Submitted March 28, 2024

CCIA Submission to the United Nations Advisory Body on Artificial Intelligence's Interim Report[1]

After reviewing the Interim Report, please provide your feedback on the following sections:

Opportunities and Enablers (Maximum 3,000 characters):

CCIA strongly agrees with the opportunities presented by AI technology and its applications, as reflected by Paragraph 16.[2] Of particular interest to the UN is the current and potential power of AI to accelerate meeting the UN's 17 Sustainable Development Goals (SDGs) including addressing urgent issues of poverty, hunger, health care, equality in gender, clean energy, safe and clean water, economic inequality, and education by 2030. Further, AI technologies carry great promise to accelerate scientific progress in the coming years, if overly burdensome requirements and oversight measures are avoided.

This report offers a significant opportunity for the United Nations to put forward inclusive principles and promote institutional functions that serve both high-income nations and the broader Global South. Given the broad range of economic, infrastructure, and societal circumstances of the Member States to the United Nations compared to other governing bodies analyzing the issue of AI, such as the OECD and the G7, the final report should ensure that the principles promulgated are of utility to both countries seeking to untap the full potential of AI to solve long-standing challenges and countries that have already began to use AI in essential industries and everyday life. In particular, promoting access to government-generated information, and networks and technologies that enable its efficient distribution, is a key area where UN expertise can leverage AI's potential.

As such, the approach highlighted by the UN General Assembly's recently-approved resolution, "Seizing the opportunities of safe, secure and trustworthy artificial intelligence systems for sustainable development,"[3] should be instructive in the UN Advisory Body's efforts. Calls for UN Member States to expand "participation of all countries, in particular developing countries, in digital transformation to harness the benefits and effectively participate in the development, deployment and use of safe, secure and trustworthy artificial intelligence systems, including by

---

[1] "Interim Report: Governing AI for Humanity," United Nations Advisory Body on Artificial Intelligence (Dec. 2023) https://www.un.org/sites/un2.un.org/files/un_ai_advisory_body_governing_ai_for_humanity_interim_report.pdf (herein referred to as "Interim Report").

[2] Interim Report at 5. ("AI has the potential to transform access to knowledge and increase efficiency around the world. A new generation of innovators is pushing the frontiers of AI science and engineering. AI is increasing productivity and innovation in sectors from healthcare to agriculture, in both advanced and developing economies.").

[3] https://undocs.org/Home/Mobile?FinalSymbol=A%2F78%2FL.49&Language=E&DeviceType=Desktop&LangRequested=False.

capacity building relating to artificial intelligence systems, recognizing that promoting knowledge sharing activities and the transfer of technology on mutually agreed terms is an important aspect of building capacity, stressing the need to close the artificial intelligence and other digital divides; and increase digital literacy."[4]

Risks and Challenges (Maximum 3,000 characters):

As the Advisory Body develops Box 3 into a "risk assessment framework," it is critical that it center the concept of proportional risk, taking into account, as noted in Paragraph 40, that lack of common standards and benchmarks currently precludes precise risk evaluation. CCIA agrees that rather than seeking to list a comprehensive catalog of risks and attendant mitigations, beginning by identifying risk targets is a more useful approach, since mitigations will need to be target-specific to be effective.

Further, the risks detailed by the Advisory Body in its final report should incorporate the concept of varying responsibilities regarding such risks, by clearly distinguishing the differing roles of Developers, Deployers, and End Users, and the different responsibilities with respect to risks assigned to them. The specific harms that arise from the application of AI technologies must be clearly linked to the corresponding point in the AI deployment chain to ensure that harms and responsibility are not inappropriately assigned to the wrong party. The OECD's AI Principles and the report, "Advancing accountability in AI: Governing and managing risks throughout the lifecycle for trustworthy AI,"[5] provide helpful insight for developing metrics to measure risks throughout an entire AI technology application's Life-Cycle.

CCIA would recommend shifting the focus in Paragraph 34 away from "lack of transparency, access, compute and other resources, and understanding" towards an emphasis on testing AI models and applications.[6] The framing of this Paragraph could lead to the adoption of a broad set of norms or requirements that over-requests data that could undermine competitive advantage for businesses and privacy for individuals. Further, the impacts of AI technologies would be more apparent to governing bodies through testing than through over-inclusive demands for datasets. Datasets, algorithms, and other pieces of AI systems and services would not necessarily reveal how the AI model works in practice, it would only reveal how the AI model was developed. This would do little to mitigate harms that arise from the application of AI technologies.

CCIA agrees that a fragmented approach to AI governance would be problematic for both the oversight of the technology and its cross-border application, but also notes that introducing rigid and prescriptive governance would be both premature and ineffective in a rapidly evolving

---

[4] https://undocs.org/Home/Mobile?FinalSymbol=A%2F78%2FL.49&Language=E&DeviceType=Desktop&LangRequested=False.

[5] https://www.oecd-ilibrary.org/docserver/2448f04b-en.pdf?expires=1710529358&id=id&accname=guest&checksum=6FF5639062318D7DC89288FB83CC1DF0.

[6] The U.S. NIST has led work on methods to effectively test AI products and services through its AI Measurement and Evaluation Projects; https://www.nist.gov/ai-measurement-and-evaluation.

domain.[7] The current phrasing of Paragraph 35, "Mapping, avoiding, and mitigating risks will require self-regulation, national regulation, as well as international governance efforts," could be more of a hindrance if the threats of duplicative or inconsistent regulatory approaches are not made clearer.

CCIA strongly agrees with the points laid out in Paragraph 38 and appreciates their inclusion, and would urge the Advisory Body as it moves forward with this report and related initiatives to emphasize the potential danger of "missed uses."[8]

Guiding Principles to guide the formation of new global governance institutions for AI (Maximum 3,000 characters):

CCIA appreciates the issues regarding harms highlighted in Guiding Principle 2, Paragraph 49. However, as broad-based use cases are still in their infancy, AI products and applications are currently most effectively governed through adherence to high-level principles and voluntary commitments, with the adoption of binding obligations reserved for cases that are carefully calibrated and based solely on tangible harms and well-established risks, rather than potential future applications of the technology. As best practice, flexible and international technical standards-based approaches to the governance of AI are crucial to supporting AI innovation and diffusion. This is particularly important given the rapid development of AI and its nature as a general-purpose technology. Adopting overly-prescriptive rules while the technology still develops risks both slowing innovation becoming outdated quickly as global standards are themselves still in development, as are the real-world applications of the technology. As such, any "binding norms" should themselves be flexible enough to allow the technology to evolve and the use cases to develop while also addressing potential harms.

CCIA agrees with Guiding Principle 3 which highlights the importance of encouraging the "development of public data commons."[9] The Data Commons for the SDGs has proved to be a particularly helpful resource for advancing private sector projects to address the UN's SDGs,[10] and similar public data commons could provide the on-the-ground information needed to train burgeoning AI technologies to generate new solutions. This would be particularly helpful as a

[7] Interim Report at 10 ("Despite AI's global reach, governance remains territorial and fragmented. National approaches to regulation that typically end at physical borders may lead to tension or conflict if AI does not respect those borders. Mapping, avoiding, and mitigating risks will require self-regulation, national regulation, as well as international governance efforts. There should be no accountability deficits.").

[8] Interim Report at 11 ("Besides misuse, we also note countervailing worries about missed uses — failing to take advantage of and share the benefits of AI technologies out of an excess of caution. Leveraging AI to improve access to education might raise concerns about young people's data privacy and teacher agency. However, in a world where hundreds of millions of students do not have access to quality education resources, there may be downsides of not using technology to bridge the gap. Agreeing on and addressing such trade-offs will benefit from international governance mechanisms that enable us to share information, pool resources, and adopt common strategies.").

[9] Interim Report at 14 ("The development of public data commons should also be encouraged with particular attention to public data that is critical for helping solve societal challenges including climate change, public health, economic development, capacity building and crisis response, for use by multiple stakeholders.").

[10] https://blog.google/technology/ai/google-ai-data-un-global-goals/.

handful of U.S. companies have dedicated resources and investment to tackle the SDGs, as evidenced by the ministerial side event at the United Nations' 78th Session High Level Week on Sep. 18, 2023.[11] This report may also benefit from referencing the recent introduction in trade disciplines of the promotion of Open Government Data, such as through the USMCA, the APEC, and the OECD.[12]

CCIA strongly agrees with Guiding Principle 4.[13] It is crucial that as AI technology develops and governments seek to simultaneously leverage its potential and mitigate for emerging harms that the UN's efforts bring together industry, civil society, academia, and governing bodies. Such multi-stakeholder models of cooperation were central to the development and flourishing of the internet as a communication and commercial platform and should be informative as similar frameworks are developed for the AI age. Indeed, various cultural contexts for AI development and use should be instructive for UN activity, as some countries will be seeking to build capacity for AI development and deployment with much fewer resources than other jurisdictions that are already integrating AI technologies into standard commercial activities and everyday life.

Institutional Functions that an international governance regime for AI should carry out (Maximum 3,000 characters):

CCIA notes Function #3, where the UN could serve as a convener of nations to harmonize standardization—in this regard, CCIA would recommend that the Advisory Body emphasize work that is already being done in this field rather than potentially duplicate standards that could lead to a fragmentation of the technology and its use. For example, the International Organization for Standardization (ISO) has already published and continued developing standards within the field of AI, such as AI concepts and terminology,[14] the implications of AI governance for use by organizations,[15] guidance for AI risk management,[16] and AI management systems.[17]

CCIA agrees with Function #5, and believes the UN should leverage its ability to: "Promote international collaboration on talent development, access to compute infrastructure, building of diverse high-quality datasets and AI-enabled public goods for the SDGs." The UN can use its

[11] https://www.state.gov/artificial-intelligence-for-accelerating-progress-on-the-sustainable-development-goals-addressing-societys-greatest-challenges/.

[12] https://ustr.gov/sites/default/files/files/agreements/FTA/USMCA/Text/19-Digital-Trade.pdf; https://www.apec.org/docs/default-source/publications/2023/11/appendix-14-non-binding-principles-for-facilitating-access-to-open-government-data-in-the-apec-region.pdf?sfvrsn=ce3d72ca_2; and https://www.oecd.org/gov/digital-government/open-government-data.htm.

[13] Interim Report at 14 ("Such an AI governance framework can draw on best practices and expertise from around the world. It must also be informed by understanding of different cultural ideologies driving AI development, deployment, and use. Innovative structures within this governance framework would be needed to engage the private sector, academia, and civil society alongside governments. Inspiration may be drawn from past efforts to engage the private sector in pursuit of public goods, including the ILO's tripartite structure and the UN Global Compact.").

[14] https://www.iso.org/standard/74296.html.

[15] https://www.iso.org/standard/56641.html.

[16] https://www.iso.org/standard/77304.html.

[17] https://www.iso.org/standard/81230.html.

role as a convener of a broad set of countries to harness the resources and talents available on the global stage to spread the impact of AI systems and technologies in development to a wide range of contexts. By focusing on collaborative solutions at this institutional level, the UN can bring together stakeholders to work in tandem to develop and deploy urgently-needed AI technologies that can target long-standing challenges to realizing the SDGs in the next six years.

However, CCIA would recommend removing Function #7,[18] which refers to duties more applicable to individual governments than the UN, which should avoid wading into adopting binding enforcement mechanisms. As stated elsewhere, the UN can play a key role coordinating with organizations such as the OECD, the Partnership on AI, and the Council of Europe to align policies and guidelines that assess AI systems to help ensure effective and responsible use of AI technology. Such collaboration can subsequently influence the development of international norms and international standards (such as the ISO or IEC). The Internet Governance Forum can serve as a strong model to bring together the relevant stakeholders.

Other comments on the International Governance of AI section (aside from Principles and Functions, covered in above questions) (Maximum 3,000 characters):

CCIA agrees with the following piece of the interim report, from the conclusion, and would urge for the Advisory Body to apply this thinking throughout this report and follow-up workstreams: "To be effective, the international governance of AI must be guided by principles and implemented through clear functions. These global functions must add value, fill identified gaps, and enable interoperable action at regional, national, industry, and community levels. They must be performed in concert across international institutions, national and regional frameworks as well as the private sector." Broadly speaking, the UN Advisory Body should advance and promote multi-stakeholder approaches as it proceeds with this initiative, as buy-in from the industry and governing bodies will be paramount for these entities to adopt the report's principles. Proactive support from the private and public sectors will be necessary to ensure the conclusions and initiatives of the UN Advisory Body enjoy sufficient longevity.

CCIA would recommend adding to the "risks and challenges" section of this report barriers hindering the voluntary sharing of data. Such barriers include data localization measures, restrictions on cross-border data flows, and approaches to government data that do not allow for open access for AI researchers and developers. In particular, for AI technologies to train systems and models on representative datasets to address urgent (and often, locality-specific) problems such as those in healthcare, agriculture, and critical infrastructure, it is essential to enable access to data across the widest possible set of contexts. Further, encouraging open government data should be included as part of this report, as allowing AI systems and models access to government-held statistics regarding health and disease, infrastructure projects, agriculture yields and weather patterns, and educational outcomes would be helpful to building technologies to bridge divides in these issue areas as the SDGs aspire to do. This would also help bridge the gap between data representation between higher-income countries

[18] Interim Report at 18 ("We cannot rule out that legally binding norms and enforcement would be required at the global level.").

and the "Global South," where there are fewer tech startups engaging in the same innovation as richer markets.

Any other feedback on the Interim Report (Maximum 3,000 characters):

CCIA appreciates the opportunity to offer feedback on the UN Advisory Body on Artificial Intelligence's Interim Report, which is an important beginning to establishing alignment on this topic among UN Member States. For overall input on how AI fits into international trade rules and cooperation facilitating cross-border delivery of services, please refer to CCIA's White Paper, "Trade Principles for AI,"[19] linked to here, and attached to the online version of this submission.

Further, as the Advisory Body pursues governance principles, CCIA would generally urge for the incorporation of specific definitions of and delineations of responsibilities between developers, deployers and end users, as previously mentioned. International approaches such as the Hiroshima AI process should be instructive in this regard. For a broader overview of best practices to pursue in devising governance of AI, please refer to CCIA's White Paper, "Understanding AI: A Guide to Sensible Governance,"[20] linked to here, and attached to the online version of this submission.

[19] https://ccianet.org/library/trade-principles-for-ai/.
[20] https://ccianet.org/library/understanding-ai-guide-to-sensible-governance/.

Computer & Communications Industry Association

Open Markets. Open Systems. Open Networks.

ccianet.org

Rules of the Road

Trade Principles for a Competitive Global AI Market

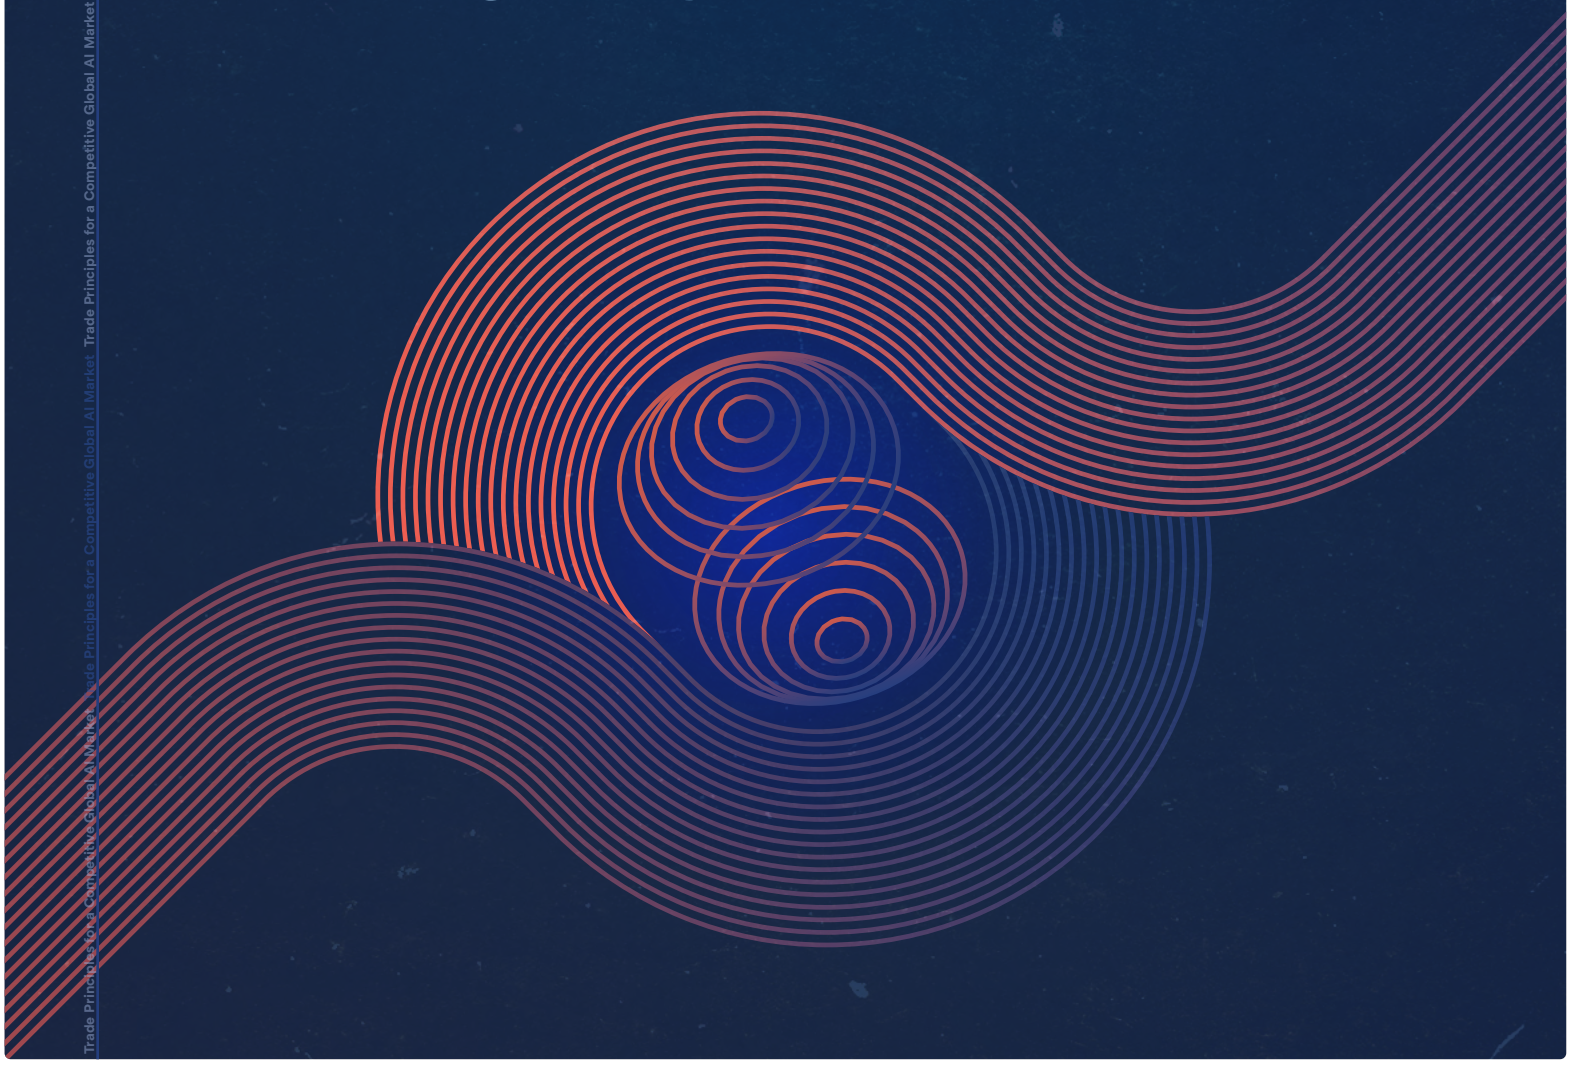

Executive Summary

Hardly a day goes by without new evidence of AI's potential to increase productivity, the long-term source of an economy's ability to generate broad-based growth and lay a foundation for sustainable wage gains. With most advanced countries, including the United States, facing the demographic challenge of declining working-age populations, such technological advances will be a key source of maintaining a high standard of living.

For the United States and other countries' ability to fully leverage these capabilities, however, the technology needs to scale—both domestically and globally—and thus be able to be fully integrated into the breadth of sectors where its benefits can take root.

A key obstacle to such scalability is the risk of a fragmented global market. One source of such fragmentation is skepticism about a technology when all risks have yet to be identified or mitigated—a challenge of any new technology, but one that thoughtful policymakers are well-positioned to address. But a more intractable source of fragmentation is the patchwork of market restrictions that have hindered digital trade generally. This occurs in foreign markets where, as with many internet applications generally, incumbents have often focused on slowing down a competitive threat or protecting local market advantage, rather than embracing openness to technological opportunity. Addressing such restrictions should be a priority if the U.S. interest in advancing AI is to succeed. Many of the same rules that have helped safeguard an open digital ecosystem to date, now time-tested, if not broadly adopted, will be key to the success of AI as well.

In this environment, navigating the trade impact, both in terms of addressing barriers countries might seek to erect as well as promoting supportive policies enhancing its benefits (including through binding and enforceable trade rules), will be core to expanding U.S. economic interests and those of our close partners.

Introduction

With recent rapid advances in AI technology, and the recognition of its transformational potential across almost every economic sphere, investment in AI has surged: new capital invested in 2022 is estimated to be over 90 billion dollars, more than half of which is from the United States.[1] Motivating this intense interest is growing evidence for how AI can be integrated into myriad use cases contributing to enhanced economic and welfare gains, across manufacturing, services, and agriculture.

While the rapid advances in AI have engendered vigorous debate about potential risks, appropriate regulation, and how to address inevitable labor market dislocations, the promise of this technology now appears undisputed: in addition to potential benefits to humanity in addressing challenges in areas as diverse as healthcare, climate and agriculture, AI also has the potential to usher in unprecedented productivity gains. This latter attribute is particularly relevant for countries' competitiveness, both domestically and internationally: without such gains, countries such as the United States, recently suffering from prolonged declines in productivity growth[2] may fail in generating broad-based, sustainable growth that can support long-term wage gains. If the Biden Administration's pursuit of a worker-centered trade policy is to have any long-term meaning, such productivity gains, addressing one of the key critiques of trade policy to date (trade's contribution to wage stagnation), should take center stage. Accordingly, while most policy debate has revolved around domestic responses to the anticipated effects of the technology, the intersection between domestic policy responses and core trade principles merits attention.

The stakes in this policy arena could not be higher. On the one hand, they involve strategic competition with China, whose leadership has set 2030 as a target for China becoming the world's "primary" AI innovation center[3] through aggressive governmental support. On the other hand, the EU's attempt to seize the opportunity in quickly instituting a comprehensive, top-down governance model in the forthcoming AI Act, another notch in its belt as a presumptive "regulatory superpower."

Apart from scale (an internet user base of over one billion served by some of the world's biggest digital companies), many of China's advantages in generating data are not ones democracies would emulate—e.g., over a billion surveillance cameras, and the world's most extensive system for monitoring speech on the

1 Stanford University, Artificial Intelligence Index Report 2023, https://aiindex.stanford.edu/wp-content/uploads/2023/04/HAI_AI-Index-Report_2023.pdf.

2 See Bureau of Labor Statistics, The U.S. Productivity Slowdown: An Economy-Wide and Industry-Level Analysis (April 2021), https://www.bls.gov/opub/mlr/2021/article/the-us-productivity-slowdown-the-economy-wide-and-industry-level-analysis.htm.

3 See Graham Allison and Eric Schmidt, Is China Beating the U.S. to AI Supremacy? (Aug. 2020), https://www.belfercenter.org/publication/china-beating-us-ai-supremacy#footnote-046-backlink.

internet, a key focus of China's early AI efforts.[4] In the case of the EU, however, the lack of domestic players has led to a focus on prescriptive regulation that appears aimed not only on addressing potential social risks but also blunting competitive challenges, rather than incubating productivity growth responses. This may well contribute to Europe falling further behind in the development and use of AI. In contrast, many other countries, particularly in the Indo-Pacific, are looking to institute policies that, while cognizant of risks, also look to ensure that they can attract similar investment and participate in this transformation.

The Economic Promise of AI

AI is not a new phenomenon, having been the subject of research and development for decades, with "deep learning" and the development of "expert systems" being advanced four decades ago by noteworthy pioneers.[5] But the confluence of interrelated developments (mainly, breakthroughs in machine learning models, computing power that has grown exponentially in the past 5 years,[6] and the unprecedented generation and availability of training data[7]) has propelled AI to the forefront of global attention with its commercial impact now undisputed. This revolution is now being led by private companies, whose massive investments and fierce competition has contributed to both scale of innovation and speed of development, a significant change from just a decade ago when advances were concentrated in academia.[8]

The commercial opportunities appear vast. McKinsey recently estimated that one subset of AI alone, generative AI, will contribute across 63 use cases up to $4 trillion in added value annually to the global economy in the next decades across all sectors—with a particular impact on banking, high technology and life sciences, in areas as diverse as customer support, marketing and sales, and software development. [9] McKinsey estimates that this evolution, if successful, could enable labor productivity growth of 0.1 to 0.6 percent annually through 2040, a potentially extraordinary achievement.[10]

4 As summarized by RAND, "China has an advantage over the United States in the area of big data sets that are essential to the development of AI applications. This is partly because data collection by the Chinese government and large Chinese tech companies is not constrained by privacy laws and protections." RAND, Maintaining the Competitive Advantage in Artificial Intelligence and Machine Learning (2020), https://www.rand.org/pubs/research_reports/RRA200-1.html at 1.

5 See The History of Artificial Intelligence, Harvard University Science in the News Blog (Aug. 28, 2017), https://sitn.hms.harvard.edu/flash/2017/history-artificial-intelligence/.

6 See Stanford University, Artificial Intelligence Index Report, supra note 1 at 56.

7 Statista, Volume of data/information created, captured, copied, and consumers worldwide (Jun. 2021), https://www.statista.com/statistics/871513/worldwide-data-created/.

8 Since 2014, 32 of 35 significant models have been developed by companies. Stanford University, Artificial Intelligence Index Report, supra note 1 at 50.

9 McKinsey, The Economic Potential of Generative AI: The Next Productivity Frontier (Jun. 2023), https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-AI-the-next-productivity-frontier#/.

10 Id. at 45.

While arguably dominant, the United States is certainly not alone in this transformation, with its key geopolitical rival also targeting this technology as a strategic imperative for both economic as well as less benign reasons. Nonetheless, U.S. strengths are compelling. As of now, most of the foundation models (large language models, and multimodal models) were developed in the U.S; U.S. research, based on citations, is unrivaled; U.S. hardware, particularly in chip and system design and deployment of cloud computing is unparalleled; and U.S. industries have been global leaders to seek to integrate AI advances into a wide variety of use cases, new case for which appear daily.[11]

However, scalability, the precondition of a successful transition to AI-enhanced economic activity, will often require a global footprint, since even smaller suppliers seeking to deploy niche applications may require a customer base beyond any one country to justify the large capital expenditures many AI applications will require. Cataloged below are a number of the challenges to that path, and the trade rules that can help address these challenges.

Why Trade Matters

As with digital technology generally, a big part of the promise of AI is its inherent scalability, paralleling the software applications and services that currently thrive on the internet. This refers to the ability to quickly spread and integrate into a mature global digital ecosystem that links communications, computer power and software applications with businesses and consumers throughout the world. This characteristic, where high and/or risky capital investments can be recovered through a globally-addressable market, is at the heart of why reasonable trade rules are inextricably related to the potential success and possible constraints on this technology: without predictable, consistent rules, scalability founders. And, such scalability is not just with respect to the companies directly offering AI applications, but also with the companies integrating AI into their traditional businesses. For them as well, whether an airline, bank, automobile manufacturer or drug developer, often with a global footprint, integrating AI across this footprint is critical to making the investment worthwhile.

Scalability captures, in particular, both the promise and challenge this technology represents for small- and medium-sized enterprises (SMEs), the "unsung heroes" of the digital economy.[12] SMEs are potentially one of the AI's biggest beneficiaries. Their participation in the digital economy has already been revolutionized by ready availability of globally-accessible cloud computing power, (one of the foundations of AI) and AI offers unrivaled opportunities for these

11 AI Is So Hot Even KFC and Williams-Sonoma Execs Are Talking About It, Wash. Post (Aug. 24, 3023), https://www.washingtonpost.com/technology/2023/08/24/ai-corporate-hype/.

12 Digital Technology: The Unsung Hero of Small Business, Disruptive Competition Project (Jul. 10, 2023), https://www.project-disco.org/innovation/digital-technology-the-unsung-hero-of-small-businesses/.

smaller players to expand through the use of technology once restricted to larger companies.[13] In fact, adoption of AI by SMEs in the United States is impressive. A recent survey by a major marketing firm reports that 91 percent of surveyed SMEs claim use of AI made their business more successful, cutting costs, time, avoiding mistakes, and helping their businesses grow.[14] For such users, for whom AI can enhance their export competitiveness, integrating the technology necessitates extending its use beyond the domestic market, and thus curbs on AI in destination markets will curb these SMEs growth.

Apart from users, a key feature of the current marketplace is SMEs themselves as developers—contributing to the start-up renaissance that AI has fostered.[15] For both sets of players—SME deployers and developers—trade rules may be even more important than for bigger companies, as the burdens of navigating fragmented global markets with inconsistent rules may be, for such smaller players, insuperable.

The impact of AI and its effect on trade may well spawn new approaches on trade rulemaking, as evidenced in some early efforts to address regulatory challenges in the Digital Economy Partnership agreement concluded recently between New Zealand, Chile, and Singapore.[16] But even before looking to create novel rules, ensuring that existing frameworks apply to AI may be equally if not more important to developing a sustainable framework for trustworthy growth, by ensuring that obvious frictions are minimized. Many specific trade provisions that will be critical to ensuring that AI applications and services thrive globally are well established and it is their expansion that may be the most important first step in advancing useful guardrails. Such provisions include:

- Ability to move data into and outside a jurisdiction (cross-border data transfer rules);
- Ability to rely on computing facilities outside a specific jurisdiction;
- Protections against unwarranted disclosure and transfer of commercially sensitive resources (source code and algorithms);
- Affirmation of copyright exceptions and limitations critical for machine learning;

13 As noted by the OECD, "SMEs can source external AI expertise and solutions from knowledge markets that typically compensate for a lack of internal capacity. Cloud computing-based Software as a Services (SaaS) and Machine learning as a Service (MLaaS) offer advantages such as the scalability of AI solutions and costs, no prerequisite of technical knowledge (for SaaS), digital security features directly embedded in the software." OECD, Digital Transformation of SMEs (2021), https://www.oecd-ilibrary.org/sites/01a4ae9d-en/index.html?itemId=/content/component/01a4ae9d-en.

14 Constant Contact Research Reveals Small Businesses Who Use AI Are More Likely Save Money And Be Successful (Aug. 9, 2023), https://www.prnewswire.com/news-releases/constant-contact-research-reveals-small-businesses-who-use-ai-are-more-likely-to-save-money-and-be-successful-301896332.html

15 40 Growing AI Companies & Startups in 2023 (Oct. 19, 2023), https://explodingtopics.com/blog/ai-startups.

16 See article 8.2 of the Agreement, available at: https://www.mfat.govt.nz/assets/Trade-agreements/DEPA/DEPA-Signing-Text-11-June-2020-GMT-v3.pdf.

- Reliance on international standards rather than country-specific technical requirements;
- Good regulatory practices in the development of obligations applicable to AI developers and implementers; and
- Non-discriminatory treatment of service suppliers (national treatment).

These rules will not solve the larger questions of what form of oversight should be applied to AI, particularly in areas deemed threats to human safety or inconsistent with core social values. But in many cases, these rules can contribute to a more trustworthy framework, helping clarify where government intervention is and is not beneficial. And where consensus emerges that a restriction is justified, such rules, common in many existing trade agreements, include exceptions, that provide for meaningful discretion in countries' need to address requirements based on local conditions or values. Importantly, trade rules provide a baseline of accountability for governments and clarity for suppliers that enable an expansion of trade and the benefits that accrue to both exporters, importers, and consumers.

Trade Rules Relevant to AI

1. Cross-Border Data Flow Rules

At the core of AI is the ability to discern patterns in varied data sets and transform learned correlations into predictive or generative outputs. One basis for AI's recent and rapid advance has been the accelerated digitalization of the economy, creating large data repositories that, subject to advanced modeling and unprecedented computing power, allow increasing accuracy and relevance of AI-generated outputs.[17] Since relevant data is not restricted by geography (and in many cases requires inputs from global sources to be comprehensive— e.g., for text and speech recognition, cybersecurity, health, climate, weather, etc.) the ability to move data cross-border is fundamental: both to train models and to interact with them once trained. And many AI models (e.g., adaptive models) are not static, but are constantly updated based on real-time feedback. The richness of cross-border data flows, accordingly, will have a significant impact on the quality, relevance and utility of many AI applications.

Largely parallel to the growth of the internet, the recent growth of such flows has been remarkable, with U.S. cross-border data flows alone, based on one representative metric (on submarine cables) tripling over the past decade from 2,000 to 6,000 petabytes per month.[18]

17 According to IBM, 90% of the world's data was created in the past two years, and this quantity continues to double every two years. See https://www.ibm.com/case-studies/hsbc-usa#.

18 Equinix, New Subsea Cable Architecture Are Carrying the World's Traffic (Mar. 16, 2020), https://blog.equinix.com/blog/2020/03/16/new-subsea-cable-architectures-are-carrying-the-worlds-traffic/.

However, this robust cross-border exchange is not assured. Increasingly, governments are seeking to restrict what can be exported from their territory (and in some cases, notably China, what can be imported) subject to an increasing array of restrictions.[19] Such measures, while clearly hurting any cross-border supplier of an AI-enabled service, will also greatly handicap the importing country and its users, since best-in-class technology might only be available on a cross-border basis. Although there have been efforts to minimize the need to move data outside of specific locations when conducting training,[20] which could help in the processing of sensitive data such as health data, such approaches may involve additional costs and performance trade-offs, and are unlikely to fully address the need for robust cross-border data flows.

Trade rules facilitating the ability of companies to move data, subject to reasonable safeguards, are not new and preceded any focus on AI. For example, recognizing that cross-border financial services are practically impossible without the movement of data, a core group of WTO members conclude an addendum to the General Agreement on Trade in Services (the GATS) in 1994 called the Understanding on Commitments in Financial Services that guaranteed financial service suppliers the ability to move data between the territories of signatory members.[21] It was not until negotiation of the Trans-Pacific Partnership that this rule was extended to other sectors,[22] but since then it has been a standard feature of high-standard trade agreements.

Importantly, this trade rule does not prohibit regulation of cross-border data transfers. Rather, it ensures that conditions attached to its transfer (e.g., to protect privacy or security) are reasonable, proportionate, and justified and can accommodate a range of transfer mechanisms. As such, this kind of rule is a foundation for a predictable framework for data-intensive industries like AI by enabling them to function globally, while allowing for accompanying safeguards to protect consumers, companies and governments against the threat of unwarranted access.

19 The EU, which has long had a restrictive data export regime based on privacy rationales, has now sought to expand restrictions into non-personal data, under its proposed Data Act; China has progressively tightened the nature and scope of data that can be exported; India, Vietnam, and many other countries are considering analogous restrictions. See ITIF, How Barriers to Cross-Border Data Flows Are Spreading Globally, What They Cost, and How to Address Them (2021), https://itif.org/publications/2021/07/19/how-barriers-cross-border-data-flows-are-spreading-globally-what-they-cost/;OECD, Digital Services Trade Restrictiveness Index, https://goingdigital.oecd.org/en/indicator/73.

20 E.g., the so-called "federated" learning model, a distributed, decentralized approach where data remains dispersed and only iterations to a model are combined centrally. See https://research.ibm.com/blog/what-is-federated-learning.

21 See Article 8 of the Understanding, available at: https://www.wto.org/english/tratop_e/serv_e/21-fin_e.htm.

22 See Article 14.11, Cross-border Transfer of Information by Electronic Means, available at: https://ustr.gov/sites/default/files/TPP-Final-Text-Electronic-Commerce.pdf.

Democratic countries can and do integrate such safeguards into their regulations on how data is treated. Cross-border data flow rules are consistent with such practices, as demonstrated by the numerous trade agreements including such provisions. This is, in fact, a critical factor distinguishing rule-of-law countries with authoritarian regimes, and a basis on which global alliances promoting trustworthy data flows can be built.

China has significant advantages in developing AI based on the sheer quantity of data its firms can access, due in no small part to its extensive surveillance practices and weak rule of law. The comparative advantage democracies can play in promoting AI, however, is their ability to work together to pool geographically diverse datasets that result in globally representative training data, making the resulting AI systems more resilient and less prone to bias or cultural/demographic errors. Accordingly, this is a critical moment to both address a geopolitical competition and lay a foundation for how trustworthy data flows are possible—by re-committing to a rule ensuring that data can flow on a cross-border basis among likeminded countries.

2. Location of Computing Facilities

One of the hallmarks of some of the most promising AI applications is their reliance on unprecedented computing power, both in processing data for training models and generating specific outputs for consumers and businesses once the model is mature.[23] This growth of computing power has encompassed both the number and capability of processing units, resulting in dedicated systems that are now at the cutting-edge of computing,[24] powering models that now incorporate hundreds of billions of parameters.[25] Since such massive computing power is not equally distributed around the world, the computing resources and human expertise supporting AI (both in training and implementing a model) will inevitably be concentrated, in the near-to-medium term, in a limited number of geographic locations.[26] Compounding this is the race to design and deploy the most advanced chipsets, demand for which has limited their availability. Accordingly, countries that require that computer processing and storage for specific applications be done locally will undermine their ability to participate in the training and implementation of relevant applications—to the detriment of foreign suppliers and their own economic development.

23 Center for Security and Emerging Technology, AI and Compute (Jan. 2022), https://cset.georgetown.edu/wp-content/uploads/AI-and-Compute-How-Much-Longer-Can-Computing-Power-Drive-Artificial-Intelligence-Progress.pdf ("Between 2012 and 2018, the amount of computing power used by record-breaking artificial intelligence models doubled every 3.4 months.").

24 Id. at 7.

25 A recent model, PaLM, advertises 540 billion parameters. Google Research Blog, Pathways Language Model (PaLM): Scaling to 540 Billion Parameters for Breakthrough Performance (Apr. 4, 2022), https://blog.research.google/2022/04/pathways-language-model-palm-scaling-to.html.

26 Stanford University, Artificial Intelligence Index Report 2023, supra note 1 at 51.

While hosting data centers, even if uneconomic,[27] has taken on aspects of national pride akin to steel mills in the last century, the high capital commitment and pace of technological development makes it unlikely that such policies will help their development, and in fact are likely to do the opposite. Large economies, such as China and the EU may have sufficient scale to attract the necessary investment (domestically or from abroad) but even for such markets, a localization mandate is bound to negatively affect growth. While bigger companies may conclude that they cannot afford to bypass such markets and may submit to such restrictions, the inevitable limiting of smaller players will mean that some of the most innovative services and applications may not be available there—potentially stymying innovations in such markets. In short, expanding, through trade rules, the principle that governments should not mandate use of local facilities is critical to ensuring that the benefits of AI can be distributed globally.

3. Protection of Source Code and Algorithms

Computing and software development has long benefited from open-source development, an approach that has stimulated broad ecosystems of co-developers and sparked an untold amount of innovation. In the earlier days of AI, the academic involvement using an open-source approach was the norm. Even now several major foundational models (e.g., Meta's LLaMA) are open source. On the other hand, for many companies, significant investment in AI is based on the goal of offering a differentiated product whose design is its competitive advantage, a business model that can also spur innovation. Accordingly, when mandated disclosure of source code, and embedded algorithms can result in competitors (or a government) appropriating that advantage, incentives to invest and innovate will be diminished.

A trade rule protecting source code and algorithms from disclosure builds on a general consensus that trade secrets should generally benefit from protection, and was first introduced as a specific trade rule by Japan in the Trans-Pacific Partnership Agreement (TPP). The motivation of the rule was largely based on experience in China, where a combination of mandated disclosure (e.g., under the equipment certification program called the Multi-Level Protections Scheme, MLPS) coincided with widespread alleged misappropriation, as borne out by cases involving companies such as varied as Motorola, Cisco, Google, Sinovel, and Tesla.

Given the competition with China in AI and the value of preserving the option of proprietary competitive advantages, any disclosure mandate could put U.S. firms at a competitive disadvantage globally. Any such trade rule will require exceptions, to address cases where evidence of illegal activity may

27 ECIPE, The Costs of Data Localisation: A Friendly Fire on Economic Recovery (2014), https://ecipe.org/publications/dataloc/.

merit investigating whether behavior was coded into the product (e.g., in the Volkswagen emissions scandal). Such a rule does not preclude robust testing and certification of products, based on the broad-based consensus that testing can be accomplished without access to source code.[28] Neither does such an approach conflict with the view that trustworthy AI systems should incorporate robust "explainability," so regulators and consumers understand the basis of automated decision- making. But explainability of a system need not include how complex algorithms are coded in software, disclosure of which is unlikely to advance that goal.

4. Reasonable Exceptions and Limitations in Copyright Regimes

A significant portion of data used to train AI models is protected by copyright, meaning that some uses of that data are restricted by copyright law, but limitations and exceptions apply. In the United States, courts have found that, to the extent that training infringes any of the uses restricted by copyright law, the mass copying of raw material to build databases for uses by AI processes is permitted under fair use. Israel's Ministry of Justice recently issued an opinion that its fair use provision, modeled on U.S. law, permits the copying of works for AI training purposes.[29] Further, the EU,[30] Singapore and Japan[31] have adopted provisions on text and data mining under their copyright laws, which would permit AI training. These provisions are all consistent with existing international IP law, which provides adequate flexibility to support both AI developers and rightsholders. Nonetheless, additional trade provisions designed to either explicitly permit AI training or to ensure that relevant exceptions and limitations are consistently maintained[32] could be helpful in maintaining a predictable legal environment for the growth of AI.

Relatedly, there have been efforts by rightsholders to bolster their ability to monetize content by seeking to impose limits on its use as training data. Since such goals can be accomplished contractually through terms of use, and by use of technical tools like robots.txt to prevent unauthorized online access and use, additional AI-specific intellectual property rights do not appear justified and should not be contemplated in trade rules.

28 See NIST, Facial Recognition Vendor Test (FRVT), https://www.nist.gov/programs-projects/face-recognition-vendor-test-frvt (designed to detect bias, which does not require access to source code).

29 Israel Ministry of Justice Issues Opinion Supporting the Use of Copyrighted Works for Machine Learning, Disruptive Competition Project (Jan. 19, 2023), https://www.project-disco.org/intellectual-property/011823-israel-ministry-of-justice-issues-opinion-supporting-the-use-of-copyrighted-works-for-machine-learning/.

30 Articles 3 and 4 of EU Directive on Copyright and Related Rights in the Digital Single Market, available at https://eur-lex.europa.eu/legal-content/en/TXT/HTML/?uri=CELEX:32019L0790.

31 European Alliance for Research Excellence, Japan Amends Its Copyright Legislation to Meet Future Demands in AI and Big Data (2018), http://eare.eu/japan-amends-tdm-exception-copyright (summarizing and explaining Copyright Act 2018); text of legislation is available at http://www.mext.go.jp/b_menu/houan/kakutei/detail/1405213.htm.

32 For example, CP-TPP's Article 18.66 includes this helpful provision: "Each Party shall endeavour to achieve an appropriate balance in its copyright and related rights system, among other things by means of limitations or exceptions that are consistent with Article 18.65 (Limitations and Exceptions), including those for the digital environment." Available at https://ustr.gov/sites/default/files/TPP-Final-Text-Intellectual-Property.pdf.

5. Reliance on International Standards and Conformity Assessment

As governments begin to regulate AI, particularly for uses deemed high-risk (i.e., uses that can significantly impact health or safety, or affect individuals' legal rights), consistency of approach will be critical to ensuring both the global acceptability of specific models and applications, and a consistent and effective mitigation of potential harms. While high-level principles for trustworthy AI have begun to emerge and gain global acceptance (e.g., the OECD's AI Principles[33]) and countries have begun to institute more granular frameworks (e.g., NIST's AI Risk Management Framework[34]), efforts to create detailed technical standards critical to achieve regulatory goals have only recently begun. Nevertheless, these efforts are well underway, mobilizing broad-based expertise in finding consensus approaches to addressing core issues. Some existing, mature standards, such as the International Standards Organization (ISO/IEC) 27001 family of standards for cybersecurity are directly relevant to AI systems. Other international standards development that is helping to building consumer trust and regulatory acceptance of the use of AI includes the following:

- ISO/IEC 42001 AI Management System (AIMS) (the first AI standard to be used for product certification)
- ISO/IEC 42005 AI System Impact Assessment
- ISO/IEC 25012 Data Quality
- ISO/IEC 42006 Certification Body Requirements for AI
- ISO/IEC 27090 AI System Security
- ISO/IEC 27091 Privacy Protection for AI
- ISO/IEC 6254 Explainability
- ISO/IEC 23894 Risk Management for AI
- ISO/IEC 17866 Guidance for mitigating ethical and societal concerns for AI
- ISO/IEC 12791 Treatment of Unwanted Bias in AI systems
- IEEE Adaptive Instructional Systems (AIS) portfolio
- MLCommons Training and Inference benchmarks

While many of these standards are still under development, the fact that consensus standards development fora have mobilized their expertise and resources to address this breadth of issues provides a clear path to consistent, globally-applicable outcomes, and the possibility of avoiding trade-restrictive fragmentation of regulatory requirements. As with standards already covered

33 OECD AI Principles, https://oecd.ai/en/ai-principles.

34 NIST, AI Risk Management Framework, https://www.nist.gov/itl/ai-risk-management-framework.

by the WTO's technical barriers to trade agreement (TBT), the existence, or imminent completion of a global consensus-based standard provides a legal basis for a preferred alternative to country-specific requirements, proliferation of which could deal a major blow to the ability of AI applications to scale globally. And, as noted earlier, global standards are one of the only ways smaller companies and smaller countries can navigate a path to global relevance, critical where risky investment is at stake.

6. Good Regulatory Practices

AI applications are already prevalent in many regulated industries, and are beginning to be integrated into mandatory conformity assessment procedures. For example, the Federal Drug Administration (FDA) has long regulated certain software as a medical device. The FDA has now developed a specific program for AI and machine-learning (AI-ML) enabled devices which has already approved 178 applications.[35]

Leveraging the domain expertise of existing regulators, rather than creating general AI-specific rules, may be the most effective way to address needs reflected in their existing mandate (e.g., to ensure safety, privacy, fairness etc.). Given this, having open, transparent, and accountable processes, with broad-based stakeholder input in the development of regulations is particularly important when a new and rapidly-evolving technology is involved. Trade rules have recently begun focusing on the importance of consistent procedures in the development of regulations and these are particularly relevant to AI. Accordingly, rules incorporating such practices,[36] (e.g., Good Regulatory Practices chapters of TPP and USMCA, and under negotiation in the U.S.-Taiwan Initiative and the Indo-Pacific Economic Partnership) are a significant step forward and should be encouraged.

7. National Treatment of Service Suppliers

Given the nature of the internet, digital services, including AI-enabled services, will be generally available wherever internet access is a reality. Nevertheless, whether through discriminatory standards or a perceived need to promote local suppliers at the expense of competing foreign services, trade-restrictive measures remain a constant threat. AI-enabled services will generally benefit from existing commitments to national treatment, where available, given trade partners' general acceptance of the technologically-neutral nature of trade commitments. However, gaps in coverage in many countries remain, and a temptation to characterize an AI-enabled service as

35 FDA, Artificial Intelligence and Machine Learning-Enabled Medical Devices, https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-aiml-enabled-medical-devices.

36 See, e.g., USMCA's Good Regulatory Practices chapter, available at: https://ustr.gov/sites/default/files/files/agreements/FTA/USMCA/Text/28_Good_Regulatory_Practices.pdf.

novel and outside the scope of existing commitments means that expanding such commitments, ideally on a "negative-list" approach (focusing on scheduling exceptions, rather than underlying services), should remain a long-term goal for the support of AI. This is particularly important since the prospect of countries seeking to preference a national AI model, or national AI-enabled applications is a distinct possibility.

One basis for doing so would be to make arbitrary distinctions between models or applications that serve as a proxy for nationality. Such an approach could be facilitated by proposals currently under consideration in the EU,[37] for example, that seek to impose additional regulatory burdens on AI models based on the computing power necessary to develop or implement the model, or the size of data sets used to train the model. Unless well-grounded in demonstrable risks relating to specific use cases, such categories could easily be designed to simply target models of disfavored countries, an outcome that national treatment obligations should discipline—i.e., to ensure that differential treatment was not arbitrary or discriminatory, and necessary for the purpose of protecting against AI harms.

37 See page 19 of Annex 1 of the draft AI Act at https://table.media/europe/wp-content/uploads/sites/9/2023/10/2023-10-17-conseil-ia-mandat-de-negociation-10412dc9fadd4e4fa9b0360960fd13af.pdf.

Computer & Communications Industry Association

Open Markets. Open Systems. Open Networks.

ccianet.org

2023

Understanding AI

A Guide To Sensible Governance

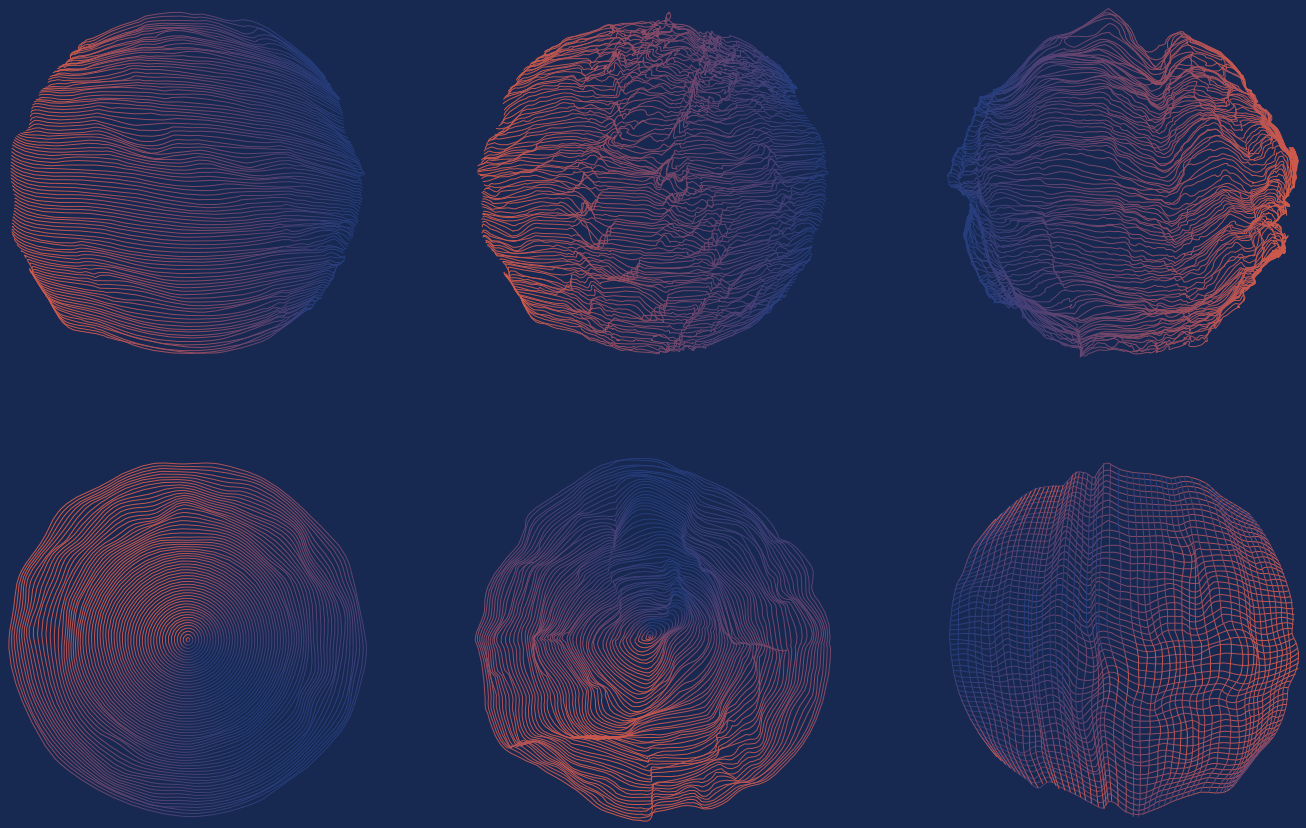

Executive Summary

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.