

Executive Summary

In today's rapidly evolving technological landscape, artificial intelligence (AI) has emerged as a powerful force with the potential to reshape various aspects of society, from economic prosperity to national security. However, only through careful consideration and a deliberate approach to regulation can we harness the benefits of AI and mitigate its potential risks. Critically, AI is not a single technology but rather a family of related, but distinct, technologies, each of which may be applied in significantly different contexts. Applying rules designed for one type of AI or one context to another situation can hinder the development of new forms of AI and create, rather than reduce, harms.

To ensure effective regulation and self-governance of AI, a multistakeholder approach is vital. Drawing from the successes of the broader internet governance ecosystem, a similar framework can be applied to AI governance. Such an approach allows for diverse perspectives, fosters innovation, and accommodates the evolving nature of AI technologies.

Existing laws can address aspects of AI that are not unique to the technology. Whether performed by a human or an AI, illegal discrimination already violates federal and state laws, for example. Allowing existing law to cover AI overall, while also identifying the limited instances where AI introduces unique challenges that may require discrete additions to existing law, will result in a predictable and stable environment for AI investment, limit duplicative regulation and regulatory arbitrage, and ensure that the benefits of AI flow to Americans while mitigating potential harms.

Regulation will also play a vital role in engendering trust in AI systems. By establishing clear guidelines and standards for transparency and accountability, regulation can help address concerns related to privacy, bias, and accountability. But overly prescriptive approaches, like those under the EU's AI Act, may hamper the development of the next generation of AI technologies. And regulation of AI can also create outcomes that are antithetical to the U.S. system of democratic institutions, as with China's draft law requiring AI services to obtain political pre-approval.

Regulating AI Requires Understanding AI

AI has already become an integral part of our lives. Technologies like speech and facial recognition and machine translation are forms of AI that are already widely used. While recently developed technologies like Large Language Models and transformer-based image generators have drawn recent attention, regulation of AI must avoid unintended consequences by taking into account these other forms of AI, as well as the rapid pace of advancement in AI technology. New types of AI are continuously being developed, making it challenging to predict the precise direction of advancement in AI technology. To foster innovation and progress, it is important not to implement rigid regulations that rely on the present mechanisms by which AI operates, but rather to take approaches that manage overall risk in a way that incorporates the context in which each AI system operates. One example of such an approach is the National Institute of Standards and Technology (NIST) AI Risk Management Framework, which was created per Congressional direction.

Among the existing types of AI, there are several prominent examples worth mentioning:

- ❖ **Automated Decision-Making (ADM):** Algorithms autonomously make decisions based on predefined rules and data. Existing practical applications of ADM are nearly endless, with ADM used in diverse fields from scaling content moderation tools to increasing access to financial credit.
- ❖ **Machine Perception:** Enables machines to understand sensory inputs. This includes computer vision and speech recognition. Practical applications of machine perception can be seen in Shopify's automatic product description generation, making it easier for businesses to create detailed product listings, and in accessibility tools that automatically describe images for visually impaired individuals.
- ❖ **Natural Language Processing (NLP):** A form of AI that focuses on machine understanding of human language. NLP is often combined with machine perception to enable a machine to interact with humans more naturally. Applications like Google Translate and natural language search engines such as Google and LexisNexis exemplify the capabilities of NLP, and voice assistants like Siri, Alexa, and Google Assistant apply a combination of NLP and machine perception to listen to, understand, and respond to human requests.
- ❖ **Machine Learning (ML):** A technique for creating various forms of AI, including some of those used in NLP or machine perception. ML involves training algorithms with large datasets to recognize patterns and make predictions or decisions. Generative models and Large Language Models (LLMs) are examples of ML-based AI systems that have gained significant attention recently. These models have demonstrated impressive capabilities in generating realistic text, images, and even entire stories.

While these applications of AI may not hold the same level of attention as recent generative AI tools, they have already solved real problems. Translation allows people to access documents that were created in languages they don't speak. Image recognition has been used to detect potholes in roads and to improve weather forecasting. And automated decision-making techniques have helped to modernize occupational license processing and to make water management decisions more quickly and with better outcomes. These existing applications hint at the tremendous potential AI holds, if implemented responsibly with appropriate risk management.

Developing AI Responsibly Requires Flexible Regulation

In the rapidly advancing landscape of AI, responsible development and deployment are paramount. However, it is crucial to strike a balance between regulation and flexibility, avoiding overly prescriptive principles that may stifle innovation. To achieve this delicate equilibrium, the principles of responsible AI should be considered in designing thoughtful, adaptable regulation that can be applied in all contexts. Rather than being overly prescriptive, the focus should be on designing AI systems for the benefit of society while proactively analyzing and mitigating risks during the development and deployment processes.

One significant consideration is guarding against overbreadth in definitions. Regulation should focus on high impact decisions where AI plays a crucial role. Clear delineations must be established to distinguish between AI as a contributing factor in decision-making and instances where AI makes decisions without human review. By doing so, we can ensure that appropriate oversight is in place while avoiding unnecessary constraints on AI development.

Similarly, caution should be exercised to prevent overbreadth in implementation strategies. Human guardrails may be beneficial in certain cases, providing necessary checks and balances. However, it is essential to recognize that no single approach will always be correct. Flexibility is key when determining the level of human involvement, ensuring that the level aligns with the unique characteristics and requirements of each AI system.

Broad agreement exists among leading AI developers and researchers, including CCIA's members, that responsible AI development requires the following:

- ❖ Design for social benefit.
- ❖ Design to avoid unfair outcomes.

- ❖ Analyze and minimize risks as you design.
- ❖ Consider the risks to third parties from AI systems during design, but also the benefits.
- ❖ Use up-to-date safety, security, and privacy best practices.
- ❖ Monitor and govern identified risks in deployed systems.
- ❖ Provide appropriate disclosures for deployed AI systems.

While these principles may be expressed in different ways, any responsible AI framework will incorporate them. CCIA's members have engaged in responsible AI development, ranging from developing and applying their own responsible AI principles to conducting academic research that promotes privacy-by-design and the hardening of AI against motivated attackers seeking to extract training data, among other valuable contributions.

These high-level principles, applied in the context of any given application, provide the necessary flexibility to manage risks while providing the benefits AI can deliver. In high-risk applications, such as medical diagnostics, human supervision and significant disclosure of the AI would be appropriate; in lower risk applications, such as content moderation or video games, there may be little or even no need for human review.

AI Warrants Only Targeted Regulation Combined With Considered Application Of Existing Law

Rather than rushing to create new laws, it is essential to evaluate whether existing laws at the federal, state, and local levels adequately address the concerns posed by AI. In general, there should be little to no difference whether an act is performed by a person or by an AI system. This can be achieved by writing and applying law and regulation in a way that constrains outcomes, while maintaining neutrality as to the process by which those outcomes are created. For example, instead of creating a new law requiring AI systems to operate in a non-discriminatory fashion, existing discrimination laws should be applied to AI systems. By leveraging established legal frameworks, we can address these types of concerns without burdening the regulatory landscape with unnecessary redundancy. Using established legal frameworks and applying them evenhandedly to AI and human systems alike will also avoid regulatory arbitrage by ensuring there will be neither a legal advantage nor a disadvantage to operating a system as an AI system versus via human action.

The effective application of existing laws, such as intellectual property (IP) laws and product liability laws, will also address the vast majority of concerns that have prompted calls for the regulation of AI systems. Recent statements by officials from the FTC, DOJ, EEOC, and CFPB emphasize exactly this approach. These technologically neutral laws should be the first line of defense, addressing common legal issues when they arise in the context of AI applications. But where AI-specific distinctions exist, or when a failure of existing law emerges, new regulations tailored to that unique situation should be created.

Moving Towards A Risk-Based Framework For AI

Comprehensive regulation of AI should employ a risk-based framework rather than a prescriptive framework requiring specific mechanisms. National standards such as the NIST AI Risk Management Framework and international standards such as ISO/IEC 23894 and ISO/IEC 42001 may be relevant to refer to in the development of risk-based approaches. Policy-makers should focus on identifying and addressing the concerns associated with AI development and deployment. This approach empowers developers to find appropriate solutions within the defined limits while not limiting room for new technologies and experimentation.

The level of acceptable risk, required guardrails, and potential impacts should be evaluated based on the specific context. For applications with lower impact, higher tolerable risk levels and fewer guardrails may be acceptable. Conversely, applications with higher impact demand lower tolerated risk and more robust guardrails. This approach allows flexibility and adaptability, catering to the diverse nature of AI technologies.

Appropriate levels of transparency and disclosure are also crucial aspects of AI regulation. While they may not impact benefits or harms, they are essential to engendering trust in AI systems. People should have access to relevant information about how an AI system was designed and trained, as well as how it operates. This knowledge fosters accountability and user trust, enabling individuals to understand the basis of AI-driven decisions.

While transparency is important, it must be appropriate and relevant. Context is the key factor in determining the needed level of transparency, with riskier AI systems requiring higher levels and potentially more human involvement. An AI system that directs the movement of pallets in a warehouse should require significantly lower levels of transparency than an AI system that makes lending decisions. Additionally, protection of proprietary knowledge and confidential

business information is critical. Striking a balance between transparency and confidentiality is vital to promote investment in innovation while maintaining ethical and accountable AI practices.

Addressing Specific Issues That Have Received Attention

A. Determining responsibility for AI outputs

There are a number of different entities involved in any given AI system, including the provider who trained the AI model, the deployer who applies that model to a specific task, the compute provider who provides the hardware the AI system runs on, and the user who ultimately is utilizing the AI system. Basic legal principles of agency can serve as a starting point for determining responsibility. The developer, deployer, user, and compute resources involved in an AI system might each bear responsibility, depending on the circumstances.

Compute resources, typically acting as intermediaries or common carriers, should generally not be held responsible for AI outputs. On the other hand, trainers of a model may be held accountable if defects are inherent to the design of the AI system. For instance, if a model developer intentionally creates an AI that consistently ranks people of color as less creditworthy, they should bear responsibility for that, not just the operator of the system.

Similarly, operators of AI systems, while they may be generally responsible for the usage of the technology, should not be held liable for inherent design flaws or the actions of users if users can interact with the operator. For example, if a user instructs an AI to generate defamatory content, the operator should not be liable for that content.

This division of responsibility will ensure that liability lies in the most appropriate place, with the actor most capable of minimizing harm and most responsible for any harms that ensue.

B. Determining regulatory responsibility

While a governmental coordination role might be useful, creation of a new department or similar bureaucracy is likely to lead to regulatory duplication and stifle investment in and development of AI systems. In most cases, existing agencies responsible for specific areas of law are equipped to oversee regulation of AI that falls within their area of responsibility. Leveraging the



expertise and jurisdiction of these agencies will ensure a coherent regulatory landscape. For example, housing discrimination law would fall under the purview of agencies like the Department of Housing and Urban Development (HUD) or the Office of Fair Housing and Equal Opportunity (FHEO).

Similarly, coordinating the regulatory efforts and fostering industry development of best practices across various domains could be the role of the National Institute of Standards and Technology (NIST) or a similar entity; another potential model is the role of the IP Enforcement Coordinator in the IP ecosystem. Such coordination ensures consistency of the overall approach while allowing domain experts to ensure effective regulation of AI systems within their agency's expertise.