FutureMatters™
Centre of Excellence

PRIMER ON ARTIFICIAL INTELLIGENCE: GOVERNING SELF-LEARNING INTELLIGENT SYSTEMS

A Knowledge Guide to Forum Programming

March 2024

MAIN CONTRIBUTOR

Jaskaran Singh Bhalla
Head of Content, Elevandi

CONTRIBUTOR

Rebecca Martin
Senior Director, Elevandi

POINT ZERO FORUM™
Policymakers | Leaders | Investors
1 - 3 JULY 2024
Zurich, Switzerland

TABLE OF CONTENTS

FutureMatters™
Centre of Excellence

INTRODUCTION

Since the coining of the term Artificial Intelligence (AI) by John McCarthy in 1956, the field has evolved incrementally over the 7 decades that followed through predictive algorithms that can recognise patterns and predict future events. In that context, **the relatively recent introduction of the text, image, audio, and video creating Generative AI (GenAI) has been truly revolutionary**: GenAI has the potential to increase the GDP of the world by USD 6.1 – 7.9 trillion[1], increase developer productivity by 30%[2] and is on track to becoming an integral element in our professional lives with multiple Copilots[3] launched in 2023-2024.

The technology and banking sectors are in pole position to benefit from GenAI[4]:

Generative AI productivity impact by industry and business function

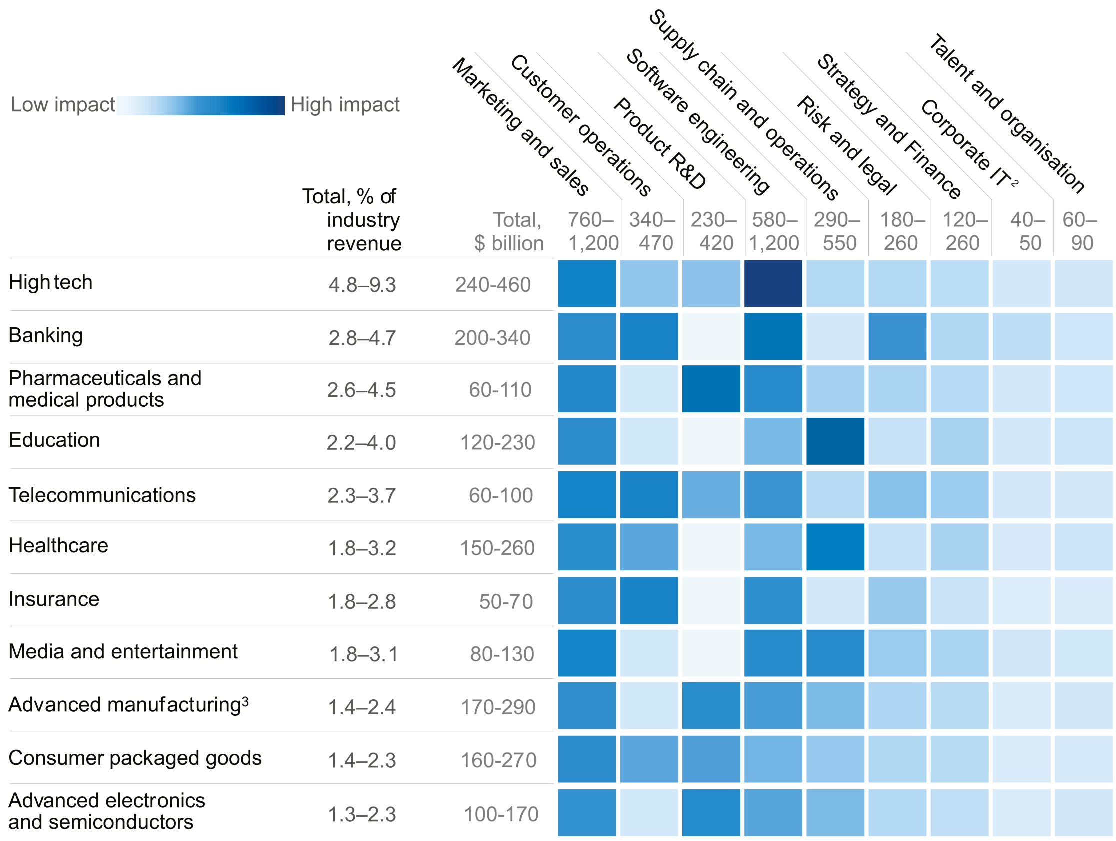

Figure 1: GenAI's impact on business functions across industries[5]

However, the full realisation of the potential of GenAI is obstructed by a myriad of challenges. **Questions remain regarding the veracity, availability and ownership of the data required to power GenAI models, the imperative for dependable and equitable access to computing resources, and the looming spectre of job displacement** (projected at 40% for all AI use cases by the International Monetary Fund[6]).

CHAPTER 1: The Development of AI

Brief History of AI (1950 – 2010)

Since the introduction of AI in the 1950s, the field has gone through ebbs and flows in the 7 decades that followed in its pursuit of replicating the cognitive power of humans[7]:

Brief History of Artificial Intelligence (1950 – 2010)[8]

1950 - 70: AI emerges as a field due to the convergence of **algorithms, hardware and automation** (which were all accelerated due to World War II)

1950: Alan Turing proposes the question, "Can machines think?" in his article, *Computing machinery and intelligence.*

1956: MIT's John McCarthy assigns the term AI to computer programs that can engage in task requiring high-level mental processing. Following earlier funding from US Gov in 1963, MIT's Joseph Weizenbaum creates the first chatbot ELIZA in 1966.

1959: An IBM researcher, Arthur Samuel, coined the term *Machine Learning* for the ability of computers to learn without being explicitly programmed, inspired by the first self-learning algorithm created in 1958.

1970 - 90: Period of **AI winters** where the initial hype driven by scientist predictions could not survive the poor results produced by very expensive and ineffective early AI systems. This led to a combination of decline in funding and failed commercial projects.

1973: British Science Research Council published the Lighthill report which presented AI in a pessmistic light and stated that "In no part of the field have discoveries made so far produced the major impact that was then promised", in response, UK and US cut funding.

1981-92: Japan spent more than USD 400 million on a commercially failed project called Fifth Generation Computer Systems which aimed to develop computers with reasoning ability.

1990 - 2010: AI Expert systems, created to mimic decision making process of human experts

1997: World Chess Champion Gary Kasparov was defeated by IBM's Deep Blue.

2003: NASA landed two rovers onto Mars that navigate the surface of the planet without human intervention.

Evolution of Computer Hardware[9]

Initially, all computer processors had their logic programming custom-built into their physical architecture, an approach that made their research and development prohibitively expensive, which created a shortage of engineers and made devices comparatively large and cumbersome.

1965: Intel co-founder Gordon Moore predicted that processor capacity will approximately double every 24 months.

1971: Intel launched the first microprocessor **Central Processing Unit (CPU)** built on a single chip Intel 4004, which delivered the same computing power as the first electronic computer which had filled an entire room. This general-purpose chip could be mass produced and then programmed, making computer cheaper, more powerful and smaller.

1999: NVIDIA releases the world's first Graphics Processing Unit (GPU). While a CPU handles all tasks for running computer software, a GPU is more efficient at handling complex mathematical operations by running them in parallel. In 2009, Stanford scientist Andrew Ng showed that using GPUs is 70 times faster than using CPUs for certain AI tasks.

2002: Amazon launched Amazon Web Services (AWS) which offered cloud based storage and computing power, lowering the costs of data storage and computation. Google and Microsoft release their cloud offerings in 2010.

The success of **rule-based AI expert systems** from 1990 onwards was mainly due to the continued improvements in **computing hardware** that enabled the storage and execution of extensive libraries of instructions. For example, IBM's Deep Blue evaluated 200 million chess positions per second due to its brute computing prowess.[10] Even though these AI expert systems perform well in automating human tasks, they have key limitations, such as the inability to learn independently and the need for structured data.

Evolution of Computer Software

The interest in **self-learning algorithms** dates back to 1950-65, a period which saw great enthusiasm in **Machine Learning** (ML). Inspired by how a network of neurons in the human brain functions[11], psychologist Franck Rosenblatt in 1958 proposed the perceptron algorithm which is a simple network of a single neuron and could distinguish cards marked on the left side from those marked on the right side[12]. In 1965, scientists Ivakhnenko and Lapa proposed the first general working learning algorithms with multilayer artificial neural networks[13]. Till date, the Artificial Neural Network (ANN) is the most popular ML model. Even though the algorithmic concepts for machines to self-learn were introduced in the 1960s, the data, techniques, and computing power to train these networks were not ready till an inflexion point around 2010.

The Deep Learning Era (2010 onwards)

The introduction of the Internet in 1991, followed by the mass proliferation of broadband, the personal computer, and the smartphone in the 2000s, solved the problem of **data** needed to train self-learning ML models. Advancements in data availability and hardware capacity fuelled the realisation of the subset of ML called **Deep Learning**. Deep learning uses layers of connected neurons (or software-based calculators[14]) that are especially effective at learning from unstructured data such as images, text, and audio. **Deep learning models solve the problems posed by AI expert systems: they don't need to be built as an extensive rule library and can use unstructured data to solve previously unseen problems.** Deep learning models accelerated the ability of AI to recognise and understand language, handwriting, speech, and image, and now AI systems can perform these tasks better than humans (benchmarked using standardised tests):

Language and Image Recognition Capabilities of AI Systems[15]

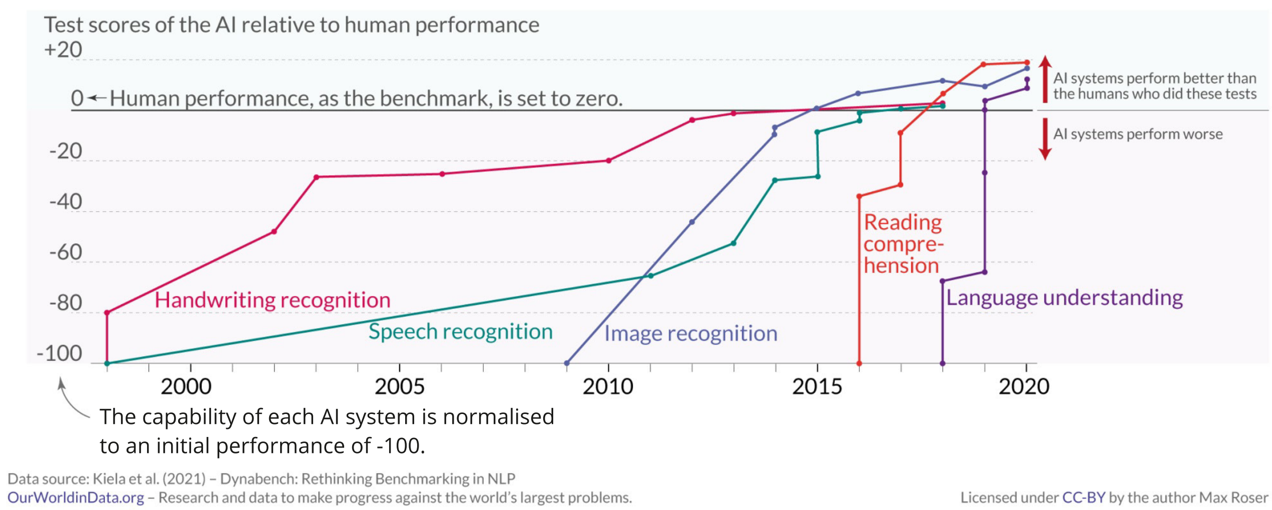

DEEP LEARNING ERA

FutureMatters™ Centre of Excellence

GenAI emerged as a subset of deep learning, where **Foundation Models (FM)** (such as GPT-4) trained on very vast and varied types of unstructured data can now generate text, image, audio, and video. FMs, different from earlier deep learning models that can often perform only one task, are versatile and allow businesses to solve various business cases with the same one model[16]. A subset of FMs is a **Large Language Model** (LLM) that can generate text by processing massive amounts of text and learning the relationships between words. **Fine-tuning** is the process of adapting pretrained FMs to perform better by understanding nuances, terminology, and specific patterns for a specific use case.

Relationship between AI, ML, Deep Learning and GenAI

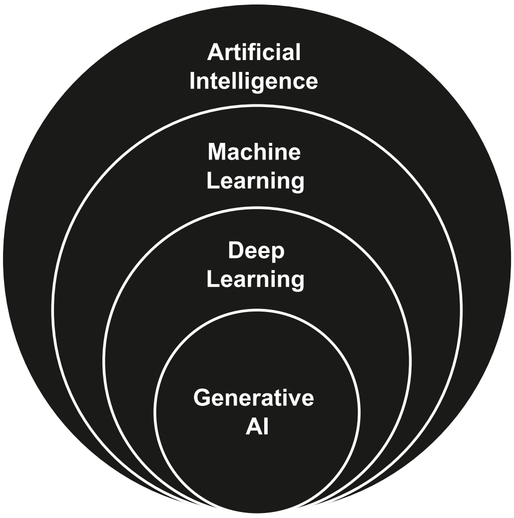

Progress of Deep Learning Models and Their Training Computation (2010 onwards)[17]

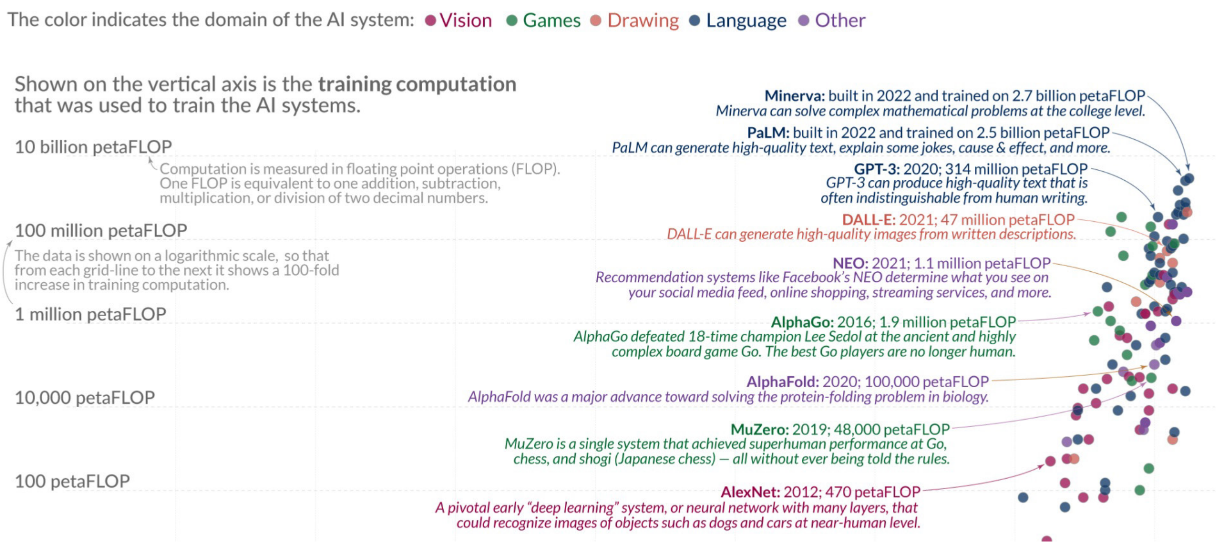

The Data and Hardware Supply Imbalance

With various open-source FMs available for use, a greater detail of attention is being paid to the other two enabling factors: the data and the hardware. **It is critical to have a resilient supply of both 1) data and 2) GPUs to drive the adoption of deep learning models across industries and sectors.**

However, there are challenges with both:

1) The high-quality language data to train deep learning models is predicted to run out by 2026[18]. This supply challenge is further exacerbated by questions over which data can be used to train the models and who owns it. Synthetic data is emerging as an avenue to supply reliable data without high-quality real data.

2) NVIDIA has a monopoly over GPUs, which has quadrupled its stock since the start of 2023. On 15 February 2024, the company rose to number four on the world's most valuable companies, overtaking Amazon and Alphabet. Against the backdrop of the macroeconomic and geopolitical landscape of 2024, the strategic significance of GPU supply has been further heightened, emphasizing the importance of resilient supply chains and the mitigation of concentration risks.[19]

NVIDIA's meteoric rise since January 2023 [19]

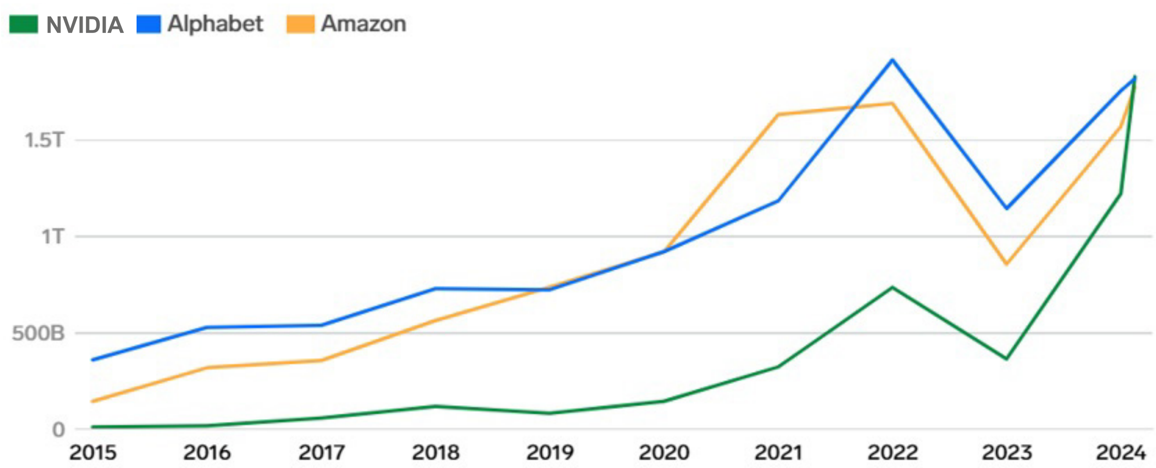

There is an increased spotlight on key global manufacturing hubs that create GPU components (like semiconductors) and trade routes that ship them between stages of the manufacturing lifecycle. To counter these forces, actions are underway to diversify the supply chain and bake resiliency into the availability of 'the gold of AI'.[20]

The Global Semiconductor Trade[21]

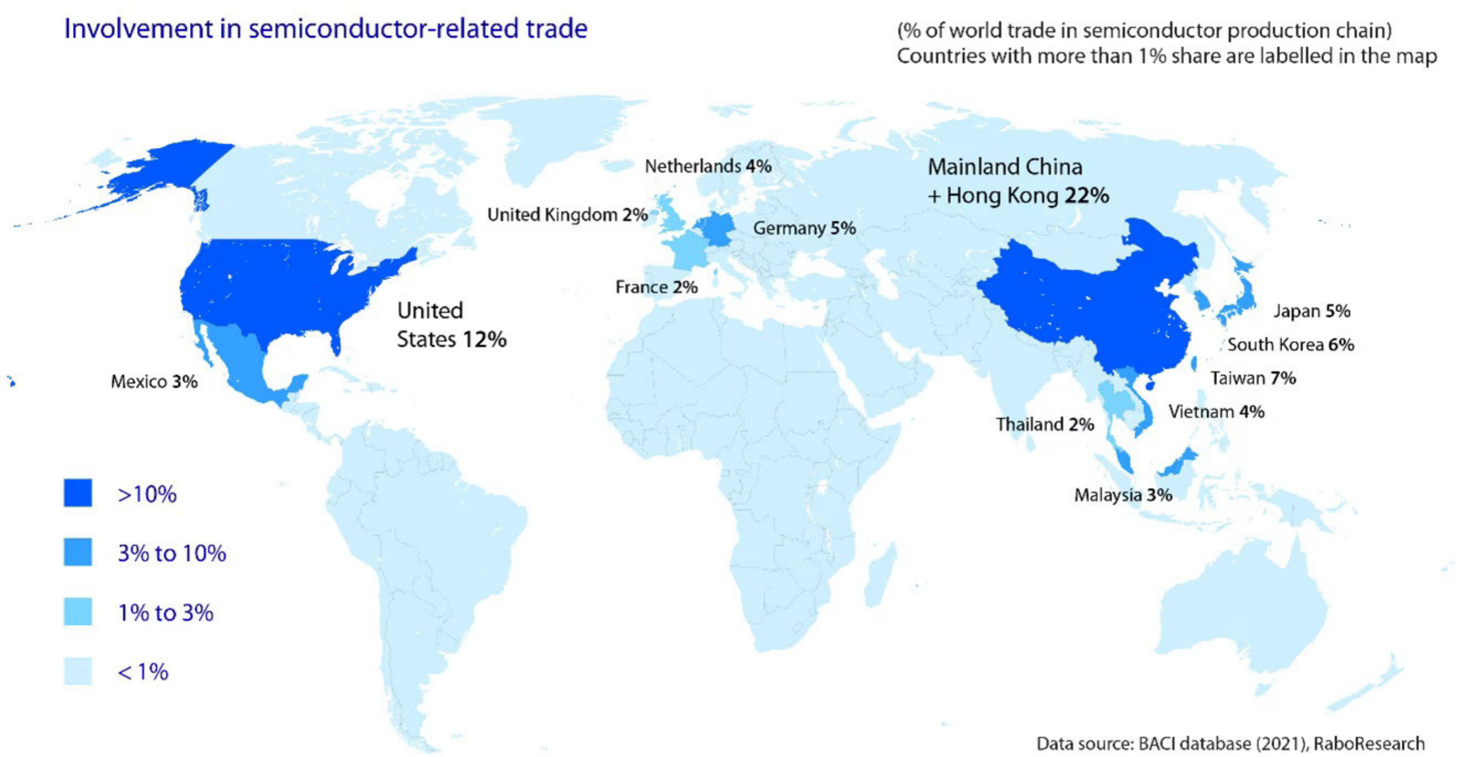

CHAPTER 2: Globally divergent approaches to AI governance

With the expansion of global dependence on AI, there has been a naturally correlated increase in policymaker interest in AI – the number of bills containing AI increased from 1 in 2016 to 37 in 2022.[22] The differing approaches to governing AI include risks-based, rules-based, principles-based, and outcomes-based.[23]

Divergent approaches to AI governance across the globe[24]

	Risks-based	Rules-based	Principles-based	Outcomes-based
Definition	Focuses on classifying and prioritising risks in relation to the potential harm AI systems could cause	Lays out detailed and specific rules, standards and/or requirements for AI systems	Sets out fundamental principles or guidelines for AI systems, leaving the interpretation and exact details of implementation to organisations	Focuses on achieving measurable AI-related outcomes without defining specific processes or actions that must be followed for compliance
Benefits	– Tailored to application area – Proportional to risk profile – Flexible to changing risk levels	– Potential reduction of complexity – Consistent enforcement possible	– Intended to foster innovation – Adaptable to new developments – Can encourage sharing of best practices	– Can support efficiency – Flexible to change – Intended to foster innovation – Compliance can be cost-effective
Challenges	– Risk assessments can be complex – May create barriers to market entry in high-risk areas – Assessment and enforcement can be complex	– Rigidity can increase compliance costs – May be unreliable to enforce	– Potential inconsistencies with interpretation of principles – Unpredictable compliance and impractical enforcement – Potential for abuse by bad actors	– Scope of measurable outcomes can be vague – Potential for diffused accountability – Limited control over process and transparency
Example	EU: Artificial Intelligence Act, 2023 (provisional agreement)	China: Interim Measures for the Management of Generative AI Services, 2023	Canada: Voluntary Code of Conduct for Artificial Intelligence, 2023	Japan: Governance Guidelines for Implementation of AI Principles Ver. 1.1, 2022

While there are differing approaches to governing AI, the policymaking community is increasingly converging on a set of core principles[25]:

1. **Safety and security**: Focused on physical safety, prevention of harm and protection against malign attacks (notable incident: USD 25 million lost to a deepfake attack[26])
2. **Inclusivity and Fairness**: Prevention of unjust bias against one or more groups (notable incident: Google pauses Gemini image creation[27])

3. **Transparency and explainability**: Transparency and explainability of an AI model's purpose, scope, limitations and how it is trained (notable incident: US Senate AI Insight Forum Written Statement in Nov 2023[28])

4. **Respect for Intellectual Property (IP) considerations**: Respect for IP rights over the data used for model training and clarity on IP rights of the output from the AI model (notable incident: New York Times sues OpenAI for use of copyrighted data[29])

5. **Ensuring fair competition**: Avoiding excessive dependency on specific providers of AI models or associated infrastructure (notable incident: US FTC launches inquiry into the impact of Generative AI Investments and Partnerships of Microsoft, Amazon, and Google on the competitive landscape[30])

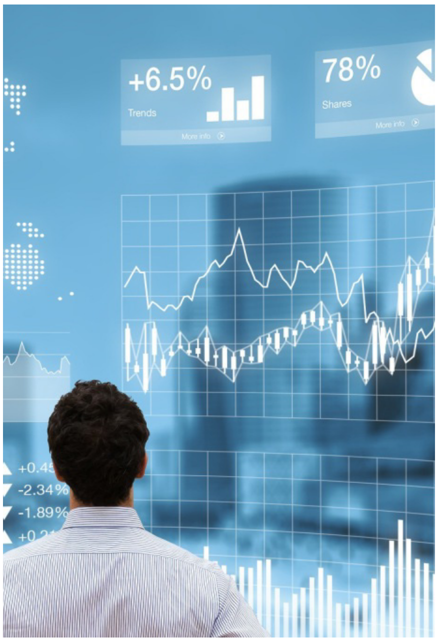

Notable mention of regulatory action: EU is leading the way with its AI Act which assigns applications of AI to three risk categories. First, applications and systems that create an unacceptable risk, such as social scoring, are banned. Second, high-risk applications, such as a CV-scanning tool that ranks job applicants, are subject to specific legal requirements. Third, applications not explicitly banned or listed as high-risk are largely left unregulated.

Notable mention of harmonisation efforts: The UK organised the AI Safety Summit on the 1 and 2 November 2023 at Bletchley Park, Buckinghamshire and brought international governments, leading AI companies, civil society groups and experts together to discuss how risks rising from the "Frontier" of AI (i.e. GenAI) can be mitigated through internationally coordinated action. The Bletchley Declaration was signed by 28 attending major countries to support an internationally inclusive network of scientific research on frontier AI safety that encompasses and complements existing and new multilateral, plurilateral and bilateral collaboration.[32]

How do deep learning models impact the governance discussion?

Given deep learning models (like FMs) utilise vast amounts of unstructured data, there are certain nuances to the governance of self-learning intelligent systems[33] (as compared to AI expert systems):

1. Bias risk increases if training data is biased, or the underlying FM is developed unfairly.
2. Explainability and transparency risks increase as models are not based on a set of rules.
3. Reliability risk increases as models can produce different answers to the same prompt.
4. Privacy risks increase as models could include prompt information in output for other users (as they self-learn and improve), enabling the identification of users.
5. IP risks increase as vast quantities of data are used to train and fine-tune models.

CHAPTER 3: Adoption of GenAI in Financial Services

The adoption of GenAI within financial services is broadly still at the level of experiments and testing, mainly due to a lack of regulatory clarity and concerns around security, reliability, and bias.

GenAI use cases at financial institutions with indicative examples[34]

COMMERZBANK	**Improving customer experience:** Commerzbank will use Microsoft's Azure OpenAI service to build and launch an **AI Banking Avatar** which combines GenAI and Avatar technology for customer facing applications.
Morgan Stanley	**Creating digital experts:** Morgan Stanley launched the **"AI @ Morgan Stanley Assistant"** in September 2023. The tool gives **16,000** financial advisers the ability to quickly find and synthesize answers reducing their time spent on documentation and enhancing operational efficiency.
Goldman Sachs	**Delivering software faster:** Goldman Sachs is working on a GenAI project that writes code in English-language commands and generates documentation.
BIS	**Obtaining real-time estimates of inflation expectations:** The Bank of International Settlements is working on a Central Bank Language Model (CB-LM) project which adapts LLMs to central bank terminology and increase accuracy from **50–60%** to **90%.**
	Enhancing real-time transaction monitoring: Mastercard's "Decision Intelligence Pro" technology uses GenAI to scan one trillion data points in real time to predict whether a transaction is likely to be genuine or not - this boosts fraud detection rates on average by as high as **300%** in some instances.
EUROPEAN CENTRAL BANK	**Aiding central bank employees to create and assess information quickly:** The European Central Bank (ECB) is exploring multiple use cases of GenAI for internal efficiencies.
Capital One J.P.Morgan	**Augmenting fraud and suspicious activity detection systems:** Capital One and JPMorgan Chase have leveraged GenAI to significantly reduce false positives, improve detection rate, reduce costs, and improve customer satisfaction.

CHAPTER 4: The AI Opportunity in 2024

To truly harness the potential of AI and drive the deep and pervasive adoption within the financial services industry, the following actions are required:

1. **Policy Action**

   Global organisations need regulatory clarity to implement AI roadmaps: Various jurisdictions have issued AI guidelines and regulations (both generic and sector-specific), resulting in organisations facing diverging requirements before implementing AI.

2. **Technology Action**

   Clear technology roadmaps are needed to overcome legacy data and technology infrastructure debt: To utilise internal data to build industry specific use cases and integrate AI with existing highly customised technology infrastructure (like cloud and databases), organisations will need to devise implementation roadmaps which overcome legacy challenges.

3. **Hardware Supply Resilience**

   Regional, national, and business strategies need to create resilient chip supply: There is an opportunity to create resilient supply chains that ensure access to chips and their components in the face of 1) concentration risks and 2) geopolitical, macroeconomic and trade tensions. (notable investments: The US Congress announced the CHIPS Act of 2022 to strengthen semiconductor manufacturing, research, and design[35], TSMC opens first Japan Factory[36], Japan to invest USD 67 billion in chips[37], Intel to invest over EUR 30 billion in Germany for chip making[38]).

4. **Alignment on Data Usage**

a) **Data owners need to work with policymakers as well as business and technology leaders to firm up data usage guidelines.** Clear guidelines need to be in place for technologists to be able to use both internal and external sources of data without worrying about Intellectual Property infringement.

> "Data owners need to work with policymakers as well as business and technology leaders to firm up data usage guidelines"

b) **Technology strategies will need to incorporate synthetic data**: In the absence of new sources of real data, the demand for synthetic data will go up.[39] Apart from using synthetic data for training, trained GenAI models are well-suited to produce synthetic data for all analytical use cases.

5. Ensuring Safe, Inclusive and Ethical Adoption of AI

a) **All LLMs should be fair and transparent by design:** Technologists should ensure that the models created have sufficient transparency and explainability. Policymakers should ensure that in cases of unethical usage, there is a clear matrix of accountability and ownership.

b) **Greater cybersecurity efforts are required to protect the digital economy against deepfake and other fraud**: GenAI has bolstered cyber attackers by providing them the ability to convincingly impersonate personalities and use deepfakes across new channels (such as conference calls). These attacks grew by over 700% from the first to the second half of 2023[40].

c) **Ecosystem strategies need to provide options for smaller organisations to implement AI:** To ensure that smaller organisations with limited computational power can also harness the full potential of AI, there is a need to work on public utilities[41] and software which increase efficiency of existing hardware.[42]

d) **The AI revolution needs to include the Global South**: Efforts need to be made that the global south has a fair access to AI infrastructure and is fairly represented in the data that is used for training LLMs. *[This topic will not be a key discussion item at Point Zero Forum and will instead be discussed at Elevandi's platforms in Africa – 3i Africa Forum in Ghana (May 2024) and Inclusive FinTech Forum (August 2024).]*

"Greater cybersecurity efforts are required to protect the digital economy against deepfake and other fraud"

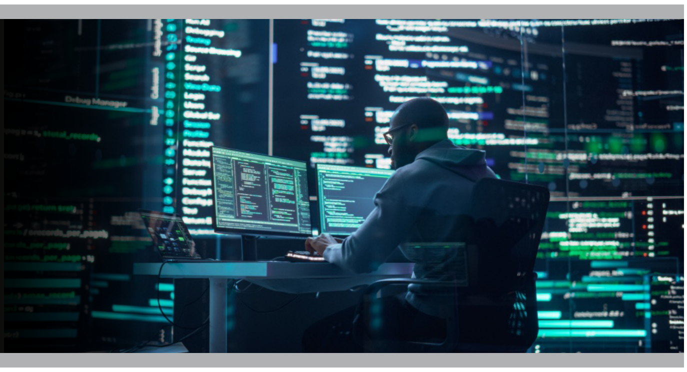

6. **Alignment of Long-Term Policy, Business, Technology and Capital Roadmaps**

   Global public and private sectors need to collaborate on investing capital and resources in AI research: While understanding near-term roadmaps will help organisations move forward, there needs to be collaboration on the long-term roadmap to harness the full potential of AI.

CHAPTER 5: The Future of AI

Foreseeing the future of any technological force is a challenge and when a field like AI is evolving at an unprecedented pace, it becomes even more daunting. Instead of peering through a crystal ball and forecasting what AI will look like 10 or 20 years from now, below are some factors to keep in mind and best prepare for a *VUCA* (volatile, uncertain, complex, and ambiguous) future:

1. **There is a need to balance optimism with caution.** We have a precedent: scientists like Marvin Minsky (an early AI pioneer) inaccurately predicted in 1970 that "*from three to eight years we will have a machine with the general intelligence of an average human being*"[43] and heralded an **AI winter** due to this over-optimistic promise being unfulfilled. Hence, **there is a risk of a GenAI winter being ushered in due to the costs of implementing GenAI outweighing the positive impact promised**.
2. **AI governance efforts will need to converge and might be driven by global organisations like the United Nations**[44], as AI's perceived threat becomes far-reaching and is considered by some sections of society an existential risk possible of causing human extinction at par with a pandemic or a nuclear war.[45]
3. **AI is a field made possible by the amalgamation of various independent realms** such as hardware, software, and data, coming together. Just like the varying degrees of evolution of these spheres brought to us the last 7 decades of AI, the next decades will depend upon how they each shape up and how their intersections play out. For example, **quantum**

"AI has the opportunity to bridge the current digital divide"

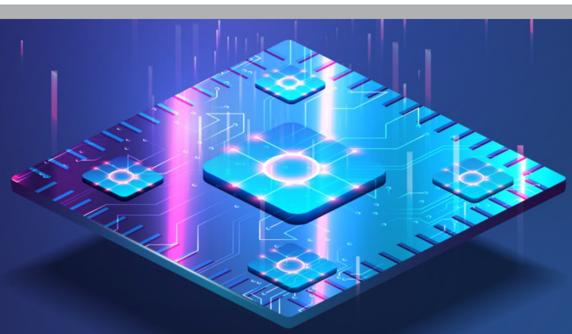

computing might unlock the next step change in computing power to make deep learning models even easier to train by reducing the need for data by 10 times to 300 times.[46]

4. **AI will continue to exist in a world constantly evolving** and there will be dependencies on the forces such as blockchain and the transition to a net-zero world. At a more philosophical level, there will also be a need to further collaborate with fields such as psychology and neuroscience to understand better AI's impact on the human mind and society.

5. **AI has the opportunity to bridge the current digital divide** caused by the internet being run in a few languages like English, Spanish, and Mandarin. The quality, quantity, and type of data used to train deep learning models are key for AI's success and will remain key. As the data generated on the internet used to train deep learning models is limited to a few languages, there is a need to ensure greater linguistic representation of the world's 7000 languages[47].

6. **The future of the workforce calls for augmenting human capabilities through partnerships with machines, not replacement.** Even though the early GenAI models have shown promise by passing the bar exam, negotiating contracts, and answering medical questions[48], there is still a need to ensure these outputs are trustworthy, safe, and ethical.

CHAPTER 6: Coverage at Point Zero Forum 2024

POINT ZERO FORUM™
Policymakers | Leaders | Investors

GenAI coverage at Point Zero Forum 2024

Masterclass on AI: From Turing to Transformer

- How did Deep Learning emerge from the evolution of computer hardware, software and data?
- What is the impact of geopolitics, macroeconomics and trade fractures on tech adoption?

Masterclass on Quantum Technologies: Decoding the Future of Computing

- What are the fundamental principles of quantum computing, and how has its evolution unfolded to its current state?
- How are quantum computing applications currently shaping financial services, and what developments should organisations consider for strategic future-proofing?

AI Under the Scanner: Global Policymakers' Dialogue on Creating Guardrails for AI

- What is causing policymakers to consider regulations?
- What efforts are underway to harmonise globally divergent AI regulations?

Intelligent & Auto-Learning Financial Systems

- What are the challenges faced by financial institutions in the current implementation of AI, and what strategies are being employed to overcome these hurdles?
- How can financial institutions develop comprehensive roadmaps and blueprints to maximise the benefits of AI while addressing regulatory, ethical, and operational considerations?

A Philosophical Dialogue on the Impact of AI and Quantum Computing on Human Society

- How do advancements in AI and quantum computing challenge traditional notions of human agency and decision-making processes?
- How should society navigate the potential convergence of AI and quantum computing to ensure responsible outcomes?

ABOUT THE AUTHORS

FutureMatters™
Centre of Excellence

Jaskaran Singh Bhalla
Head of Content, Elevandi
Based in Singapore

I write about Financial Services, Policy, and Technology

Rebecca Martin
Senior Director, Elevandi
Based in Singapore

I write about Entrepreneurship, Financial Services, and Technology

The views expressed here are their own and do not necessarily reflect the views of Elevandi or its staff.

REFERENCES

FUTUREMATTERS™
CENTRE OF EXCELLENCE

[1] "The economic potential of generative AI: The next productivity frontier", McKinsey, 14 June 2023, https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier

[2] "The economic impact of the AI-powered developer lifecycle and lessons from GitHub Copilot", GitHub, 27 June 2023, https://github.blog/2023-06-27-the-economic-impact-of-the-ai-powered-developer-lifecycle-and-lessons-from-github-copilot/

[3] "Bringing the full power of Copilot to more people and businesses", Microsoft, 15 January 2024, https://blogs.microsoft.com/blog/2024/01/15/bringing-the-full-power-of-copilot-to-more-people-and-businesses/ , "Bard becomes Gemini: Try Ultra 1.0 and a new mobile app today", Google, 08 February 2024, https://blog.google/products/gemini/bard-gemini-advanced-app/ , "Apple Readies AI Tool to Rival Microsoft's GitHub Copilot", Bloomberg, 15 February 2024, https://www.bloomberg.com/news/articles/2024-02-15/apple-s-ai-plans-github-copilot-rival-for-developers-tool-for-testing-apps

[4] "The economic potential of generative AI: The next productivity frontier", McKinsey, 14 June 2023, https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier - GDP growth is based on two main types of improvements: new targeted applications of GenAI to specific business challenges and GenAI's potential impact on labour productivity across current work activities performed by the global workforce.

[5] Comparative Industry Service (CIS), IHS Markit, Oxford Economics, McKinsey Corporate and Business Functions database

[6] "AI Will Transform the Global Economy. Let's Make Sure It Benefits Humanity.", International Monetary Fund, 14 January 2024, https://www.imf.org/en/Blogs/Articles/2024/01/14/ai-will-transform-the-global-economy-lets-make-sure-it-benefits-humanity

[7] "History of Artificial Intelligence", Council of Europe, https://www.coe.int/en/web/artificial-intelligence/history-of-ai and "The History of Artificial Intelligence", Rockwell Anyoha, 28 August, 2017, https://sitn.hms.harvard.edu/flash/2017/history-artificial-intelligence/ 8 Alan Turing's "Computing machinery and intelligence", DARPA is the US agency "Defense Advanced Research Projects Agency", NY Times' article on Japan closing down FGCS, The Lighthill report on AI in 1973.

[9] Intel's History Webpage https://www.intel.com/content/www/us/en/history/virtual-vault.html 10 IBM's website page on Deep Blue, How NASA used AI for its Mars Rovers in 2003, Paper on the first chatbot ELIZA created by Joseph Weizenbaum in 1966, IBM paper on Machine Learning

[11] "Artificial Neural Networks – The Rosenblatt Perceptron", David Ibañez, 02 August 2016, https://www.neuroelectrics.com/blog/2016/08/02/artificial-neural-networks-the-rosenblatt-perceptron/

[12] "The brief history of artificial intelligence: The world has changed fast – what might be next?", Max Roser, Our World in Data, 06 December 2022, https://ourworldindata.org/brief-history-of-ai

[13] "Deep Learning in a Nutshell: History and Training", NVIDIA, 16 December 2015, https://developer.nvidia.com/blog/deep-learning-nutshell-history-training

[14] "An executive's guide to AI", McKinsey, https://www.mckinsey.com/capabilities/quantumblack/our-insights/an-executives-guide-to-ai

[15] "The brief history of artificial intelligence: The world has changed fast – what might be next?", Max Roser, Our World in Data, 06 December 2022, https://ourworldindata.org/brief-history-of-ai

[16] "An executive's guide to AI", McKinsey, https://www.mckinsey.com/capabilities/quantumblack/our-insights/an-executives-guide-to-ai

[17] "The brief history of artificial intelligence: The world has changed fast – what might be next?", Max Roser, Our World in Data, 06 December 2022, https://ourworldindata.org/brief-history-of-ai

[18] "State of AI Report 2023", Nathan Benaich and the Air Street Capital, https://www.stateof.ai/

[19] Chart: George Glover, Sources: CompaniesMarketCap and Dow Jones

[20] GPUs or Chips are also referred to as the gold of AI because they're foundational for GenAI.

[21] "Mapping Global Supply Chains – The Case of Semiconductors", Rabobank, 14 June 2023, https://www.rabobank.com/knowledge/d011371771-mapping-global-supply-chains-the-case-of-semiconductors

[22] "Artificial Intelligence Index Report 2023", Stanford University Human-Centred Artificial Intelligence, https://aiindex.stanford.edu/wp-content/uploads/2023/04/HAI_AI-Index-Report_2023.pdf

[23] "Generative AI Governance: Shaping a Collective Global Future", World Economic Forum, 18 January 2024, https://www.weforum.org/publications/ai-governance-alliance-briefing-paper-series/

[24] "Generative AI Governance: Shaping a Collective Global Future", World Economic Forum, 18 January 2024, https://www.weforum.org/publications/ai-governance-alliance-briefing-paper-series/

[25] "GLOBAL AI GOVERNANCE - CONVERGENCE OR FRAGMENTATION?", FutureMatters, Shameek Kundu, March 2024, https://www.elevandi.io/insights/global-ai-governance-covergance-or-fragmentation

[26] "Finance worker pays out $25 million after video call with deepfake 'chief financial officer'", CNN, 04 February 2024, https://edition.cnn.com/2024/02/04/asia/deepfake-cfo-scam-hong-kong-intl-hnk/index.html

[27] "Google pauses Gemini AI image generator after it created inaccurate historical pictures", CNBC, 22 February 2024, https://www.cnbc.com/2024/02/22/google-pauses-gemini-ai-image-generator-after-inaccuracies.html

[28] "Making AI more explainable to protect the public from individual and community harms" Brookings Institution, 29 November 2023, https://www.brookings.edu/articles/making-ai-more-explainable-to-protect-the-public-from-individual-and-community-harms/

[29] "The Times Sues OpenAI and Microsoft Over A.I. Use of Copyrighted Work", New York Times, 27 December 2023, https://www.nytimes.com/2023/12/27/business/media/new-york-times-open-ai-microsoft-lawsuit.htm

[30] "FTC Launches Inquiry into Generative AI Investments and Partnerships", US FTC, 25 January 2024, https://www.ftc.gov/news-events/news/press-releases/2024/01/ftc-launches-inquiry-generative-ai-investments-partnerships

[31] "What is the EU AI Act?", Future of Life Institute's EU AI Act Website, https://artificialintelligenceact.eu/

[32] "The Bletchley Declaration by Countries Attending the AI Safety Summit, 1-2 November 2023", The UK Government, 01 November 2023, https://www.gov.uk/government/topical-events/ai-safety-summit-2023

[33] "An executive's guide to AI", McKinsey, https://www.mckinsey.com/capabilities/quantumblack/our-insights/an-executives-guide-to-ai

[34] "Commerzbank will realize AI Banking Avatar", Commerzbank, 17 November 2023, https://www.commerzbank.com/en/hauptnavigation/presse/pressemitteilungen/archiv1/2023/quartal_23_04/presse_archiv_detail_23_04_109450.html "Morgan Stanley's generative AI-based virtual assistant will enhance the group's operational efficiency", Private Banker International, 10 October 2023, https://www.privatebankerinternational.com/comment/morgan-stanleys-generative-ai-based-virtual-assistant-will-enhance-the-groups-operational-efficiency, "Goldman Sachs developing dozen generative AI projects -exec", Reuters, 10 November 2023, https://www.reuters.com/technology/reuters-next-goldman-sachs-developing-dozen-generative-ai-projects-exec-2023-11-09/ , "Mastercard supercharges consumer protection with gen AI", Mastercard, 01 February 2024, https://www.mastercard.com/news/press/2024/february/mastercard-supercharges-consumer-protection-with-gen-ai/ Gambacorta, L, B Kwon, T Park, P Patelli and S Zhu (2024): "CB-LMs: language models for central banking", BIS Working Papers, forthcoming, "Careful embrace: AI and the ECB", ECB, 28 September 2023, https://www.ecb.europa.eu/press/blog/date/2023/html/ecb.blog230928~3f76d57cce.en.html

[35] "FACT SHEET: CHIPS and Science Act Will Lower Costs, Create Jobs, Strengthen Supply Chains, and Counter China", The US White House, 09 August 2022, https://www.whitehouse.gov/briefing-room/statements-releases/2022/08/09/fact-sheet-chips-and-science-act-will-lower-costs-create-jobs-strengthen-supply-chains-and-counter-china/

[36] "Chip giant TSMC opens first Japan factory as it diversifies away from Taiwan amid U.S.-China tensions", CNBC, 26 February 2024, https://www.cnbc.com/2024/02/26/tsmc-opens-new-plant-in-japan-as-it-diversifies-away-from-taiwan.html

[37] "Japan to invest over $67 billion in domestic semicon, other computer chip production, to take on Taiwan, China" Firstpost, 21 February 2024, https://www.firstpost.com/tech/japan-to-invest-over-67-billion-in-domestic-semicon-other-computer-chip-production-to-take-on-taiwan-china-13740170.html

[38] "Intel, German Government Agree on Increased Scope for Wafer Fabrication Site in Magdeburg", Intel, 19 June 2023, https://www.intc.com/news-events/press-releases/detail/1628/intel-german-government-agree-on-increased-scope-for-wafer

[39] "2024 Tech Trends", CB Insights, 06 February 2024, https://app.cbinsights.com/research/report/top-tech-trends-2024/

[40] "Deepfake 'face swap' attacks surged 704% last year, study finds", The Next Web, 07 February 2024, https://thenextweb.com/news/deepfake-face-swap-attacks-increase

[41] "Build trust with AI Verify", AI Verify Foundation, https://aiverifyfoundation.sg/what-is-ai-verify/

[42] "2024 Tech Trends", CB Insights, 06 February 2024, https://app.cbinsights.com/research/report/top-tech-trends-2024/

[43] "The History of Artificial Intelligence", Rockwell Anyoha, 28 August 2017, https://sitn.hms.harvard.edu/flash/2017/history-artificial-intelligence/

[44] "Should the UN govern global AI?", Cameron F. Kerry, Joshua P. Meltzer, Andrea Renda, and Andrew W. Wyckoff, The Brookings Institution, 26 February 2024, https://www.brookings.edu/articles/should-the-un-govern-global-ai/

[45] "An AI Pause Is Humanity's Best Bet For Preventing Extinction", Otto Barten and Joep Meindertsma, TIME, 20 July 2023, https://time.com/6295879/ai-pause-is-humanitys-best-bet-for-preventing-extinction/

[46] "Today's GenAI models are cost and energy pigs. Quantum tech could fix that.", Business Insider, 19 January 2024, https://www.businessinsider.com/lower-cost-gen-ai-models-quantum-zapata-2024-1

[47] "How language gaps constrain generative AI development", Regina Ta and Nicol Turner Lee, The Brookings Institution, 24 October 2023, https://www.brookings.edu/articles/how-language-gaps-constrain-generative-ai-development/

[48] "How to Prepare for a GenAI Future You Can't Predict", Amy Webb, Harvard Business Review, 31 August 2023, https://hbr.org/2023/08/how-to-prepare-for-a-genai-future-you-cant-predict