AI productivity tools by offering AI training and education for new and current employees, and government should invest in workforce development and education programs related to AI. Organizations should also establish clear policies and procedures regarding the use of AI at work, and perform periodic impact assessments to understand the effects, positive and negative, of AI in the workplace. Incentives for companies that adopt AI responsibly, and guidance on how to do so such as introducing voluntary standards, could expedite AI adoption and mitigate potential negative impacts.

V. Resolving the Tensions between Data Privacy and AI Innovation

Personal information has emerged as one of the world's most valuable resources and a critical asset in the digital economy, particularly in the development of AI. As public awareness of information collection practices grows, consumer expectations around privacy and information security also evolve. The pace of AI development cannot be determined just by how fast the technology itself develops, but by the willingness of consumers and businesses to use it and of investors to invest in it.[11] A 2024 survey found that trust in AI is higher in much of the developing world than it is in the U.S.[12]

By embedding robust, responsible AI practices into their operations, companies can not only ensure compliance with applicable laws but also enhance consumer trust and long-term engagement. Well-thought out, flexible, and risk-based AI policies and governance frameworks can enable organizations to generate meaningful innovations without impinging on the privacy rights of individuals. These governance mechanisms are not an impediment to AI development, use, and innovation. Instead, they provide the strategic foundation for the long-term viability of AI. An AI policy that is focused on resolving the tensions between the advancement of AI and data privacy principles will go far to close the "trust gap" between the

---

[11] Ina Fried, "Exclusive: Trust in AI is Much Higher in China than in the U.S.," *Axios*, February 13, 2025, https://www.axios.com/2025/02/13/trust-ai-china-us?utm_term=emshare.

[12] Ibid.

U.S. and the rest of the world. The next section offers more specific recommendations on how to achieve this goal.[13]

VI. KEY COMPONENTS OF A DATA AND PRIVACY GOVERNANCE FRAMEWORK FOR AI

AI requires access to large and diverse datasets to reach its full potential. However, in the absence of clear privacy protections, consumers and businesses may be reluctant to share data, thereby limiting the availability of the critical inputs required for the training of AI. An AI Action Plan can provide new guidance or recommendations for policies and best practices ensuring that data is collected, shared, and used in an accountable manner to foster responsible acceleration in the development and adoption of AI.

Robust protections for data privacy are critical for building trust in AI, but if interpreted too rigidly, longstanding data privacy principles can inhibit innovation. This is especially true in the context of AI. Under a U.S. AI Action Plan, privacy principles should be designed or interpreted to promote effective data use for AI development, while also addressing the unique challenges or risks that AI poses. To ensure that data privacy rules enable rather than hinder development and adoption of beneficial AI technologies, CIPL recommends an AI Action Plan that reflects the following privacy principles:[14]

- An AI action plan should **facilitate lawful mechanisms for the use of personal information in model training**. Where existing laws or frameworks address permissible uses of personal information (such as in U.S. state laws), lawmakers, regulators, and enforcement agencies should avoid legal interpretations that are unduly restrictive regarding the use of personal information in AI model training, development, and deployment.

- Different data privacy rules, considerations, and mitigations apply in different phases of the AI lifecycle – data collection, model training, fine tuning, and deployment. Regulators and organizations should **interpret data privacy principles separately in the context of each relevant phase of the AI technologies**.

- Data privacy principles should be drafted or interpreted to **recognize and enable the processing and retention of sensitive personal information for AI model training**, as this is necessary to ensure accurate outputs. In addition, sensitive personal information may be necessary for the training and development of certain AI systems whose sole purpose is based on the processing of sensitive personal information or to deliver benefits to specific categories of individuals (such as accessibility tools).

---

[13] CIPL, "Artificial Intelligence and Data Protection in Tension," October 29, 2018, https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl_first_ai_report_-_ai_and_data_protection_in_tension__2_.pdf

[14] CIPL, " Ten Recommendations for Global AI Regulation," October 10, 2023, https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl_ten_recommendations_global_ai_regulation_oct2023.pdf.

- Developers should be incentivized to employ **privacy-enhancing and privacy-preserving technologies** (PETs/PPTs), such as synthetic data and differential privacy. This would enable AI models to have the rich datasets they need during training while reducing the risks associated with the use of personal information.

- Privacy laws should enable the collection and use of data that is **necessary and appropriate** for building high-quality AI models. Some U.S. state privacy laws and proposals for federal laws contain rules around "data minimization" and related concepts around purpose limitation and use limitation.[15] These concepts should be understood **contextually** and should not be interpreted to limit the collection and use of data that is necessary for the intended purpose (*e.g.*, model training, model fine-tuning, or model deployment for a particular purpose). As such, they should not automatically be understood to prohibit or conflict with the collection and use of large volumes of data.

- Policymakers should be thoughtful about **the extent to which individuals can control the use of Information from and about them.** For example, where appropriate and practicable, individuals should be able to request that information they enter into an AI system via input prompts, and the outputs they receive in response, not be used for model training or fine-tuning. When consumers can trust that their information is going to be handled ethically and securely, they are more likely to engage with AI-driven products and services. Increasing consumer trust ultimately increases the volume and diversity of data available for AI training, leading to the creation of more robust, fair, and accurate AI models.

- At the same time, policymakers should recognize where it is appropriate to set bounds on individuals' ability to exercise certain requests, such as when those requests could create an unreasonable burden with respect to model retraining. For example, in the case of web scraped data that is used during training but not catalogued or further filtered to identify personal information, it may be unreasonable or impractical for a developer to respond to requests for erasure. Developers may be able to apply alternative measures, such as output filters, to satisfy an individual's request.

- Transparency in the context of AI models should be contextually appropriate and meaningful, while also fulfilling transparency requirements under applicable laws and regulations. **Transparency should not come at the expense of other important factors,** such as usability, functionality, and security, nor should it create additional burdens for users.

- Policymakers should consult with developers and deployers of AI systems to **clarify the distinctions between duties and responsibilities** across the phases of AI development.

---

[15] The term "data minimization" describes the practice where organizations only collect and store the minimum amount of personal information necessary for its intended use.

For AI to reach its full potential and to ensure the U.S.'s continued preeminence in AI, policymakers must enable data- and privacy governance frameworks that protect individuals' legitimate privacy and other interests, as well as foster general societal trust in AI.

VII. ENSURING GLOBAL LEADERSHIP

The use of AI spans national borders, as do its challenges, which include the consideration of data privacy, security, and intellectual property rights, and effective access to data across borders. These challenges require a coordinated global response and mutual interoperability between jurisdictions. Priorities of the AI Action Plan in this area should include:

- Creating efficient and reliable cross-border data transfer mechanisms to enable access to robust, globally sourced datasets for AI development.
- Guaranteeing global interoperability between AI governance frameworks. This broadens access to data-driven AI products and services, reduces compliance costs for organizations, and increases legal certainty for businesses and consistent protections for individuals. Continued U.S. participation and leadership within G7 and OECD working groups focused on AI will be vital to advance this goal.
- Ensuring that restrictions to cross border flows of personal information are only instituted when necessary and are proportionate to the risks presented by the transfer, taking into account the type of data being transferred and the broader context of the cross-border transfer.

The U.S. has a strategic opportunity to take a leadership role in the creation of effective and globally interoperable AI governance frameworks that maximize both AI innovation and appropriate risk mitigation. Indeed, the U.S. already has been a leader in the creation of the only multilateral data transfer and accountability-based data governance framework through the Global CBPR Forum, which will be instrumental for enabling AI development and deployment as well as U.S. AI leadership going forward.[16]

[16] The Global Cross-Border Privacy Rules (Global CBPR) are a multilateral accountability-based privacy and data transfer certification, developed by the U.S. and a growing number of international partner countries and jurisdictions. See www.globalcbpr.org.

Appendix A – Emerging Best Practices in Accountable AI Programs, Mapped to the CIPL Accountability Framework

The following table contains a sample of emerging best practices and examples from accountable AI programs used by organizations from different sectors, geographies, and sizes. These practices are mapped to the corresponding element of the CIPL Accountability Framework. The practices are not intended to be mandatory industry standards but rather serve as examples of how companies are implementing specific practices to foster accountability in their development, deployment, and use of AI technologies. Each of the following should be calibrated based on risks, industry context, business model, size, and maturity level of the organization.

Accountability Elements	Related Practices
Leadership and Oversight	• Establishing "tone from the top" and demonstrating a commitment to advance ethics, values, and specific principles in AI development, deployment, and use • Implementing systematic processes and escalation pathways for AI-related decision making • Establishing AI ethics oversight bodies or committees (internal or external) to review risky AI use cases and promote ongoing improvements to AI practices • Appointing a board member for AI oversight • Appointing a responsible AI lead, AI officer, or AI champion • Setting up an internal interdisciplinary AI board or AI committee • Establishing organization-wide AI ethics principles • Ensuring that AI development, product, and governance teams include diverse experiences, perspectives, and qualifications. • Creating a centralized governance framework with oversight from the top that still provides flexibility within internal teams • Expanding the remit of privacy teams to include AI-related responsibilities • Leveraging the expertise of other relevant teams (*e.g.*, engineering, data science, legal, ethics and compliance, etc.) to ensure multidisciplinary, cross-functional AI teams

	• Encouraging employee reporting throughout all levels of the organization by offering escalation pathways to resolve potential AI-related issues
Risk Assessment	• Developing algorithmic impact assessments or fairness assessment tools to monitor and continuously test algorithms to avoid human bias, unfair discrimination, and "concept drift" throughout the entirety of the AI lifecycle • Requiring AI risk assessments at multiple points throughout the AI lifecycle, particularly for new or updated use cases or applications • Creating ethics, human rights, and/or data protection impact assessments • Creating a risk taxonomy that categorizes AI-related risks and allows for uniform assessment • Keeping a centralized repository of all risk assessment documentation • Developing standardized risk assessment methodologies that consider the benefits of the AI application or use, the likelihood and severity of risk factors on individuals and/or society, the level of human oversight needed for individually automated decisions with significant impact (*e.g.*, legal ramifications), the ability to explain the technology in the appropriate context, and the ability to audit its effectiveness • Documenting considerations (*e.g.*, accuracy, data minimization, security, transparency, scope of impact, benefits to society) for high-risk processing • Assessing data quality against key performance indicators (KPIs) • Evaluating the data vis-à-vis the purpose of its use (*i.e.,* the quality of the data, its provenance, whether it is personal, synthetic, in-house, or externally sourced) • Developing frameworks for data preparation and model assessment – including feature engineering, cross-validation, back-testing, standardized KPIs • Enabling close collaboration between business and data experts (*e.g.,* data analysts, data engineers, IT, and software engineers) on a regular basis to assess accuracy, ensure appropriate outputs, and allow for proper use of the model

	• Using privacy enhancing technologies (PETs) to preserve the privacy and security of AI systems • Outlining escalation pathways to send AI-related issues to an AI ethics council or other oversight body • Evaluating and testing models in specific application contexts prior to widespread deployment
Policies and Procedures	• Adopting specific AI policies and procedures on how to develop, deploy, or sell AI • Drafting policies on the application of privacy and security by design principles throughout the AI lifecycle • Setting rules on the level of verification for data input and output • Requiring pilot testing of AI models before release • Specifying the use of protected data (*e.g.*, encrypted, pseudonymized, tokenized, or synthetic data) in training AI models • Creating a glossary of AI-related terms for internal use and reference • Promoting the use of smaller, higher quality datasets • Cleaning and curating datasets before model training through automated or manual checks • Considering relevant and appropriate use of PETs and PPTs to integrate privacy and security controls into AI models • Outlining special considerations for organizations creating and selling AI models, software, applications • Developing a fairness or AI impact assessment to analyze and mitigate AI-related risks • Creating due diligence/self-assessment checklists or tools for business partners deploying AI • Clearly defining escalation steps for reporting high-risk AI issues • Implementing an ideation phase with all stakeholders (*e.g.*, data scientists, business, final user, control functions) where needs (including explainability), outcomes, validations rules, maintenance, and budget are discussed • Implementing specific policies for internal use

	• Requirement to incorporate diverse perspectives, experiences, and qualifications in relevant teams and business functions • Implementing internal policies in parallel with forthcoming AI regulation • Translating internal principles-based policies to third-party vendor agreements, language, and due diligence processes • Creating processes for review of high-risk AI use cases by an AI ethics board or council
Transparency	• Tailoring transparency measures for the different needs of end users, regulators, business partners, and internal stakeholders at all stages of the AI lifecycle • Communicating disclosures in a simple, easy-to-understand manner • Considering how AI disclosures can be accessible for those with special needs/disabilities • Establishing a transparency trail to explain automated decision-making and broad workings of algorithms • Providing notice when the system relies on AI/ML • Providing counterfactual information (*e.g.,* how different inputs can affect the output of an AI model) • Understanding customers' expectations and deploying AI technologies based on their readiness to embrace AI • Implementing tiered transparency • Defining criteria for internal deployment of AI technologies based on usage scenarios and communicating them to users • Publishing model or system cards (*i.e.*, short documents accompanying AI models that describe the context in which a given model is intended to be used and how the model performs in a variety of conditions) • Creating a data hub for information regarding data governance, data accessibility, data lineage, data modification, data quality, etc. • Tailoring transparency to the identified risk (*e.g.,* using watermarking for generative AI output) where possible and appropriate • Participating in benchmarking opportunities, public engagement, and regulatory sandboxes

	• Using visualization tools to depict difficult, technically complex concepts to end users
Training and Awareness	• Providing specific training for data scientists and engineers, including how to address relevant ethical issues (*e.g.,* how to limit and address bias) • Creating opportunities for cross-functional training (*e.g.*, between privacy professionals and AI engineers) • Tailoring trainings regarding ethics and fairness in AI for relevant teams • Compiling and making available AI use case information where relevant risks have been mitigated or deployment has been halted • Creating a "translator" role that helps explain the impact and technical capacities and limitations of AI • Sharing case studies to help employees learn how to address potentially complex, ethically challenging AI cases • Incentivizing compliance with completing ethics training by pairing it with eligibility for bonuses, pay raises, and/or promotions, or incorporating it into other mandatory training activities
Monitoring and Verification	• Incorporating "human in the loop" (HITL) in design, oversight, and redress • Identifying and understanding which business functions are using AI • Providing the capability for human audit of input and output • Ensuring human review of individual decisions with legal or similarly significant effects • Monitoring the data ecosystem—from data flow in, through data process, to data flow out • Using different auditing techniques • Deploying counterfactual testing techniques • Pre-defining AI audit controls • Creating an internal audit team with expertise in AI and other emerging technologies • Monitoring AI models (*e.g.*, back-testing and feedback loop) and conducting ongoing maintenance • Red teaming and adversarial testing of AI models
Response and Enforcement	• Enabling redress mechanisms to remedy an AI decision • Permitting redress through a human, not to a bot

	• Developing communication channels for internal (*e.g.,* for employees) and external (*e.g.,* end users, business customers) to report and address feedback, complaints, requests, etc.