

EAC-PM/WP/26/2024

A Complex Adaptive System Framework to Regulate Artificial Intelligence



January 2024

Sanjeev Sanyal, Pranav Sharma, and Chirag Dudani

Contents

Executive Summary	3
I. Introduction	5
II. Policy Analysis.....	8
a. United States.....	9
b. United Kingdom	12
c. The European Union.....	14
d. China	17
e. The United Nations	19
III. AI: A Complex Adaptive System.....	20
IV. Contours of AI Regulation: A Complex Adaptive System Approach.....	21
a. Instituting Guardrails and Partitions to ‘Prevent Wildfire’	21
b. Ensuring Human Control through Manual ‘Overrides’ and ‘Authorization Chokepoints’	22
c. Transparency and Explainability	22
d. Distinct Accountability	23
e. Specialized, Agile Regulatory Body	24
V. Financial Markets as a Complex Adaptive System.....	26
Conclusion	27

Sanjeev Sanyal is Member, Economic Advisory Council to the Prime Minister (EAC-PM), Pranav Sharma is Consultant (Science Diplomacy), Indian National Science Academy, and Chirag Dudani is an Assistant Consultant, Economic Advisory Council to the Prime Minister (EAC-PM). The contents of the paper including facts and opinions expressed are sole responsibility of the authors. EAC-PM or the Government of India does not endorse the accuracy of the facts, figures or opinions expressed therein. The authors would like to thank Dr. Bibek Debroy (Chairman, EAC-PM), Kris Gopalakrishnan (President, Infosys Science Foundation), Ashutosh Sharma (Professor, IIT-Kanpur), Uday B. Desai (Professor Emeritus, IIT-Hyderabad), Santanu Chaudhury (Director, IIT-Jodhpur), Balaraman Ravindran (Professor, IIT-Madras), Joanna Bryson (Professor of Ethics and Technology, Hertie School), Sanghamitra Bandyopadhyay (Director, Indian Statistical Institute), and Rahul Matthan (Partner, Trilegal).

Executive Summary

The field of Artificial Intelligence (AI) has been making significant strides, with advancements in computer vision, language processing, robotics, and many other emergent areas. This rapid development, however, has raised substantial concerns about the risks associated with uncontrolled AI propagation. The increasing capabilities of AI systems bring potential dangers, such as the prospect of "runaway AI" (under-controlled AI), where systems might recursively self-improve beyond human control and misalign with human welfare. This and other more substantiated concerns have led to a growing demand for regulatory frameworks that ensure AI development and deployment are conducted safely, ethically, and transparently, safeguarding against risks like systemic breakdowns, loss of privacy, national security threats, and compromised critical infrastructure.

The current regulatory approaches, which typically rely on ex-ante impact analysis and risk assessment, face challenges in effectively governing AI. Traditional methods fall short due to AI's non-linear, unpredictable nature. AI systems are akin to Complex Adaptive Systems (CAS), where components interact and evolve in unpredictable ways. This complexity results in butterfly effects, where minor changes can lead to significant, unforeseen consequences. Consequently, standard regulatory methods that work for static, linear systems are inadequate for AI, necessitating a different approach that acknowledges and accommodates its complex and adaptive nature.

The current landscape of policy frameworks has a broad spectrum of control, each with its own associated issues. This includes a 'hands-off' & self-regulation approach (USA) and a pro-innovation laissez-faire approach (UK), which has significant associated dangers of regulation without teeth given the far-reaching and unpredictable consequences of AI. The risk-based approach (EU) inherently assumes that a classification of risks in a non-deterministic system is possible. Similarly, the use of direct bureaucratic control (China) can fail and have cascading effects, as demonstrated by the likely lab-leak origin of Covid-19.

The aforementioned approaches do not fully address the CAS characteristics of AI, underlining the need for regulatory frameworks that set boundary conditions, enable real-time monitoring, and guide the evolution of AI systems. Our analysis suggests a complex adaptive system (CAS) approach to AI regulation.

To effectively regulate AI (algorithm, training data sets, models, and applications), a novel framework based on CAS thinking is proposed, consisting of five key principles:

- **Establishing Guardrails and Partitions:** Implement clear boundary conditions to limit undesirable AI behaviours. This includes creating "partition walls" between distinct systems and within deep learning AI models to prevent systemic failures, similar to firebreaks in forests.
- **Mandating Manual 'Overrides' and 'Authorization Chokepoints':** Critical infrastructure should include human control mechanisms at key stages to intervene when necessary, emphasizing the need for specialized skills and dedicated attention without limiting automation of systems. Manual overrides empower humans to intervene when AI systems behave erratically or create pathways to cross-pollinate partitions. Meanwhile, multi-factor authentication authorization protocols provide robust checks before executing high-risk actions, requiring consensus from multiple credentialed humans.
- **Ensuring Transparency and Explainability:** Open licensing of core algorithms for external audits, AI factsheets, and continuous monitoring of AI systems is crucial for accountability. There should be periodic mandatory audits for transparency and explainability.

- **Defining Clear Lines of AI Accountability:** Mandate standardized incident reporting protocols to document any system aberrations or failures. Establish predefined liability protocols to ensure that entities or individuals are held accountable for AI-related malfunctions or unintended outcomes. This proactive stance inserts an ex-ante "Skin in the Game," ensuring that system developers and operators remain deeply invested and accountable for AI outcomes.
- **Creating a Specialist Regulator:** Traditional regulatory mechanisms often lag the rapid pace of AI evolution. A dedicated, agile, and expert regulatory body with a broad mandate and the ability to respond swiftly is pivotal to bridging this gap, ensuring that governance remains proactive and effective. This would also entail having a national registry of algorithms as compliance and a repository of national algorithms for innovations in AI.

The regulation of financial markets offers valuable insights for AI regulation. Financial markets exemplify CAS with emergent behaviours arising from diverse interacting agents. Individual traders adapt strategies to respond to price signals, news, and each other's actions. This feeds back in nonlinear ways, resulting in volatile, unpredictable market dynamics. Yet proactive governance demonstrates feasible regulation approaches. Establishing dedicated oversight through bodies like SEBI in India and SEC in the United States of America (USA) provides specialized expertise attuned to nuances.

Requirements like financial statements, auditing and mandated incident reporting enhance transparency and traceability - akin to explainability standards and auditing for AI algorithms. Circuit breakers act as control chokepoints, halting trading when indices cross threshold triggers to prevent cascading crashes. Margin requirements, settlement periods and other boundaries constrain destabilizing behaviours. Mechanisms like bankruptcy and insolvency provide ultimate overrides when entities become non-viable. Fixing liability through personal accountability and 'skin in the game' holds individuals responsible for corporate actions. While not a perfect parallel, insights from governing chaotic markets can inform AI regulation. By applying similar strategies, such as partitions, transparency standards, control points, and accountability measures, AI's development can be steered responsibly, ensuring safety and ethical deployment akin to the orderly functioning of financial markets.

This proposed approach aims to navigate the complexities of AI regulation by learning from other CAS and adapting to the unique challenges posed by AI's evolving nature. Implementing prudent measures today can help ensure that AI develops in a way that is beneficial and safe for society.

I. Introduction

Artificial Intelligence (AI) has transitioned from academic and technological experimentation to widespread application in the past few years, where advanced AI frameworks have been deployed exponentially with multilayered application modalities. Governments across the world have been continuously making efforts to support and regulate AI technology and applications¹. Key to the regulatory discourse is the ‘black box problem’ in AI models. They are end products of an algorithm (set of procedures) that has been trained on a large set of examples (data). To keep the intellectual property intact, generally, models are hidden or are put in a ‘black box’². Consumption-oriented opaqueness, where the model can give an output to the user, has increasingly become challenging and requires adequate policy attention. In addition, even developers are facing difficulty explaining these models, especially the ones that use emergent deep learning algorithms.

AI models are fundamentally driven by data. The quality, quantity, and diversity of data directly impact the performance, accuracy, and reliability of AI models. Robust and representative datasets play a pivotal role in developing effective AI solutions, particularly in the realm of machine learning and data-driven decision-making. The quality and diversity of data used to train AI models are crucial determinants of their performance, accuracy, and fairness. High-quality data is free from errors and inconsistencies. According to IBM, poor data quality costs the US economy around \$3.1 trillion annually³. AI models trained on high-quality data are more reliable and make more accurate predictions. Datasets must represent the full spectrum of scenarios and demographics the AI system will encounter in the real world. A study by MIT showed that facial recognition systems had higher error rates for darker-skinned females, primarily due to underrepresentation in training datasets⁴. In AI systems, the Butterfly Effect highlights how minor biases or skewed data inputs can lead to significant, unfair outcomes, amplifying inherent biases and creating vulnerabilities. Addressing these issues is crucial for promoting fairness and responsible AI development⁵.

The real problem is that AI can iteratively improve its capabilities beyond normal human comprehension if left unregulated. The state of AI poses significant risks, and concerted efforts are needed to make these systems more transparent and their decisions more interpretable to those who rely on them^{6,7} (one such effort is the explainable artificial intelligence (XAI)⁸ initiative). The lack of clear tangible explanations can severely hinder our ability to evaluate, regulate, debug, and improve AI systems, especially when failures occur.

¹ Regulate AI? Here’s What That Might Mean in the US. (2023). Bloomberg.com. [online] 30 Oct. Available at: <https://www.bloomberg.com/news/articles/2023-10-30/ai-regulation-what-biden-s-new-rules-might-mean-in-the-us#xj4y7vzkg>.

² Bagchi, S. (2023) Why we need to see inside AI’s Black Box, Scientific American. Available at: <https://www.scientificamerican.com/article/why-we-need-to-see-inside-ais-black-box/> (Accessed: 06 November 2023).

³ Redman, T.C. (2016). Bad Data Costs the U.S. \$3 Trillion Per Year. [online] Harvard Business Review. Available at: <https://hbr.org/2016/09/bad-data-costs-the-u-s-3-trillion-per-year>.

⁴ Hardesty, L. (2018). Study finds gender and skin-type bias in commercial artificial-intelligence systems. [online] MIT News. Available at: <https://news.mit.edu/2018/study-finds-gender-skin-type-bias-artificial-intelligence-systems-0212>.

⁵ Ferrara, E. (2023). The Butterfly Effect in Artificial Intelligence Systems: Implications for AI Bias and Fairness. Social Science Research Network. doi:<https://doi.org/10.2139/ssrn.4614234>.

⁶ Yampolskiy, R. (07 2020) ‘Unexplainability and Incomprehensibility of AI’, Journal of Artificial Intelligence and Consciousness, 07, pp. 1–15. doi: 10.1142/S2705078520500150.

⁷ sdgs.un.org. (n.d.). Beyond a black-box approach to artificial intelligence policy – a simple guide to definitions, functions and technology types, by Richard A Roehrl (DESA) | Department of Economic and Social Affairs. [online] Available at: <https://sdgs.un.org/documents/beyond-black-box-approach-artificial-intelligence-policy-simple-guide-definitions> [Accessed 6 Nov. 2023].

⁸ Johal, A. (2023) From black box to glass box: The evolution of explainable AI, Medium. Available at: <https://medium.com/geekculture/from-black-box-to-glass-box-the-evolution-of-explainable-ai-c4b932c9fe94> (Accessed: 06 November 2023).

The standard regulatory approaches centred on ex-ante impact analysis and risk assessment seem unlikely to be sustained. For effective regulation, it is essential to recognize that AI constitutes a Complex Adaptive System (CAS) that does not follow a predictable and fundamentally deterministic path. AI systems combine elements of CAS - where multiple components interact and evolve in nonlinear ways. Much like chaos theory and the butterfly effect, small initial changes in complex AI systems can cascade into disproportionate impacts potentially going out of control. The evolutionary trajectory of AI and the unintended consequences it may engender cannot be understood or fully anticipated a priori through a reductionist approach, which is why most of the reactionary and conservative regulations keep spiralling in a revision loop.

As AI rapidly approaches and potentially surpasses human-level capabilities, an array of profound dangers arises. One concern is the loss of control over recursively self-improving AI that eclipses human comprehension, also called runaway AI^{9,10,11}. In addition to this, interconnected critical infrastructure could face disruption through hacking of systems such as power grids. The integration of AI into defence, space, and biotechnology could also escalate cyber, outer space, nuclear, and biowarfare risks through engineered pathogens, autonomous weapons, etc. AI-powered mass surveillance^{12,13} risks the erosion of privacy, free discourse disruption, and manipulation of information.

However, more subtle dangers also arise from advanced AI's ability to manipulate perceptions and simulate realities¹⁴. Through a combination of surveillance, persuasive messaging and synthetic media generation, malevolent AI could increasingly control information ecosystems and even fabricate customized deceitful realities to coerce human behaviour, inducing mass schizophrenia. As realities increasingly blur between virtual and physical, developing oversight and security protocols will be critical to preserving human agency and autonomy.¹⁵ Indeed, given the unintended consequences of self-generation, the AI need not be particularly malevolent, but in fact, attempting to be benevolent (eradication of cancer from the world could mean eradication of the disease or the patients).

This paper undertakes a comparative analysis of key AI regulations and distils fundamental learnings from global practices and ethical charters. The analysis demonstrates significant gaps vis-a-vis the parameters above in existing regulatory regimes. A better understanding of AI regulation and its ability to regulate would also help us effectively regulate quantum computing before it blows out of proportion¹⁶.

We propose a tailored and forward-looking regulatory framework with a new lens of AI as a CAS. It recommends an adaptive approach centred on setting ethical boundaries, mandating human oversight for high-risk applications, partitioning AI systems to limit contagion, and governing AI as a living

⁹ Ebbhardt, J. D. (2018). The Threat of Artificial Superintelligence. *Exigence*, 2 (1). Retrieved from <https://commons.vccs.edu/exigence/vol2/iss1/1>

¹⁰ MIT Technology Review. (2016). Yes, We Are Worried About the Existential Risk of Artificial Intelligence. [online] Available at: <https://www.technologyreview.com/2016/11/02/156285/yes-we-are-worried-about-the-existential-risk-of-artificial-intelligence/>.

¹¹ Davidson, T. (2023) The Danger of Runaway AI. *Journal of Democracy*, 34 (4), 2023, p. 132-140. Project MUSE, <https://doi.org/10.1353/jod.2023.a907694>.

¹² Feldstein, S. (2019). The Global Expansion of AI Surveillance. [online] Carnegie Endowment for International Peace. Available at: <https://carnegieendowment.org/2019/09/17/global-expansion-of-ai-surveillance-pub-79847>.

¹³ Baptista, E. (2022). China uses AI software to improve its surveillance capabilities. Reuters. [online] 8 Apr. Available at: <https://www.reuters.com/world/china/china-uses-ai-software-improve-its-surveillance-capabilities-2022-04-08/>.

¹⁴ Mustak, M. et al. (2023) 'Deepfakes: Deceptions, mitigations, and opportunities', *Journal of Business Research*, 154, p. 113368. doi: 10.1016/j.jbusres.2022.113368.

¹⁵ Emsley, R. (2023) ChatGPT: these are not hallucinations – they're fabrications and falsifications. *Schizophr*, 9,(52) . <https://doi.org/10.1038/s41537-023-00379-4>

¹⁶ Kop, V.W., Mauritz (2022). Why Quantum Computing Is Even More Dangerous Than Artificial Intelligence. [online] *Foreign Policy*. Available at: <https://foreignpolicy.com/2022/08/21/quantum-computing-artificial-intelligence-ai-technology-regulation/>

technology that evolves in ways hard to fully anticipate today. An adaptive framework also strikes a balance between mitigating risks and fostering innovation, by evolving in tandem with technological advancements.

The aim is not to delineate forecasting and meticulously regulate AI's developmental arc over decades; instead, it is to institute ethical guardrails today for foreseeable risks, mechanisms to dynamically adapt safeguards as needed, and sufficient human oversight¹⁷ throughout AI life cycles.

¹⁷ Bryson, J.J., Theodorou, A. (2019). How Society Can Maintain Human-Centric Artificial Intelligence. In: Toivonen, M., Saari, E. (eds) Human-Centered Digitalization and Services. Translational Systems Sciences, vol 19. Springer, Singapore. https://doi.org/10.1007/978-981-13-7725-9_16

II. Policy Analysis

In the broadest sense of AI regulation, including data protection and privacy laws that indirectly affect AI development and deployment, one could argue that the earliest forms of AI-related regulation date back to the data protection laws of the 1970s and 1980s. For instance, the OECD's Guidelines on the Protection of Privacy and Transborder Flows of Personal Data, established in 1980, set out principles that are fundamental to the use of AI, such as data minimization and purpose specification.¹⁸

Specific to AI regulation, which is designed with the express purpose of governing AI¹⁹, is a much more recent development, largely emerging in the 21st century as AI technology began to proliferate in modern living and global development. For example, in 2019, the High-Level Expert Group on Artificial Intelligence, which was appointed by the European Commission, released Ethics Guidelines for Trustworthy AI²⁰, which outlines a framework for achieving ethical AI that is lawful, ethical, and robust. The European Union has been particularly proactive in this area, proposing the Artificial Intelligence Act in April 2021²¹, which is often cited as one of the first comprehensive legislative proposals to regulate AI specifically.

From a regulatory standpoint, it is crucial to recognize that each nation possesses distinct capabilities and faces unique challenges, leading to a diverse landscape of legal responses. Consequently, we observe a variety of legal frameworks across different countries. Even within unified jurisdictions such as the European Union, there is a noticeable variation in the emphasis and application of enforcement measures. This diversity reflects the complex and multifaceted nature of AI technology and its implications, underscoring the need for tailored regulatory approaches that account for specific national circumstances and priorities. These regulatory approaches can vary significantly, ranging from a mix of intermittent frameworks with elements of self-regulation, to more conservative regulations focused on fundamental rights, and up to frameworks where the state exercises complete control.

In this section, we survey a range of different approaches to AI regulation across the world.

¹⁸ Kirby, M. (2010). The history, achievement and future of the 1980 OECD guidelines on privacy. *International Data Privacy Law*, 1(1), pp.6–14. doi:<https://doi.org/10.1093/idpl/ipq002>.

¹⁹ It's important to recognize that our regulatory focus is not on the AI artifacts themselves, but rather on the individuals and entities involved in their creation and application. The essence of these regulations is to guide and manage the processes of developing and deploying AI technologies.

²⁰ European Commission (2021). Ethics guidelines for trustworthy AI | Shaping Europe's digital future. [online] [digital-strategy.ec.europa.eu](https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai). Available at: <https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>.

²¹ European Parliament (2023). EU AI Act: First Regulation on Artificial Intelligence | News | European Parliament. [online] www.europarl.europa.eu. Available at: <https://www.europarl.europa.eu/news/en/headlines/society/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence>.

a. United States

The 2023 Update on The National Artificial Intelligence R&D Strategic Plan²² emphasises the role of the federal government in ensuring the responsible development of AI technologies and making a variety of thoughtful investments across diverse AI R&D and deployments for the public good. It reaffirms the key eight strategies of the 2016 draft²³ and 2019 update²⁴.

The US strategy emphasizes long-term investments in responsible AI research to advance foundational capabilities while mitigating risks from generative systems. It also focuses on effective human-AI collaboration through understanding complementarity, developing competencies, and measuring team performance. Additional priorities include expanding access to high-quality training data and environments, establishing multi-dimensional evaluation frameworks for AI systems standards and benchmarks from the Administration’s Blueprint for an AI Bill of Rights²⁵ and AI Risk Management Framework (RMF)²⁶, building an AI-ready workforce through education, and promoting public-private partnerships for sustained R&D investment and practical transition of advances. The overarching aim is to support innovation in beneficial AI while strengthening oversight to minimize harm.

The US approach to regulation has evolved from a completely ‘hands-off’ one initially to the recent shift towards ‘self-regulation’ and ‘voluntary commitments’. The US has a plethora of regulations on AI (Federal, Congressional, State, and Industry Self-Regulation)²⁷, which has made their regulatory stance laissez-faire for a while, given that the Federal administration does not want to contain the development of technology. They have asserted that the US administration would make efforts for *“maintaining American leadership in AI requires a concerted effort to promote advancements in technology and innovation, while protecting American technology, economic and national security, civil liberties, privacy, and American values and enhancing international and industry collaboration with foreign partners and allies”*²⁸.

Following this, the US Administration issued a Memorandum in 2019 for the Heads of Executive Departments and Agencies issuing Guidance for Regulation of Artificial Intelligence Applications²⁹ that stated:

“As stated in Executive Order 13859, “the policy of the United States Government [is] to sustain and enhance the scientific, technological, and economic leadership position of the United States in AI.” The deployment of AI holds the promise to improve safety, fairness, welfare, transparency, and other social goals, and America’s maintenance of its status as a global leader in AI development is vital to preserving our economic and national security. The importance of developing and deploying AI requires a regulatory approach that fosters innovation, growth, and engenders trust, while protecting core American values, through both regulatory and nonregulatory actions and reducing unnecessary barriers to the development and deployment of AI.

²²<https://www.whitehouse.gov/wp-content/uploads/2023/05/National-Artificial-Intelligence-Research-and-Development-Strategic-Plan-2023-Update.pdf>

²³ https://www.nitrd.gov/pubs/national_ai_rd_strategic_plan.pdf

²⁴ <https://www.nitrd.gov/pubs/National-AI-RD-Strategy-2019.pdf>

²⁵ <https://www.whitehouse.gov/wp-content/uploads/2022/10/Blueprint-for-an-AI-Bill-of-Rights.pdf>

²⁶ <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>

²⁷ www.morganlewis.com. (n.d.). The United States’ Approach to AI Regulation: Key Considerations for Companies. [online] Available at: <https://www.morganlewis.com/pubs/2023/05/the-united-states-approach-to-ai-regulation-key-considerations-for-companies>.

²⁸ Executive Order 13859 of February 11, 2019; Federal Register. (2019). Maintaining American Leadership in Artificial Intelligence. [online] Available at: <https://www.federalregister.gov/documents/2019/02/14/2019-02544/maintaining-american-leadership-in-artificial-intelligence>.

²⁹ The White House, Vought, R. T.(2019). MEMORANDUM FOR THE HEADS OF EXECUTIVE DEPARTMENTS AND AGENCIES FROM. [online] Available at: <https://www.whitehouse.gov/wp-content/uploads/2020/11/M-21-06.pdf>

To that end, Federal agencies must avoid regulatory or non-regulatory actions that needlessly hamper AI innovation and growth. Where permitted by law, when deciding whether and how to regulate in an area that may affect AI applications, agencies should assess the effect of the potential regulation on AI innovation and growth. Agencies must avoid a precautionary approach that holds AI systems to such an impossibly high standard that society cannot enjoy their benefits.”

The aforementioned memorandum stated that in cases where existing regulations suffice, or the costs of new regulations outweigh their benefits for AI applications, agencies might opt for non-regulatory approaches. This emphasizes avoiding hampering innovation and assessing the regulatory impact on AI growth. This aligns with laissez-faire principles prioritizing technological advancement over precaution. Additionally, agencies are encouraged to prefer voluntary consensus standards developed by the private sector or independent organizations for managing AI risks, considering these standards' robustness before proposing new regulations or compliance programs.

The approach marked a shift in July 2023 when the Biden-Harris Administration secured Voluntary Commitments from Leading Artificial Intelligence Companies to Manage the Risks Posed by AI³⁰. While the “*voluntary commitments*” and self-regulation from companies³¹ gesticulate a positive intention toward public safety and ethical considerations, relying solely on these pledges presents a significant risk in itself³² and is being argued as pro-business^{33,34}. The commitments to pre-release testing, information sharing, and investment in security are steps in the right direction, but without a binding regulatory framework, these promises lack enforceability and accountability. History has shown that self-regulation in rapidly advancing technological domains can lead to gaps in safety and ethics, as corporate interests might overshadow public welfare. The very nature of AI as an evolving and complex field makes it difficult to anticipate all potential risks, and voluntary measures are inadequate to cover unforeseen vulnerabilities. The US's regulatory approach to AI, with its emphasis on voluntary industry commitments, carries significant risks and may be inadequate to address the full range of potential issues posed by the development and deployment of AI technologies.

This self-regulatory approach was further enhanced with proposed limited oversight in the latest OMB Memorandum³⁵ that was sent out for public review in October 2023. It reflects a regulatory stance focused on promoting responsible AI innovation while rigorously managing its risks through “*specific minimum risk management practices for uses of AI that impact the rights and safety of the public, especially those affecting public rights and safety*”. It mandates specific governance structures, strategies for AI implementation, and comprehensive risk management practices, emphasizing accountability, transparency, and public engagement. However, it does not indicate any regulatory guidance for the agencies as indicated in the

³⁰ The White House (2023). FACT SHEET: Biden-Harris Administration Secures Voluntary Commitments from Leading Artificial Intelligence Companies to Manage the Risks Posed by AI. [online] The White House. Available at:

<https://www.whitehouse.gov/briefing-room/statements-releases/2023/07/21/fact-sheet-biden-harris-administration-secures-voluntary-commitments-from-leading-artificial-intelligence-companies-to-manage-the-risks-posed-by-ai/>

³¹ Zakrzewski, C. and Tiku, N. (2023). AI companies form new safety body, while Congress plays catch-up. Washington Post. [online] 27 Jul. Available at: <https://www.washingtonpost.com/technology/2023/07/26/ai-regulation-created-google-openai-microsoft/>.

³² Alikhani, K. (2019). Council Post: Why It's Dangerous For AI To Regulate Itself. [online] Forbes. Available at: <https://www.forbes.com/sites/forbestechcouncil/2019/03/22/why-its-dangerous-for-ai-to-regulate-itself/?sh=3d79c20d7e54>.

³³ VentureBeat. (2023). Why self-regulation of AI is a smart business move. [online] Available at: <https://venturebeat.com/ai/why-self-regulation-of-ai-is-a-smart-business-move/>

³⁴ Adonis Hoffman, opinion contributor (2023). Why self-regulation is best for artificial intelligence. [online] The Hill. Available at: <https://thehill.com/opinion/4300288-why-self-regulation-is-best-for-artificial-intelligence/>

³⁵ The White House, Young, S. (2023). PROPOSED MEMORANDUM FOR THE HEADS OF EXECUTIVE DEPARTMENTS AND AGENCIES FROM. [online] Available at: <https://www.whitehouse.gov/wp-content/uploads/2023/11/AI-in-Government-Memo-draft-for-public-review.pdf>

2019 memorandum; it only suggests better governance. Since it's not a superseding document and is still open to public review, the previous guidelines do remain effective. Industry self-regulation has been seen to promote the industry's interest at the expense of the public³⁶, and its failures have been noted across disciplines^{37,38}.

Recently, the US Government took another step towards regulation through the Presidential Executive Order of 30 October 2023³⁹. The order presents a plan to address the risks of AI, requiring companies to rigorously test new products like ChatGPT and share findings with U.S. officials, a move seen as a significant regulatory step. In terms of regulation, it calls for a '*potential voluntary, regulatory, and international mechanisms to manage the risks and maximize the benefits*' of AI in its emergent regulatory framework while maintaining that the "*Administration will support programs to provide Americans the skills they need for the age of AI and attract the world's AI talent to our shores — not just to study, but to stay — so that the companies and technologies of the future are made in America*" making regulatory sector-specific regulatory authorities independent allowing them, "*...as they deem appropriate, to consider whether to mandate guidance through regulatory action in their areas of authority and responsibility*". The order includes provisions like AI-enabled product watermarks to guard against deepfakes and mandates biotech precautions, but these are advisory rather than mandatory.

The order also instructs U.S. agencies to adopt AI practices, setting an example for the private sector, with specific departments addressing AI's threat to critical infrastructure. Highlighting AI threats to critical infrastructure acknowledges interconnected systemic risks. Aligned with a complex systems perspective. One key challenge that it highlights is the risk of AI-induced biowarfare resounding the concerns of a 2017 study by John T. O'Brien and Cassidy Nelson⁴⁰.

The Executive Order moves towards more oversight by requiring testing and information sharing. This is somewhat of a departure from pure laissez-faire but remains advisory rather than mandatory regulation. The Blueprint for an AI Bill of Rights⁴¹ highlights the development of responsible automated systems and creates barriers such as the ones we also propose. Overall, recent gestures indicate movement away from pure laissez-faire, but it still relies largely on self-regulation.

³⁶ Edwards, Benjamin P., "The Dark Side of Self-Regulation" (2017). Scholarly Works. 1117.<https://scholars.law.unlv.edu/facpub/1117>

³⁷ UK Parliament, EDM (Early Day Motion)1345: tabled on 08 February 1993, FAILURE OF SELF-REGULATION OF FINANCIAL SERVICES, Available at: <https://edm.parliament.uk/early-day-motion/6448/failure-of-selfregulation-of-financial-services>

³⁸ Amit Narang, O.C. (2019). Corporate self-regulation is failing. [online] The Hill. Available at: <https://thehill.com/blogs/congress-blog/the-administration/436328-corporate-self-regulation-is-failing/>.

³⁹ The White House (2023). Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence. [online] The White House. Available at: <https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/>.

⁴⁰ O'Brien JT, Nelson C. Assessing the Risks Posed by the Convergence of Artificial Intelligence and Biotechnology. *Health Secur.* 2020 May/Jun;18(3):219-227. doi: 10.1089/hs.2019.0122. PMID: 32559154; PMCID: PMC7310294.

⁴¹ The White House (2022). Blueprint for an AI Bill of Rights. [online] The White House. Available at: <https://www.whitehouse.gov/ostp/ai-bill-of-rights/>.

b. United Kingdom

The United Kingdom adopted a cautious yet pro-innovation stance towards AI governance, aimed at supporting growth while managing risks. The 2017 Industrial Strategy White Paper⁴², now retracted, initially stimulated business environments, skills, and infrastructure. It was transitioned to the 2021 Plan for Growth, which does not explicitly mention AI⁴³. The 2018 AI Sector Deal committed £1 billion to advance AI across sectors, emphasizing research, innovation, infrastructure, and skills development⁴⁴. This aligned with the 2022 National AI Strategy's⁴⁵ core focus on sustaining UK AI leadership. The UK administration advocates an application-centric framework, with regulations tailored to specific use contexts rather than one-size-fits-all governance. Overall, the UK strategy rejects precautionary restraints on AI advancement focusing on innovation, instead preferring selective governance guided by ethical principles and aligned with national interests.⁴⁶

This is largely antithetical to our principle of 'Guardrails' advocated later in the paper. It centres on targeted support for technology growth, while instituting sector-specific regulations to address foreseeable sectoral risks. This approach enables sector-specific regulators to customize AI governance based on contextual needs, capitalizing on their specialized knowledge and experience. This is unlike the specialized common regulator that we advocate, given the pervasive nature and the risks of super-connected AI systems.

Principally, the UK articulates a quintet of foundational principles designed to scaffold the suggested regulatory apparatus⁴⁷, principles which are applicable across a broad spectrum of sectors such as safety, appropriate transparency and explainability, equity, and mandatory accountability and governance. It also advances the proposition that, where deemed pertinent, aggrieved parties and entities engaged in the AI lifecycle should have the agency to challenge AI-generated decisions or outcomes that are detrimental or pose a significant risk. While these principles, especially transparency and explainability, find resonance in the approach we propose, the principles-based approach to AI regulation shared by the UK is undermined by the ambiguity of the principles and the challenge of converting them into actionable standards, further complicated by potential conflicts with established legal frameworks. The minimal legislative efforts do not equip regulators to enforce these principles effectively.

A significant deficiency in the proposal is the lack of a detailed regulatory model, particularly in defining the operational structure and mechanisms of the central regulatory function, leaving questions about standardization, resources, accountability, and implementation. The ultimate configuration of this proposed regulatory framework, as outlined in the White Paper, may yet undergo evolution. Several updates have already emerged.

⁴² Department for Business, Energy & Industrial Strategy (2017). Industrial Strategy: building a Britain fit for the future. [online] GOV.UK. Available at: <https://www.gov.uk/government/publications/industrial-strategy-building-a-britain-fit-for-the-future>.

⁴³ GOV.UK. (2021). Build Back Better: our plan for growth (HTML). [online] Available at: <https://www.gov.uk/government/publications/build-back-better-our-plan-for-growth/build-back-better-our-plan-for-growth-html>.

⁴⁴ GOV.UK. (2022). National AI Strategy - HTML version. [online] Available at: <https://www.gov.uk/government/publications/national-ai-strategy/national-ai-strategy-html-version>.

⁴⁵ GOV.UK. (2022). National AI Strategy - HTML version. [online] Available at: <https://www.gov.uk/government/publications/national-ai-strategy/national-ai-strategy-html-version>.

⁴⁶ GOV.UK (2023). A pro-innovation approach to AI regulation. [online] GOV.UK. Available at: <https://www.gov.uk/government/publications/ai-regulation-a-pro-innovation-approach/white-paper>.

⁴⁷ Ada Lovelace Institute. Regulating AI in the UK. [online] Available at: <https://www.adalovelaceinstitute.org/report/regulating-ai-in-the-uk/>.

The UK's pro-innovation agenda in AI regulation is concerning due to its potential inadequacy in meeting the complex challenges posed by AI⁴⁸. The focus on positioning the UK as a prime location for AI companies seems to favour industrial growth at the expense of public welfare in the face of AI's risks.

They state that:

“While we should capitalise on the benefits of these technologies, we should also not overlook the new risks that may arise from their use, nor the unease that the complexity of AI technologies can produce in the wider public. We already know that some uses of AI could damage our physical and mental health, infringe on the privacy of individuals and undermine human rights.

Public trust in AI will be undermined unless these risks, and wider concerns about the potential for bias and discrimination, are addressed. By building trust, we can accelerate the adoption of AI across the UK to maximise the economic and social benefits that the technology can deliver, while attracting investment and stimulating the creation of high-skilled AI jobs. In order to maintain the UK’s position as a global AI leader, we need to ensure that the public continues to see how the benefits of AI can outweigh the risks.”⁴⁹

It goes on further to state that:

“New rigid and onerous legislative requirements on businesses could hold back AI innovation and reduce our ability to respond quickly and in a proportionate way to future technological advances. Instead, the principles will be issued on a non-statutory basis and implemented by existing regulators. This approach makes use of regulators’ domain-specific expertise to tailor the implementation of the principles to the specific context in which AI is used. During the initial period of implementation, we will continue to collaborate with regulators to identify any barriers to the proportionate application of the principles and evaluate whether the non-statutory framework is having the desired effect.”⁵⁰

The purportedly agile regulatory framework may be a pretext for possible deregulation, advocating for the removal of "barriers" without justifying the abandonment of existing safeguards amidst heightened AI risks. This could compromise the autonomy and efficacy of regulatory bodies essential for overseeing AI's multifaceted impact.

Similar to the US's 'hands-off' approach, the UK's pro-innovation approach, where principles governing AI regulation in the UK are issued on a non-statutory basis, allows existing domain-specific regulators to tailor AI governance to specific contexts, leveraging their domain-specific expertise. This method underscores a laissez-faire attitude by preferring guidance and principles over hard legislation.

⁴⁸ Charlesworth, A., Fotheringham, K., Gavaghan, C., Albert Sánchez Graells and Torrible, C. (2023). Response to the UK's March 2023 White Paper 'A pro-innovation approach to AI regulation'. Social Science Research Network. doi: <https://doi.org/10.2139/ssrn.4477368>.

⁴⁹ GOV.UK (2023). A pro-innovation approach to AI regulation. [online] GOV.UK. Available at: <https://www.gov.uk/government/publications/ai-regulation-a-pro-innovation-approach/white-paper>.

⁵⁰ *ibid*

c. The European Union

In contrast to the UK and the US, the European Union (EU) has taken a risk-based approach in recognizing the need to balance AI innovation with safeguarding fundamental rights and addressing potential risks⁵¹. The EU's approach to AI regulation is centred around the European Commission's proposal for the Regulation of Artificial Intelligence⁵². This legislative proposal is a component of a more comprehensive Artificial Intelligence strategy, inclusive of the refreshed Coordinated Plan on AI⁵³.

The cornerstone of the EU's AI regulations is the EU's AI Act⁵⁴, which aims to create a comprehensive framework for AI governance. The Act emphasizes a risk-based approach, categorizing AI systems into different levels of risk and imposing specific obligations accordingly. It prohibits AI systems that are considered "unacceptable" under EU law. In this approach, it classifies AI systems into four categories: Unacceptable Risk, High Risk, Limited Risk, and Minimal Risk.

Unacceptable Risk AI systems, such as those used for social scoring or manipulating human behaviour, are outright banned. High-risk AI systems, such as those employed in critical infrastructure, healthcare, or law enforcement, are subject to stringent requirements, including conformity assessments, documentation, and human oversight. The Act also imposes specific obligations on high-risk AI systems to ensure transparency, accountability, and robust governance. High-risk AI systems also must undergo conformity assessment procedures before being placed on the market or deployed.^{55,56}

EU after successfully deploying GDPR and GDR passed the revised draft AI Act in the European Parliament earlier in 2023 and then was sent to the European Council for negotiation with a goal to get it passed before the end of 2023.⁵⁷ This revised draft took a tough stance on AI risks and prohibited AI's use for biometric surveillance.

A risk-based system is completely dependent on the designation of the risk. However, in a complex system such as AI, it cannot be deterministically designated. What might be perceived as low, or no risk can lead to cascading high-risk outcomes given the unpredictability of the algorithm and perpetuation of processing logic through deep learning processes, even on a heterogeneous data set.

⁵¹ Feingold, S. (2023). 'On artificial intelligence, trust is a must, not a nice to have,' one lawmaker said. #AI. [online] World Economic Forum. Available at: <https://www.weforum.org/agenda/2023/06/european-union-ai-act-explained/>.

⁵² European Commission. (2023). Regulatory framework proposal on artificial intelligence | Shaping Europe's digital future. [online] Available at: <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai#:~:text=The%20Commission%20is%20proposing%20the>.

⁵³ European Commission. (2021). Coordinated Plan on Artificial Intelligence | Shaping Europe's digital future. [online] Available at: <https://digital-strategy.ec.europa.eu/en/policies/plan-ai>.

⁵⁴ European Commission (2021). EUR-Lex - 52021PC0206 - EN - EUR-Lex. [online] Europa.eu. Available at: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52021PC0206>.

⁵⁵ Title II & III of the EU AI Act

⁵⁶ Chee, F.Y., Mukherjee, S., Chee, F.Y. and Mukherjee, S. (2023). EU lawmakers vote for tougher AI rules as draft moves to final stage. Reuters. [online] 14 Jun. Available at: <https://www.reuters.com/technology/eu-lawmakers-agree-changes-draft-artificial-intelligence-rules-2023-06-14/>.

⁵⁷ European Parliament (2023). EU AI Act: first regulation on artificial intelligence | News | European Parliament. [online]. Available at: <https://www.europarl.europa.eu/news/en/headlines/society/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence#:~:text=On%2014%20June%202023%2C%20MEPs> [Accessed 6 Nov. 2023].

This problem is worsened by the fact that the European regulation with its regulatory localization in ‘European Values’ is ambiguous about what this actually means. The interpretation of the term and then eventually the law is based on entirely arbitrary and undefinable criteria.

There is a difference of opinion between the European Parliament & the European Council, and the European Commission. The "trilogue" discussions were crucial for finalizing a cohesive AI regulatory framework, with pressing issues like AI definitions, high-risk categories, and biometric identification on the table. The proposals differ in their approach to enforcement— the European Parliament had favoured a single national surveillance authority per member state for stronger, more centralized oversight, diverging from the multiple authority model proposed by the European Council and Commission. The EU AI Act was politically agreed upon in December 2023.

Fundamental challenges that are of concern in EU’s AI regulation are⁵⁸:

1. Regulation of AI Systems by Risk Category

- a) Risk-Based Approach: The AI Act categorizes AI technologies based on their risk to safety, health, or fundamental rights, sorting them into unacceptable, high, and limited or minimal risk levels.
- b) Prohibited AI Systems: AI systems considered unacceptable are banned outright under Article 5. These include systems that manipulate behaviour, exploit vulnerabilities, or perform social scoring that results in harm or discrimination.
- c) High-Risk AI Systems: High-risk AI systems are allowed but come with stringent obligations, such as informing human operators of risks like automation bias and conducting fundamental rights impact assessments.
- d) Limited or Minimal-Risk AI Systems: AI systems with lower risks have minimal obligations, primarily around transparency, to enable informed user decisions.
- e) Contention Over Risk Categories: Debate continues over which systems belong to which risk category, especially regarding biometric surveillance. The European Parliament has pushed for a broader list of prohibitions, including real-time remote biometric identification, despite attempts for derogation. Conversely, the European Council advocates for a narrower scope of prohibited systems. The trilogue discussions are expected to address these disputes and potentially alter the list of prohibitions and exceptions within Article 5.

2. High-Risk AI System Obligations and Categories

- a) Article 6 and Annex III Requirements:
 - Eight primary high-risk AI system categories are defined.
 - Categories include critical infrastructure, biometric identification, education, employment, essential services, law enforcement, border control, and administration of justice/democracy.
 - Additional sub-categories and expansions by the European Parliament include social media platforms and election influence.
- b) High-Risk Classification Ambiguity:
 - Challenges in determining high-risk status for AI systems.
 - European Commission's list revised by the Council and Parliament.
 - Risk classification studies show unclear risk levels for many AI systems.

⁵⁸ Müge Fazlioglu, (2023). Contentious areas in the EU AI Act trilogues. [online] Available at: <https://iapp.org/news/a/contentious-areas-in-the-eu-ai-act-trilogues/>.

While we agree with the setup of a separate regulator proposed in the legislation, the risk-based approach to AI regulation, which categorizes AI applications based on their potential threat to safety, health, or fundamental rights, presents several challenges. There is no consensus on what constitutes 'high risk' in AI.

While it may be obvious in a few instances (e.g., nuclear weapons), the problem is with ex-ante risk categorisation. AI technologies evolve rapidly, and a system deemed low-risk today may become high-risk tomorrow due to changes in usage, context, or capabilities. The challenge is that this approach only works for static, linear systems with predictable risks. AI combines qualities of complex adaptive systems (CAS) where components interact and evolve in nonlinear ways. This can lead to butterfly effects, where small changes can cascade disproportionately through AI systems. Similarly, its evolutionary trajectory cannot be predicted through reductionist thinking.

d. China

China was one of the earliest actors to have an AI strategic plan, but their approach is entirely based on bureaucratic oversight. They came up with a New Generation Artificial Intelligence Development Plan (2015-2030)⁵⁹. The document delineates a strategic blueprint for China's advancement in AI. Initially, by the terminus of 2020, the ambition was to narrow the innovation chasm separating China from global frontrunners in technology, elevating its AI sector to a status of international prominence. Subsequently, the intermediate milestone of 2025 targets seminal triumphs in the core research domains of AI and aims to establish the foundational structure for Next Generation Artificial Intelligence (NGAI). The culminating objective for 2030 envisions China's emergence as a global vanguard of AI innovation, complete with a comprehensive suite of regulatory, legal, and ethical frameworks to guide the responsible evolution of AI. Since China started very early in understanding, harnessing, and regulating AI, it is in a better position to regulate more complex applications than any other country⁶⁰.

In addressing the evolving landscape of algorithmic governance, China has instituted a trio of regulatory frameworks targeting artificial intelligence (AI) technologies. These include the 2021 mandate on recommendation algorithms⁶¹, the 2022 provisions governing deep synthetic media⁶², and the prospective 2023 guidelines concerning generative AI systems⁶³.

The core objective shared by these regulatory actions is to fortify information oversight, yet their scope extends beyond this focal point, incorporating various significant clauses. Specifically, the 2021 regulation on recommendation algorithms introduces measures to curb pronounced price differentiation practices and safeguards the welfare of personnel influenced by algorithmically managed work schedules. The 2022 rule targeting deep synthetic content mandates that such material be clearly marked as artificially generated, thereby ensuring transparency.

The proposed 2023 draft rules on generative AI pose a particularly stringent criterion, mandating the veracity and precision of both the input data and the outputs produced by AI, presenting a challenging standard for AI chatbot services to meet. Moreover, each set of regulations compels developers to formally register their algorithms with a novel state-operated algorithm registry. This repository is designed to catalogue detailed information regarding the training methodologies of algorithms. Additionally, developers are obliged to conduct and submit a security self-assessment, attesting to the integrity and safety of their algorithmic systems. However, note that this entire set of comprehensive rules is based on closed-door bureaucratic control rather than transparency.

⁵⁹ Webster, G., Creemers, R., Kania, E. and Triolo, P. (2017). Full translation: China's 'new generation artificial intelligence development plan' (2017). [online] DigiChina. Available at: <https://digichina.stanford.edu/work/full-translation-chinas-new-generation-artificial-intelligence-development-plan-2017/>.

⁶⁰ Sheehan, M. (n.d.). China's AI Regulations and How They Get Made. [online] Carnegie Endowment for International Peace. Available at: <https://carnegieendowment.org/2023/07/10/china-s-ai-regulations-and-how-they-get-made-pub-90117>.

⁶¹ Cyberspace Administration of China (2022). 互联网信息服务算法推荐管理规定-中共中央网络安全和信息化委员会办公室. [online] Available at: http://www.cac.gov.cn/2022-01/04/c_1642894606364259.htm.

⁶² Cyberspace Administration of China (2022). 国家互联网信息办公室等三部门发布《互联网信息服务深度合成管理规定》-中共中央网络安全和信息化委员会办公室. Available at: http://www.cac.gov.cn/2022-12/11/c_1672221949318230.htm

⁶³ Cyberspace Administration of China (2023). 生成式人工智能服务管理暂行办法_中央网络安全和信息化委员会办公室. [online] Available at: http://www.cac.gov.cn/2023-07/13/c_1690898327029107.htm.

The Wuhan experience has shown the potential pitfalls of an all-powerful centralised bureaucratic system with the likely lab-leak origin of Covid-19^{64,65,66}. Therefore, a system of opaque state control over algorithms, training sets, and models with a goal to maximize interests can lead to catastrophic outcomes. Ironically, China's proactive approach puts it in a position to create global standards much before the international community is able to build consensus on fundamental issues regarding the technology and its applications⁶⁷. The absence of feedback loops in such regulatory designs means that they are prone to catastrophic breakdowns.⁶⁸

China is aggressively advancing in the Internet of Things (IoT), aiming to become a global leader and posing significant challenges to non-Chinese economic and security interests. With its expansive growth, IoT is set to revolutionize numerous sectors with billions of connected devices. Yet, issues of operation, safety, and data security remain unresolved. China's strategy encompasses leveraging national resources and its large market for IoT promotion, resulting in a commanding presence in global markets. China's IoT strategy is self-serving, prioritizing Chinese interests and showing little regard for the economic and national security concerns of non-Chinese entities⁶⁹. While we agree with human hierarchical control, a super-connected AI system without any 'Guardrails' and 'Partitions' makes it almost impossible to govern and has a risk of a small system failure infecting the entire ecosystem.

⁶⁴ Knight, D. (2021). COVID-19 Pandemic Origins: Bioweapons and the History of Laboratory Leaks. *Southern Medical Journal*, [online] 114(8), pp.465–467. doi: <https://doi.org/10.14423/smj.0000000000001283>.

⁶⁵ Ciuriak, D. (2021). Why Wuhan? What the Circumstantial Evidence Says About the Origins of COVID-19. [online] *papers.ssrn.com*. Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3763542.

⁶⁶ Cohen, J. (2022). Do three new studies add up to proof of COVID-19's origin in a Wuhan animal market? [online] *www.science.org*. Available at: <https://www.science.org/content/article/do-three-new-studies-add-proof-covid-19-s-origin-wuhan-animal-market>.

⁶⁷ China Briefing News. (2022). China to Regulate Deep Synthesis (Deepfake) Technology from 2023. [online] Available at: <https://www.china-briefing.com/news/china-to-regulate-deep-synthesis-deep-fake-technology-starting-january-2023/>.

⁶⁸ S. Zheng, J. Zhang, *Time*. (2023). China Tries to Balance State Control and State Support of AI. [online] Available at: <https://time.com/6304831/china-ai-regulations/>.

⁶⁹ Henk H.F. Smid. (2023). The Space Review: Internet of Things: the China perspective. [online] Available at: <https://www.thespacereview.com/article/4566/1>.

e. The United Nations

It should be highlighted that currently, there is no comprehensive international regulatory framework governing AI that is legally binding, neither under the auspices of the UN nor within any regional intergovernmental organizations.

The High-level Panel on Digital Cooperation outlined the Roadmap for Digital Cooperation⁷⁰ and resulted in the establishment of the Office of the Secretary-General's Envoy on Technology (OSET), which spearheads the initiatives of the UN Secretary-General and the UN system in advancing new technologies, encompassing AI. In addition to the UN Secretary-General's Roadmap for Digital Cooperation, the Common Agenda, and the Global Digital Compact—which features a principal action on fostering worldwide cooperation on AI and are all driven by the UN Secretariat—various UN entities have introduced their distinct reports and initiatives. These are aimed at tackling AI and related emerging technologies, along with the specific challenges they pose to their individual mandates and stakeholders.^{71,72,73,74,75}

UNESCO's AI Ethics global standards recommend a series of principles that align with the advancement and safeguarding of human rights. It promotes fundamental tenets, including transparency, accountability, and adherence to the rule of law in the digital realm, which have also been adopted.⁷⁶ These tenets resonate with the CAS principle of making AI systems transparent and their operations understandable, which is essential for accountability and ethical considerations.

In UNSC remarks, the Secretary General also pointed out that, *“the interaction between AI and nuclear weapons, biotechnology, neurotechnology, and robotics is deeply alarming. Generative AI has enormous potential for good and evil at scale. Its creators themselves have warned that much bigger, potentially catastrophic and existential risks lie ahead.”*⁷⁷ The United Nations' approach to AI regulation, while aspirational and aligned with global cooperation, falls short of the specificity and enforceability required by the principles proposed for regulating AI as a Complex Adaptive System.

Despite the UN providing a long list of issues pertaining to several aspects of AI, it does not provide an intellectual framework that allows for real-life policy trade-offs. Therefore, the extensive array of principles and initiatives they propose fails to translate into a viable framework for the establishment of enforceable regulations.

⁷⁰ The United Nations (2020). Secretary-General's Roadmap for Digital Cooperation. [online] Available at: <https://www.un.org/en/content/digital-cooperation-roadmap/>.

⁷¹ ITU (2022). ITU Partner2Connect Digital Coalition. [online] Available at: <https://www.itu.int/itu-d/sites/partner2connect/>.

⁷² UNICEF (n.d.). Giga. [online] Available at: <https://www.unicef.org/innovation/giga>.

⁷³ sdgailab.org. (n.d.). SDG AI Lab. [online] Available at: <https://sdgailab.org/> [Accessed 8 Nov. 2023].

⁷⁴ WHO Guidance (2021). Ethics and governance of artificial intelligence for health. [online] www.who.int. Available at: <https://www.who.int/publications/i/item/9789240029200>.

⁷⁵ AIEA (2023). 2024 5th International Conference on Artificial Intelligence and Electromechanical Automation (AIEA 2024). [online] Available at: <http://www.icaiea.com/> [Accessed 8 Nov. 2023].

⁷⁶ UNESCO (2022). Ethics of artificial intelligence | UNESCO. [online] www.unesco.org. Available at: <https://www.unesco.org/en/artificial-intelligence/recommendation-ethics>.

⁷⁷ The United Nations (2023). Secretary-General's remarks to the Security Council on Artificial Intelligence | United Nations Secretary-General. [online] Available at: <https://www.un.org/sg/en/content/sg/speeches/2023-07-18/secretary-generals-remarks-the-security-council-artificial-intelligence#:~:text=And%20the%20interaction%20between%20AI>.

III. AI: A Complex Adaptive System

Complex adaptive systems (CAS) describe a broad class of intricate, dynamic networks found throughout nature and society. Diverse agents interact and adapt in CAS, exhibiting emergent collective behaviours that are nonlinear and cannot be predicted by analysing individual components. Examples range from ant colonies to brains to cities. These decentralized, self-organizing systems evolve through feedback loops, phase transitions and sensitivity to initial conditions. Even small perturbations can cascade in chaotic, unpredictable ways. Understanding CAS requires embracing their inherent complexity rather than oversimplifying analyses. Conceptualizing artificial intelligence as sharing key qualities with CAS provides a nuanced perspective on governance to match its profound technological dynamism.⁷⁸

From the lens of AI as a complex system, the emergent behaviour from simple rules has been an intriguing aspect of study across various disciplines. AI epitomizes this phenomenon. Like the brain's neurons that modify connections based on environmental stimuli, or the financial market reflecting the collective will of its participants, AI systems learn and adapt, showcasing the power of agent-based models. In this realm, AI is akin to a self-organizing entity, where individual algorithms—acting as agents—follow basic protocols yet collectively manifest sophisticated behaviour. This adaptation is crucial for AI systems, enabling them to refine their predictive capabilities and responses to dynamic environments. Newtonian models often fall short of explaining these adaptive processes. The complexity in AI arises as these agents adjust not only their actions but also the fundamental rules guiding these actions in response to external changes. This self-modifying characteristic of AI agents is what underpins the field's complexity and challenges traditional analytical approaches, underscoring the need for new scientific paradigms to understand and harness the potential of AI as a complex adaptive system.⁷⁹

Like chaos theory's butterfly effect, small initial changes in AI systems can cascade and compound rapidly into disproportionate impacts if left unchecked. The intricacy of combinatorial explosions from multiple technology facets interacting makes modelling the long-term trajectory intractable. AI development resembles an evolutionary process, with continual interplay between exploration, selection, and propagation of technical variants.

This nonlinear complexity stems from the autonomy, interdependencies, and nesting of AI subsystems. For instance, a facial recognition algorithm may iterate through billions of parameter tweaks chosen by a separate neural architecture search algorithm. If the search algorithm evolves in unexpected ways, it dramatically alters the final facial recognition system. These cascading effects challenge traditional reductionist thinking centred on prediction and control. AI's dynamic complexity requires a new regulatory perspective focused on instituting real-time guardrails, rather than ex-ante predictions about outcomes. Setting initial boundary conditions provides a scaffolding to guide development while retaining flexibility to course correct.

The combinatorial complexity makes modelling long-term capabilities speculative. This uncertainty highlights why adaptive governance is imperative for rapidly advancing technologies like AI. Rather than prescribe rigid ex-ante rules, regulatory approaches must themselves evolve as capabilities grow. Prudent oversight mechanisms will be critical to steering inevitable AI progress in beneficial directions for humanity. In the next section, we propose a framework for AI regulation that takes into account the fundamentally unpredictable and non-deterministic nature of this technology.

⁷⁸Chan, S. (2001). Complex Adaptive Systems. [online] Available at: <https://web.mit.edu/esd.83/www/notebook/Complex%20Adaptive%20Systems.pdf>.

⁷⁹Booker, L.B. and Forrest, S. (2005). Introduction: Adaptation, Evolution, and Intelligence. Oxford University Press eBooks. doi: <https://doi.org/10.1093/oso/9780195162929.003.0004>.

IV. Contours of AI Regulation: A Complex Adaptive System Approach

As we traverse the dynamic landscape of Artificial Intelligence (AI), it becomes increasingly evident that traditional regulatory frameworks might falter under the weight of this burgeoning technology⁸⁰. AI, when viewed through the lens of Complex Adaptive System (CAS) thinking, offers new insights into how it can be effectively governed. CAS, characterized by its emergent behaviours and non-linear interactions, provides a suitable analogy for the intricate, evolving nature of AI. Building on this perspective, we propose a regulatory approach anchored in five core principles:

a. Instituting Guardrails and Partitions to ‘Prevent Wildfire’

Drawing inspiration from CAS theory, it's imperative to establish robust 'Guardrails' and 'Partitions' around AI systems. These Guardrails, in the form of boundary conditions, delineate the operational space, ensuring AI technologies neither exceed their intended functions nor encroach into potentially hazardous territories, such as nuclear armament decision-making. But setting boundaries is only one part of the solution. Secondly, to avoid the domino effect where a malfunction in one system cascades into a larger systemic failure, introducing 'partitions' becomes essential. Much like firebreaks in forests, these partitions act as barriers, preventing the spread of dysfunction. Importantly, it's crucial to underline that these partitions should be universally applied, regardless of the perceived risk of a system. By stringently segregating AI and automated/Internet of Things (IoT) systems, we insulate them from one another, thereby dramatically minimizing the contagion risk. While ‘The Internet of Things’ may exist, the partitions mean that the ‘Internet of Everything’ within statutory and geographical limits should be discouraged as a state policy.

Modern AI systems, akin to Complex Adaptive Systems (CAS), can exhibit unpredictable emergent behaviours. These require robust boundary mechanisms to regulate their operation. Guardrails provide a foundational layer, setting the rules of engagement, while partitions act as insulators, preventing systemic risks from spiralling out of control. The intention here is to prohibit a super-connected AI System or one that is either source-to-delivery or globally integrated.

The concept of a regulatory sandbox can be instrumental in supporting ongoing innovation while managing potential risks. This approach, successfully implemented in the fintech sector, allows for testing and development of AI technologies under regulatory oversight, providing a safe but real-world environment for innovation.

▪ **Boundary Setting:**

- Establish clear, technically defined thresholds within which an AI system operates.
- Regularly reassess and adjust these thresholds based on system performance and external environmental changes.

• **Partitioning Strategy:**

- Implement strict separation protocols for distinct AI systems, regardless of their perceived risk levels.
- Use containerization or virtualization techniques to encapsulate AI processes, ensuring controlled and isolated execution environments.

⁸⁰ Candelon, F., Carlo, R.C. di, Bondt, M.D. and Evgeniou, T. (2021). AI Regulation Is Coming. [online] Harvard Business Review. Available at: <https://hbr.org/2021/09/ai-regulation-is-coming>.

b. Ensuring Human Control through Manual ‘Overrides’ and ‘Authorization Chokepoints’

The human agency should never be entirely overshadowed by automation, especially in domains of critical infrastructure.⁸¹ Manual overrides empower humans as agents of intervention when AI systems behave erratically⁸². Meanwhile, multi-factor authentication authorization protocols provide robust checks before executing high-risk actions, requiring consensus from multiple credentialed humans.

Mandating manual overrides ensures that humans can intervene when AI systems begin to behave unpredictably⁸³. Implementing multi-factor authentication authorization protocols, especially for high-risk actions, provides a robust check and balance system. Multiple human validators, with the requisite credentials, must give their consensus before an action is executed. This layered approach, supported by hierarchical governance, ensures that there are multiple stages and opportunities for human intervention. However, it is pivotal to note that this strategy requires equipping humans with specialized skills to interface with and oversee AI effectively. Retaining human oversight in the AI-driven world ensures that machines serve humanity's best interests. Manual overrides and chokepoints ensure that when AI systems approach ambiguity or potential harm, human judgment intervenes.

- **Manual Overrides:**

- Identify key technical junctures within the AI processing chain where human intervention can be most effective.
- Some extreme outcome prohibitory pathways can be embedded in the algorithm in addition to the human intervention.
- Design and integrate manual override mechanisms at these junctures, enabling halt or redirection of AI processes.

- **Multi-Factor Authentication:**

- Design authentication systems that require multiple human validators for high-stake decisions.
- Ensure redundancy in validation to eliminate single points of failure.

- **Hierarchical Governance:**

- Establish a multi-tiered review process, with human decision-makers at each level equipped to assess AI decisions.
- Provide specialized training to these decision-makers, ensuring they understand AI system behaviours and the nuances of their respective oversight roles.

c. Transparency and Explainability

AI decisions can often make it a "black box". To mitigate this, emphasizing transparency and explainability is critical⁸⁴. Adopting open licensing for core algorithms fosters an environment where external experts can conduct audits, enhancing system accountability. Regular audits and assessments of AI systems for bias, privacy, and security risks are vital. In addition to this, establishing clear policies

⁸¹ McKendrick, J. and Thurai, A. (2022). AI Isn't Ready to Make Unsupervised Decisions. [online] Harvard Business Review. Available at: <https://hbr.org/2022/09/ai-isnt-ready-to-make-unsupervised-decisions>.

⁸² Firlj, M. and Taehagh, A. (2021), Regulating human control over autonomous systems. *Regulation & Governance*, 15: 1071-1091. <https://doi.org/10.1111/rego.12344>

⁸⁴ Balasubramaniam, N. et al. (2023) 'Transparency and explainability of AI systems: From ethical guidelines to requirements', *Information and Software Technology*, 159, p. 107197. doi: 10.1016/j.infsof.2023.107197.

⁸⁴ Balasubramaniam, N. et al. (2023) 'Transparency and explainability of AI systems: From ethical guidelines to requirements', *Information and Software Technology*, 159, p. 107197. doi: 10.1016/j.infsof.2023.107197.

and procedures for data quality control, auditing, and compliance with regulatory requirements helps AI data governance. Transparency in AI development is essential for public scrutiny and accountability. The foundation of AI system audits should be rooted in cybersecurity evaluations. This approach is essential because a comprehensive understanding of the cybersecurity landscape is fundamental to accurately assessing any AI system. Such an audit strategy reinforces the notion that AI systems are not abstract, untamed forces, but rather carefully designed and managed digital constructs, developed and maintained by professionals⁸⁵. This perspective is crucial in ensuring that AI systems are not only technologically sound but also secure and reliable in their applications.⁸⁶

National open algorithms and data training sets would also go far in the innovation initiative. Complementing this, "AI factsheets" can offer insights into the AI's engineering/development: the training data, metrics used, potential uncertainties, and other critical parameters. This offers users a clear snapshot of the system's capabilities and limitations. Moreover, continuously monitoring and debugging black-box systems provides a real-time pulse on their operations.

Transparency is the bedrock of trust in AI systems. By demystifying AI's operational black box, we not only foster trust but also ensure accountability at all stages.

- **Open Licensing of Algorithms:**

- Promote the use of open-source licenses for core AI algorithms to enable scrutiny by external experts.
- Encourage industry-wide adoption of open standards to facilitate audits and cross-system comparisons.

- **AI Factsheets:**

- Mandate the creation of standardized documentation detailing AI system development – whether through coding or learning: coding/revision control logs, data sources, training procedures, performance metrics, and known limitations.
- Implement a uniform format for these factsheets to ensure consistency and ease of interpretation.

- **Mandated Dynamic Monitoring:**

- Mandatory disclosures for extreme outcomes.
- There should be periodic mandatory audits of explainability and process audits for development and operations.
- Utilize AI debugging and monitoring tools to trace system decisions in real-time.
- Establish protocols for regular audits, ensuring continuous oversight and system accountability.

d. Distinct Accountability

While AI's influence permeates various sectors, it's imperative to have clear lines of responsibility⁸⁷. In a world where technology often outpaces legal frameworks, having predefined liability protocols can bridge this gap. Whether it's a malfunction, unintended consequence, or any aberration, there should

⁸⁵ Aucott, J. (2023). How to Perform an AI Audit for UK Organisations. [online] Available at: <https://www.haptic-networks.com/cyber-security/how-to-perform-an-ai-audit/>.

⁸⁶ Centre for European Policy Studies (CEPS). Technology, Governance and Policy Challenges. (2021). Available at: <https://eda.europa.eu/docs/default-source/documents/ceps-tfr-artificial-intelligence-and-cybersecurity.pdf>.

⁸⁷ Novelli, C., Taddeo, M. and Floridi, L. (2023). Accountability in artificial intelligence: what it is and how it works. AI & SOCIETY. doi: <https://doi.org/10.1007/s00146-023-01635-y>.

always be a clear trail leading to an entity or individual being held responsible. This ensures that accountability isn't an afterthought but rather an integral part of AI deployment. This proactive stance inserts an ex-ante "*Skin in the Game*," ensuring that system developers and operators remain deeply invested in and accountable for AI outcomes.

Further, a liability and traceability chain is vital for analyzing and ensuring the safety of individual apps and the broader ecosystem. This approach facilitates the development of mitigation strategies and regulatory conditions, emphasizing local management of liability within the layers of the technology stack. The safety implications are multi-faceted, encompassing the protection of citizens, ecosystems, and software components. Citizen interactions with apps are traceable and safeguarded by authentication tokens that validate identity and access permissions. The ecosystem's safety is ensured through complex adaptive system models, categorizing sectors and apps by risk level and using these classifications for system modelling. The Confidential Clean Rooms (CCR) framework's sandbox testing layer can play a crucial role in validating model safety against data bias and other risks. This comprehensive traceability and liability framework is crucial for preventing or mitigating negative societal impacts from app deployment and use.

AI's pervasiveness necessitates clear liability pathways. Ensuring that responsibility is attributed correctly is fundamental to maintaining public trust and ensuring ethical AI deployment.

- **Predefined Liability Protocols:**

- Draft clear legal frameworks detailing liability in case of AI system malfunctions or unintended outcomes.
- Clearly delineate responsibilities among developers, operators, and end-users of AI systems.
- Embedded traceability mechanisms.

- **Incident Reporting:**

- Mandate standardized incident reporting protocols to document any system aberrations or failures.
- Establish investigation mechanisms to delve into these incidents, ensuring lessons are learned, and system robustness is improved.

e. **Specialized, Agile Regulatory Body**

Given AI's rapid progression, traditional regulatory methods, often mired in bureaucratic delays, might prove inadequate. Herein lies the necessity for a specialized AI regulator, which has been central to the global conversation on AI regulation^{88,89}. This dedicated body, conceptualized as a dynamic task force, is not only equipped with the expertise but also endowed with a broad mandate. With the agility to swiftly adapt to the ever-evolving technological landscape, such a regulator can proactively identify potential pitfalls and recalibrate regulations in real-time, ensuring that governance remains both relevant and robust. Traditional regulatory mechanisms often lag the rapid pace of AI evolution. A dedicated, agile, and expert regulatory body is pivotal to bridging this gap, ensuring that governance remains proactive and effective.

⁸⁸ Levin, B. and Downes, L. (2023). Who Is Going to Regulate AI? [online] Harvard Business Review. Available at: <https://hbr.org/2023/05/who-is-going-to-regulate-ai>.

⁸⁹ Finocchiaro, G. The regulation of artificial intelligence. *AI & Soc* (2023). <https://doi.org/10.1007/s00146-023-01650-z>

- **Broad Mandate:**
 - Endow the regulatory body with a wide-ranging mandate, allowing it to swiftly address emerging challenges without bureaucratic delays.
- **Dynamic Adaptation:**
 - Equip the regulator with tools and methodologies to continuously scan the AI horizon, identifying potential pitfalls or areas requiring regulatory attention.
 - Encourage a feedback-driven approach, where the regulator engages with industry and academia to refine its directives continually.
- **Continuous Monitoring:**
 - Employ real-time monitoring tools to track AI system behaviours against the set boundaries.
 - Integrate automated alert systems to notify of potential boundary breaches or aberrant behaviours.
- **National Registry:**
 - Establishing a centralized database for AI algorithms is crucial for regulatory compliance and fostering innovation in the field of artificial intelligence.
 - A separate national registry of unforeseen or dynamic non-standard events during the functioning of AI models can be made to create necessary feedback loops for the regulator.

V. Financial Markets as a Complex Adaptive System

Financial markets exemplify intricate complex adaptive systems^{90,91}, with diverse human and algorithmic traders constantly interacting and reacting to emerging signals. As traders adapt strategies in response to price fluctuations, news developments and each other's actions, it creates nonlinear feedback loops and unpredictable turbulence. Emergent behaviours arise from diverse interacting agents^{92,93}. Yet despite this inherent chaos, proactive governance approaches demonstrate that regulating complex market dynamics is feasible.

Dedicated regulatory bodies like SEBI (Securities and Exchange Board of India) and SEC (Securities and Exchange Commission) in the United States provide specialized oversight honed over decades of financial evolution. Equipped with broad mandates and continually engaging with market complexities, they develop tailored interventions balancing stability and progress.

Strict protocols govern automated algorithmic trading to contain risks, while still permitting innovation. Requirements of transparency through regular reporting of systemic changes, quarterly financial audits, and disclosure standards enhance traceability. This is akin to mandating algorithmic explainability and audits for artificial intelligence systems.

When volatility escalates beyond thresholds, automatic circuit breakers act as control chokepoints to temporarily halt trading and regain equilibrium. Constraints like margin requirements, settlement periods and exposure limits curb excessive speculation and leverage. Graded warnings, fines and penalties deter violations. Insolvency and bankruptcy laws provide ultimate overrides when entities become non-viable.

Fixing liability is crucial; regulations hold individuals and corporate boards personally accountable for corporate actions through codes of conduct and controls like audit committees. Imposing 'skin in the game' for market participants aligns risks and rewards. Licensing and surveillance regimes weed out misconduct.

Of course, differences exist between AI systems and financial markets – technologies like recursive self-improvement have no parallel in capital markets. However, regulating complex, emergent behaviours arising from webs of interdependencies does share useful principles. Leveraging expertise, transparency, control points and accountability.

Regulation centred on expertise, transparency, control points, and accountability helps manage market complexities essential for economic functioning and has proven effective for market oversight across decades of turmoil. Adapting such insights to guide AI governance can similarly help steward progress responsibly. Building guardrails suited to AI's specific intricacies will be vital as capabilities scale exponentially.

⁹⁰ James, Paulin., Anisoara, Calinescu., Michael, Wooldridge. (2018). Agent-Based Modeling for Complex Financial Systems. *IEEE Intelligent Systems*, 33(2), 74-82. Available from: 10.1109/MIS.2018.022441352

⁹¹ M., Flood., Alexander, Lipton., D., Quintana., Antoaneta, Serguieva. (2018). Guest Editorial Special Issue on Complex Systems in Finance and Economics. *IEEE Systems Journal*, 12(2), 1087-1089. Available from: 10.1109/JSYST.2018.2817978

⁹² Markose, S.M. (2005). Computability and Evolutionary Complexity: Markets as Complex Adaptive Systems (CAS). *The Economic Journal*, 115(504), pp.F159–F192. doi: <https://doi.org/10.1111/j.1468-0297.2005.01000.x>.

⁹³ Polacek, G.A., Gianetto, D.A., Khashanah, K. and Verma, D. (2012). On principles and rules in complex adaptive systems: A financial system case study. *Systems Engineering*, 15(4), pp.433–447. doi:<https://doi.org/10.1002/sys.21213>.

Conclusion

The modernity of our age has become synonymous with the far-reaching technological advancement in Artificial intelligence (AI). While countries are attempting to regulate the untamed expansion and implications of AI through various mechanisms, static and linear regulations have shown limited scope for dynamically evolving AI.

The absence of an internationally agreed framework also leaves a lot of room for multiple regulations to emerge across the regulatory spectrum, as discussed in the case of the contemporary policy framework landscape that encompasses a wide range of control mechanisms, each presenting its own set of challenges. On one end of the spectrum, there is the 'hands-off' and self-regulation model adopted by the USA, and on the other end, the UK's pro-innovation, still a laissez-faire stance. Both approaches carry the significant risk of ineffective regulation, especially in the context of AI's extensive and unpredictable impacts. The European Union's risk-based approach presupposes the feasibility of categorizing risks in a system that is inherently non-deterministic. Concurrently, China's method of state-mandated direct bureaucratic control has great pitfalls, as evidenced by the potential lab leak scenario during the COVID-19 outbreak, which highlighted the possibility of widespread repercussions stemming from such a governance model.

As an effective approach, we consider AI to be a complex adaptive system (CAS) with unpredictable emergent behaviours arising from the interactions of multiple components like machine learning and neural networks. Just as small changes can cascade in unpredictable ways in CAS, small tweaks to AI systems can compound into disproportionate impacts that are hard to model long-term. AI development is an evolutionary process with continual interplay between variations in technical components. We are agnostic to risk assessment and predicting the AI evolution cycle. In addition to this, the burden of predictability of pattern and risk does not exist in this regulatory principle, thereby making it truly agile and independent of human-induced bias.

Viewing AI through the lens of a CAS offers valuable insights for effective governance. Effective regulation of AI as a CAS involves instituting real-time guardrails rather than attempting to predict outcomes. The proposed regulatory approach, based on CAS principles, highlights the need for robust 'Guardrails' and 'Partitions' to define operational boundaries and prevent systemic risks. These measures ensure AI technologies do not exceed their intended functions and are insulated from causing widespread dysfunction. The concept of a regulatory sandbox, as demonstrated in the fintech sector, is advocated for the safe development of AI under regulatory supervision. This includes setting clear technical thresholds and partitioning strategies for AI systems and employing containerization techniques for isolated execution environments.

Multi-factor authentication and hierarchical governance provide robust check and balance systems. The CAS-based framework emphasizes the critical role of human control in AI governance, advocating for manual overrides and authorization chokepoints to maintain human agency, especially in critical infrastructure domains. Manual overrides allow human intervention in unpredictable AI behaviours, while multi-factor authentication protocols ensure a robust check and balance system for high-risk actions. The approach also stresses the importance of transparency and explainability in AI, promoting open licensing of core algorithms and the creation of AI factsheets for system audits by external experts (demystifying the "black box" nature of AI) and accountability. In addition to this, the framework calls for distinct accountability lines and predefined liability protocols to ensure responsible AI deployment.

Governing Principles

1. Guardrails and Partitions/ Fire Breaks & Choke Points
2. Circuit Breakers/ Manual Overrides
3. Transparency & Explainability
4. Distinct Accountability
5. Specialized, Agile Regulatory Body

The necessity of a specialized, agile regulatory body is underscored, one that can adapt quickly to AI advancements, continuously monitor AI systems and maintain a national registry for regulatory compliance and fostering innovation. A dynamic, informed regulator can scan the horizon to recalibrate regulations in real-time. This comprehensive approach aims to create a governance framework that is resilient, responsive, and responsible, balancing regulatory oversight with technological progress.

Insights from governing chaotic systems like financial markets demonstrate feasible regulation approaches for complex technologies. Dedicated oversight bodies provide specialized expertise. Requirements like audited financial statements enhance transparency, similar to explainability standards for AI. Circuit breakers act as control chokepoints, while liability and skin in the game ensure accountability. This approach has many implications. For instance, we may never allow a super-connected internet of everything. Those creating AI tools will not be let off easily for supposed unintended consequences – thereby inserting an ex-ante "Skin in the Game". Humans will retain override and authorization powers. Regular mandated audits will have to enforce explainability.

We hope that this paper provides a new intellectual framework for regulating a powerful and rapidly developing field of technology. As illustrated, the CAS framework is already used to regulate other sectors with similar characteristics and has provided workable solutions.