AI and Regulation: Partners or Adversaries?

Iro Tasitsiomi, PhD

Head of AI & Investments Data Science

T. Rowe Price

Background & Context

*With the rapid rise of Generative AI (GAI), **artificial intelligence has taken center stage for both businesses and regulators** alike. As regulations proliferate, **staying informed about AI legislative and policy developments** across various jurisdictions will (has?) become **critical & often even discussed one of the challenges and risks regarding AI adoption.***

Ideally, regulating AI should unlock its potential in a clear, risk-informed manner. Will existing & upcoming regulations achieve that goal? Or do they unnecessarily (?) risk stifling innovation and hindering growth instead?

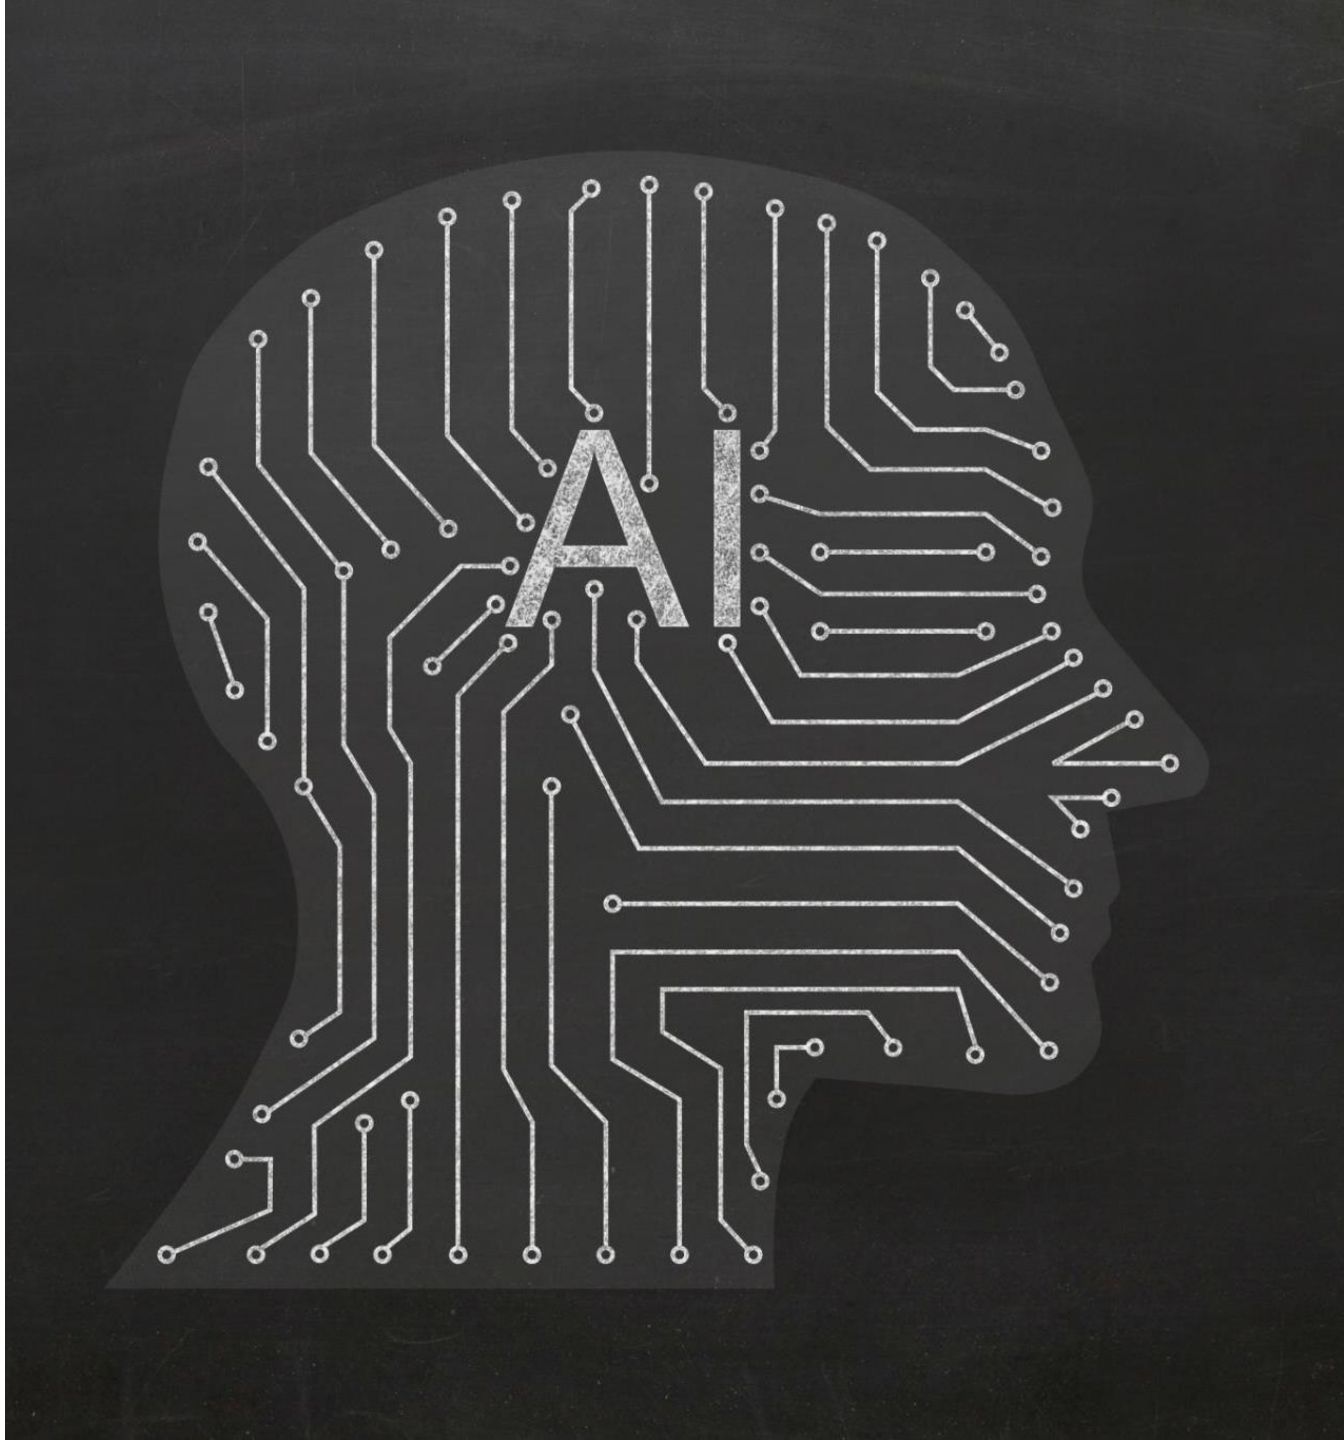

In this talk I will:

- Overview some of the AI regulatory activity we have seen in the last few years
- Identify some common themes that I think source the concerns of regulators about AI
- Discuss how several of these themes are often the result of misconceptions (we all have) about AI/GAI
- Conclude with what I think we should focus on to help regulators and businesses regulate what must and stay compliant

Current AI Regulations and Policies

Notable Jurisdictions and Their Policies

European Union AI Act

*Proposed April 2021: **first comprehensive framework for AI regulation on ethical usage, compliance and consumer protection globally**.

***Risk-based approach:** AI one of **unacceptable, high-risk, limited-risk, minimal-risk,** with regulatory measures for each.

*High-risk: **risk assessments, transparency measures, human oversight, data governance, accuracy, and robustness**. Non-compliance can result in significant fines**.**

*Creation of **European Artificial Intelligence Board (EAIB)** to ensure consistent application. National supervisory authorities will oversee compliance.

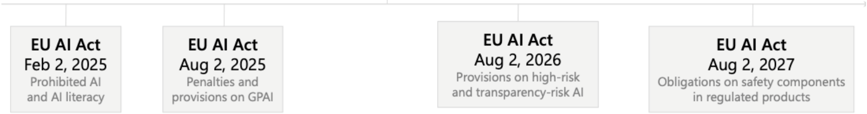

EU AI Act
Feb 2, 2025
Prohibited AI and AI literacy

EU AI Act
Aug 2, 2025
Penalties and provisions on GPAI

EU AI Act
Aug 2, 2026
Provisions on high-risk and transparency-risk AI

EU AI Act
Aug 2, 2027
Obligations on safety components in regulated products

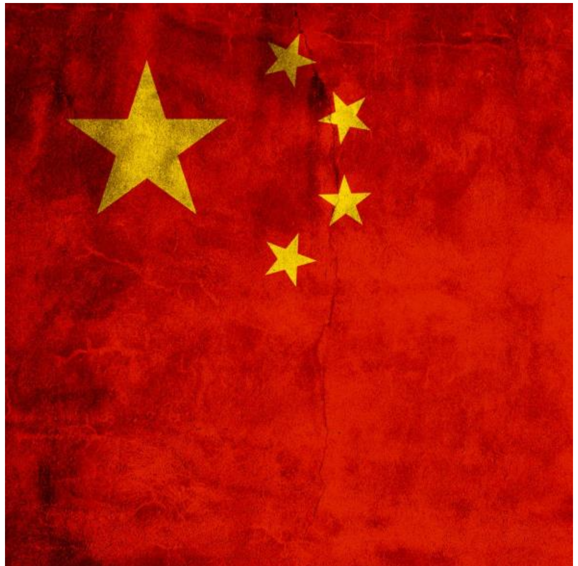

China's AI Strategy

***Three-tier system** focusing on transparency, security, and ethical alignment. Key regulations: **Recommendation Algorithm Regulation (2021)**, **Deep Synthesis Rules (2022)**, and **Generative AI Measures (2023)**.

***Centralized algorithm registry system**, **mandatory security self-assessments**, enforcement by **multiple agencies**

***Ethical review protocols, content control mandates, and transparency requirements**, aiming to balance innovation with state control, while also influencing global AI standards.

The USA landscape

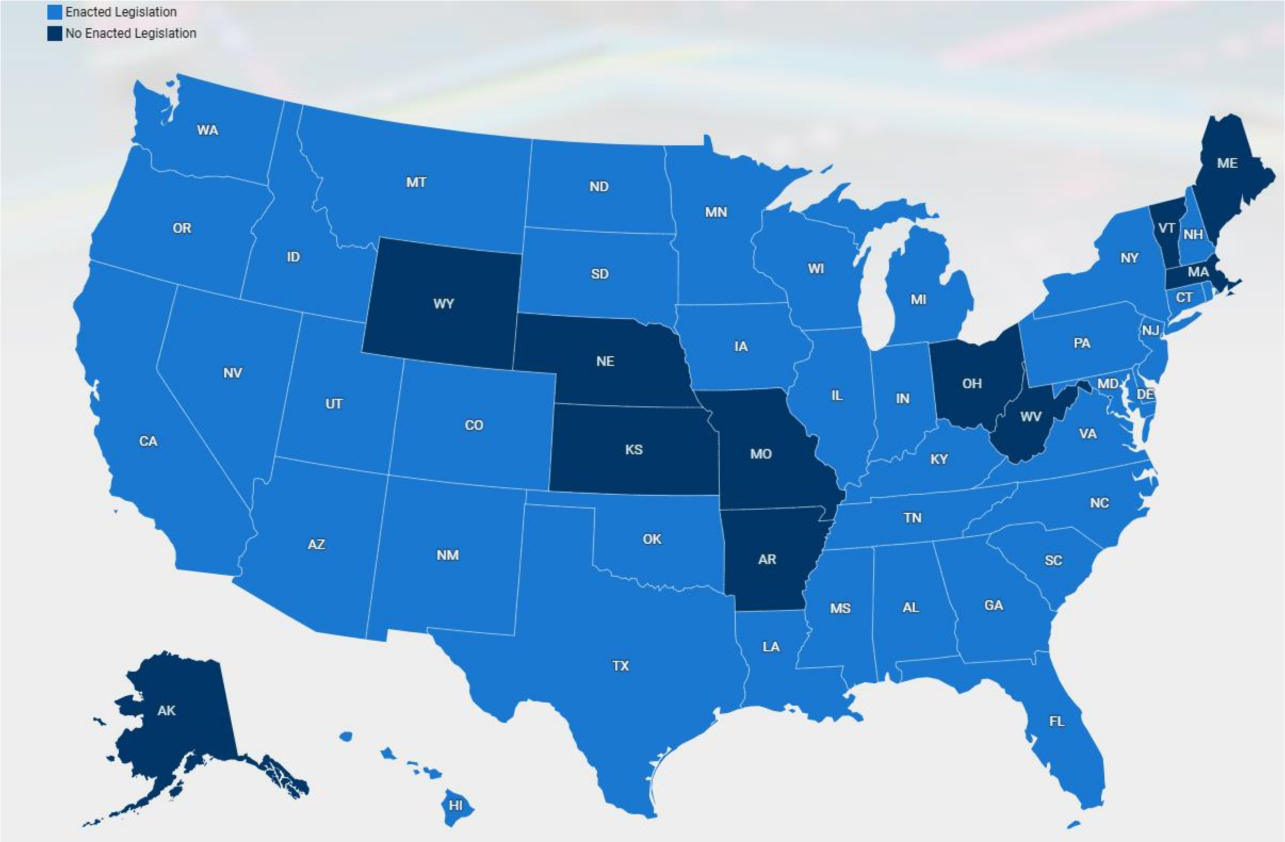

Colorado AI Act (May 17, 2024): regulates high-risk AI systems that impact **consequential decisions, such as access / denial of employment opportunities**

Utah AI Policy Act (May 1, 2024): imposes disclosure requirements for entities providing consumer facing GAI systems.

Virginia AI Act (February 20, 2025): regulates high-risk AI systems that impact **consequential decisions, such as access / denial of employment opportunities**

Illinois Predictive Analytics Act (May 24, 2024): prohibits the **consideration of race or zip codes (as a proxy for race) by an AI system in the context of employment hiring and management decisions**

NYC Local Law 144 of 2021: prohibits the **use of automated employment decision tools without a compliant bias audit.**

Maryland GAI Transparency Bill HB 823 2025 *(pending)*

Texas Responsible AI Governance Act *(pending)*

+ Numerous Existing Non-AI Specific Regulations Affecting AI Usage

- **SEC:** Best Interest Rule, Regulation S-P (Privacy of Consumer Financial Information and Safeguarding of Personal Information)
- **FINRA:** Rule 2111 (Suitability), Rule 2010 (Standards of Commercial Honor and Principles of Trade), Rule 3110 (Supervision)
- **Non-Financial Regs and Laws,** e.g., regarding discrimination
- **Enforcements**, e.g., AI-Washing Enforcements – Marketing Rule Violations

What is keeping policy makers up at night...?

Same things that keep businesses up at night!

Discrimination/Bias Concerns

- Discriminatory outputs that can harm people.

Explainability & Transparency Concerns

- The complexity of advanced models create a “black box” scenario that makes it difficult for deployers and consumers to understand how and why AI technologies work.

Data Privacy & Security Concerns

- Personal data being collected and processed in a manner not permitted.

Intellectual Property Infringement Concerns

- Training and the use of AI can violate third party copyright, trademark, publicity, trade secrets, and patent rights.

"Losing ownership/agency” & “consequential decision making” concerns

- Loss of control (of ownership of data AI is trained on, of decision making, etc.), often quoted next to AI being dynamic, autonomous, etc.

Common Themes

Ambiguous or Inaccurate Definitions & Perceptions of AI: Misunderstandings and overly broad interpretations of AI that can complicate regulation and adoption efforts.

Old “new”: legacy challenges that are not AI specific, become concerns for AI adoption & regulation, specifically.

New “old”: challenges that arise from the idiosyncrasies of GAI & either are not being emphasized enough as concerns for AI adoption & regulation or there is not much to do about them, anyway.

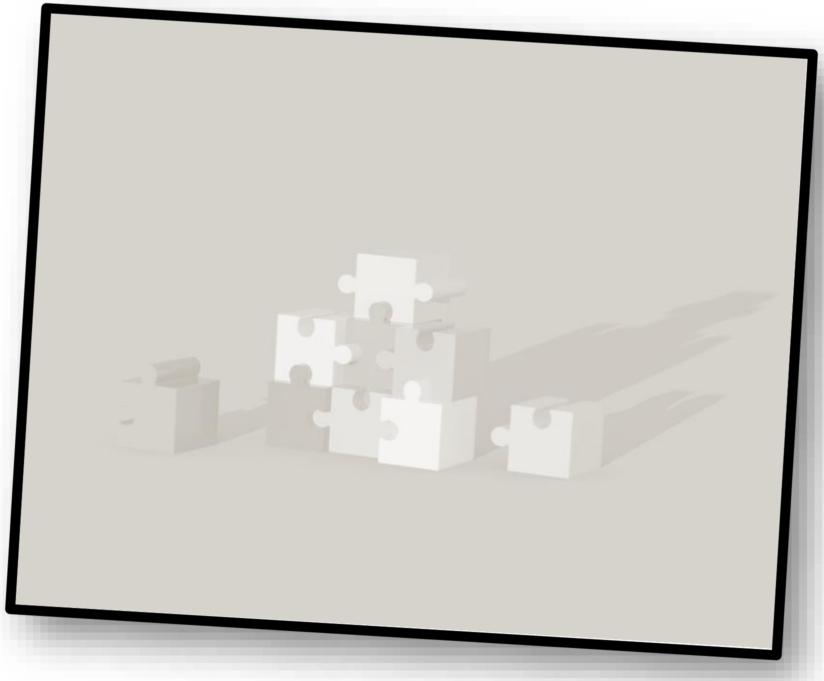

The Old “New” Challenges

Legacy challenges that are not AI specific, become concerns for AI adoption & regulation, specifically.

Old “New” Challenges: “AI Bias” & Discrimination

Biased data lead to biased results that can have societal consequences

Congratulations on your election as mayor! One of your first responsibilities is to determine a fair tax rate for the city. You plan to analyze the average cost of living and conduct a resident survey to gather income information. Your aim is to set a tax rate that supports town development while maintaining affordability for the community.

You begin distributing the survey across different zip codes. The number of survey copies is sufficient to reach all but one of the town's zip codes. Given that most zip codes are addressed, you decide it's acceptable to exclude this final one.

Your conclusion about excluding one zip code will not matter is correct only if the excluded zip code population is as diverse as the population of the rest of the town. If instead, the excluded zip code corresponds to an area where residents are mostly affluent or mostly poor compared to all other neighborhoods, then your estimated average income will be over/under-estimated, respectively.

A simple average can be biased with societal consequences. Nothing new/different in the case of AI…

Old “New” Challenges: Losing ownership & agency

Learning/Training Weights, Dynamic, Adaptive, Autonomous models, etc.

AI/GAI are like any other model. Just (typically) much more complex, especially LLMs.

They do what we program them to do; they are as adaptive, dynamic and autonomous as we program them to be. They have as much agency as we give them.

They are supposed to extract and “own” the information of data sets: that is the raison d'etre of any model, AI or not.

Why do we build models?

Will we lose ownership?

- to describe, represent/recreate our world
- information compression **is the objective**:
  - e.g., a 30-parameter model to describe a data set of 34 numbers is useless
- “weights” or “parameters”, deep learning, simple regression or just calculating and average:
  - All models – AI or not - try to capture patterns in the data and “store” them into a few(er) “weights”

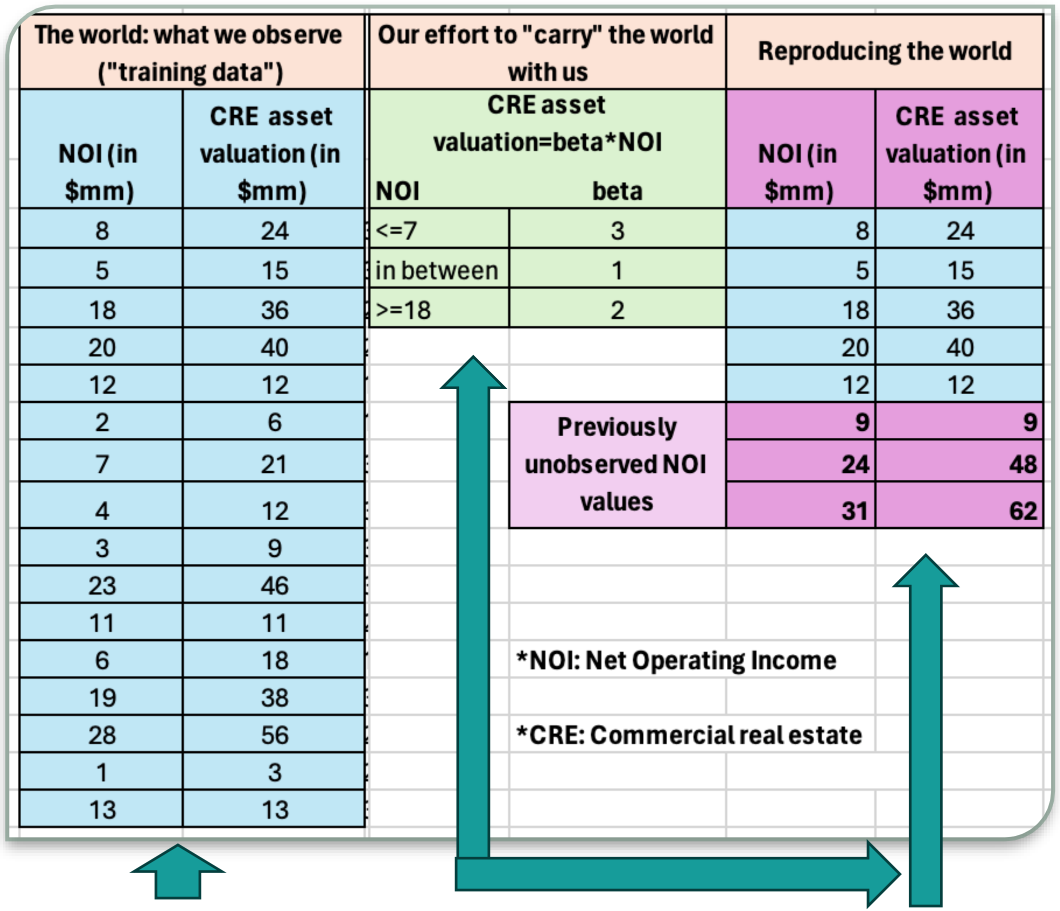

The world: what we observe ("training data")		Our effort to "carry" the world with us		Reproducing the world	
NOI (in $mm)	CRE asset valuation (in $mm)	CRE asset valuation=beta*NOI		NOI (in $mm)	CRE asset valuation (in $mm)
		NOI	beta		
8	24	<=7	3	8	24
5	15	in between	1	5	15
18	36	>=18	2	18	36
20	40			20	40
12	12			12	12
2	6		Previously unobserved NOI values	9	9
7	21			24	48
4	12			31	62
3	9				
23	46				
11	11				
6	18				
19	38				
28	56				
1	3				
13	13				

*NOI: Net Operating Income

*CRE: Commercial real estate

32 numbers

5 numbers

Can provide CRE asset value for *any number* of NOI numbers

Are these models really that intelligent?

Will we lose agency?

- Intelligent (=maximally efficient) models distill patterns in data and store it in as few parameters as possible (**maximal compression**)
- Optimal performance (best loss) at fixed compute is achieved for:
  - Data-to-model size ratio ~>20 (“Chinchilla” paper by DeepMind - Training Compute-Optimal Large Language Models (however, also see Chinchilla Scaling: A replication attempt)
  - Data-to-model size ratio ~100 (order of magnitude – Hu et al https://arxiv.org/pdf/2404.06395)

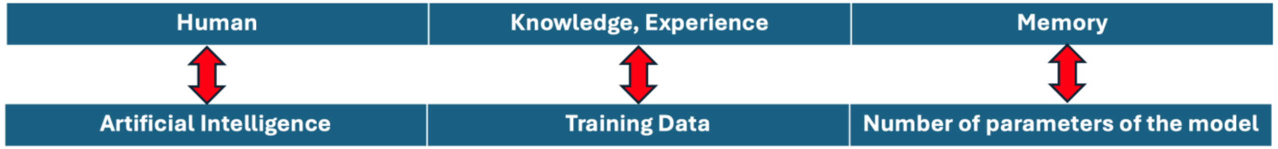

Human

Knowledge, Experience

Memory

Artificial Intelligence

Training Data

Number of parameters of the model

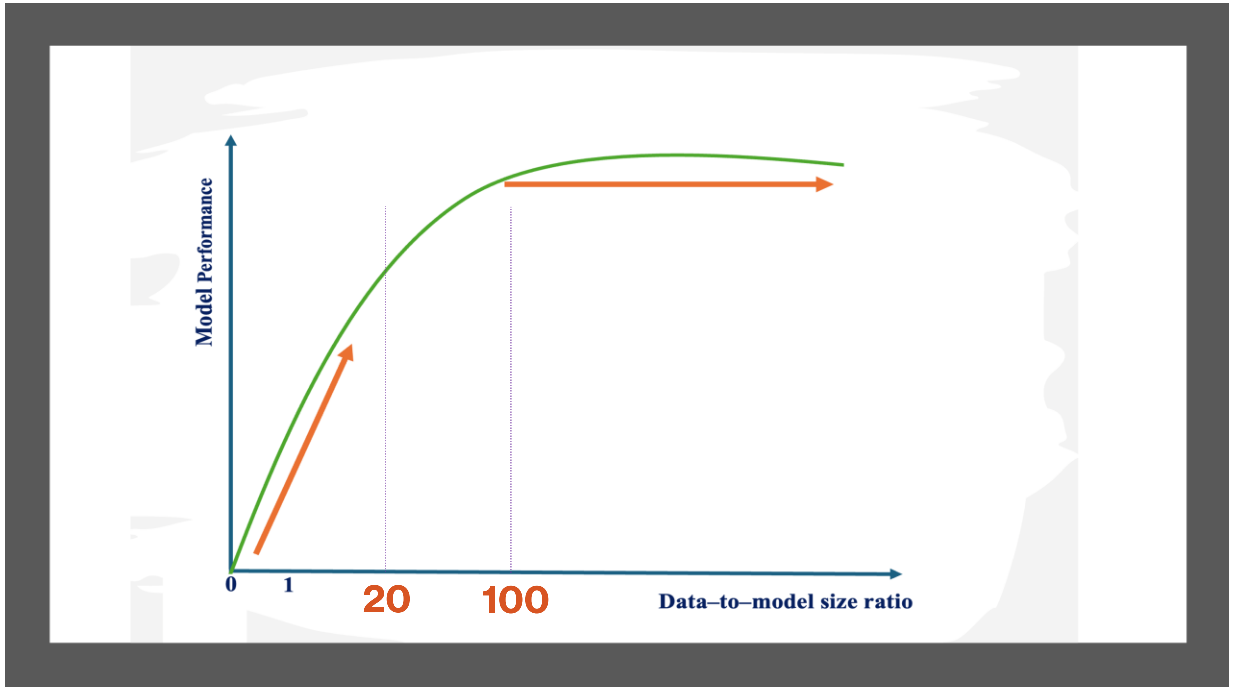

For more see: Generative AI & Large Language Models (LLMs): Will we run out of data? Part I.

Are they even that intelligent?

Will we lose agency?

- **Red: 2023**; **Green: 2024**; **Blue: 2025**
- Optimal performance (best loss) at fixed compute is achieved for:
  - Data-to-model size ratio ~>20 ("Chinchilla" paper by DeepMind - Training Compute-Optimal Large Language Models (however, also see Chinchilla Scaling: A replication attempt)
  - Data-to-model size ratio ~100 (order of magnitude – Hu et al https://arxiv.org/pdf/2404.06395)
- Most current large models (say a few hundred billions of parameters) are enormous databases and highly unlikely to claim ownership or assume agency

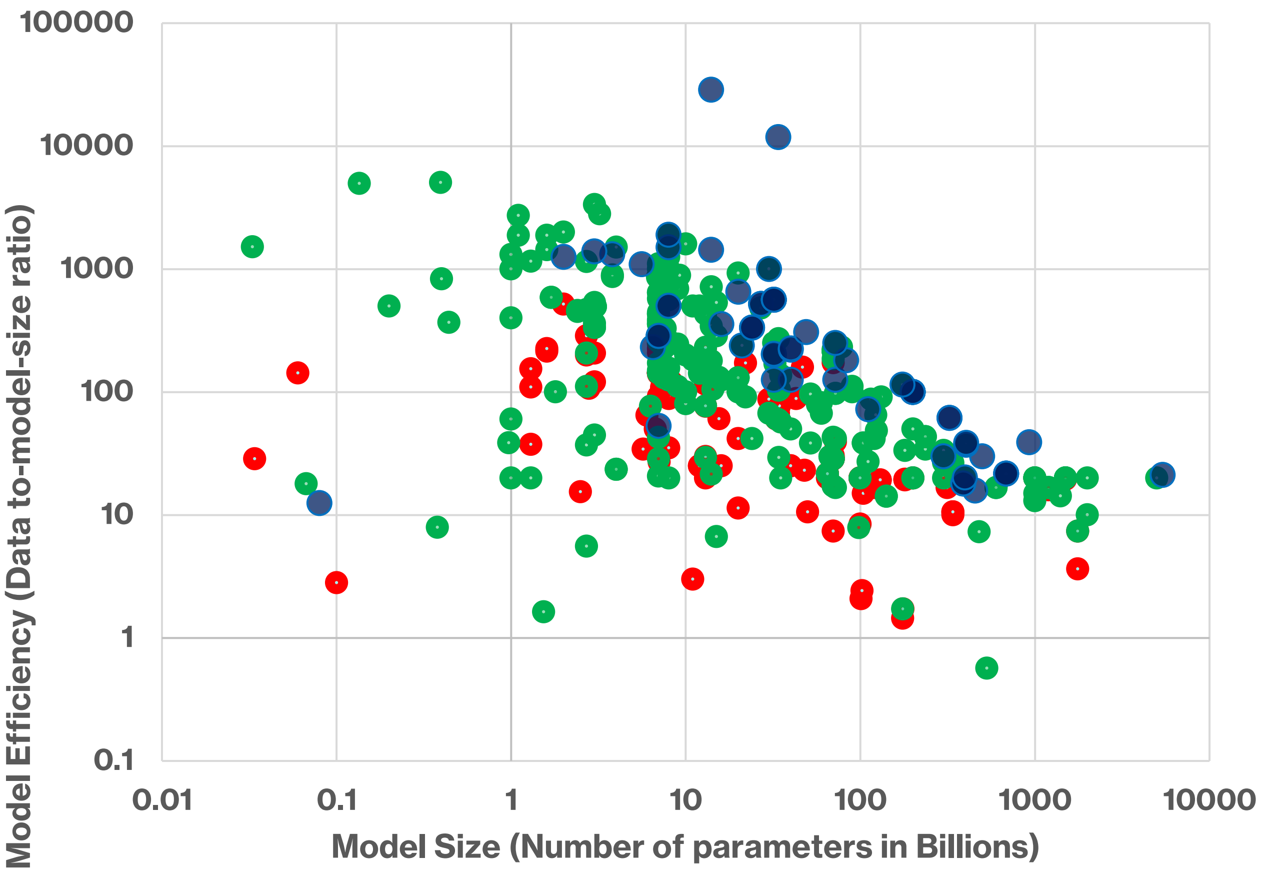

Data source: **LifeArchitect.ai.**

Old “New” Challenges: data privacy and security

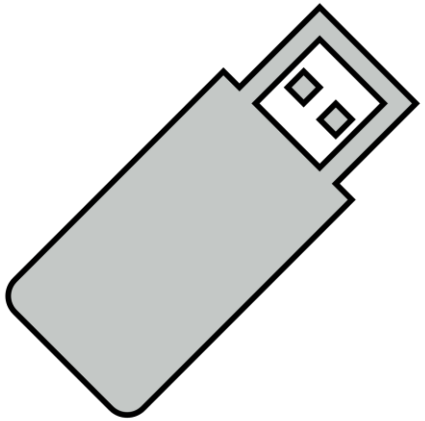

Data Privacy & Security

Most of the training for current Generative AI (GAI) and Large Language Models (LLMs) has been done using datasets selected by others.

Unless we fine-tune an LLM ourselves, we likely do not expose much private data to it.

At least, comparatively, the amount of our own data we expose when using LLMs is minimal compared to the proprietary data used to build other models from scratch – whether AI-based or not.

However, discussions about data privacy and security remain frequent and ongoing within the context of GAI & LLMs...

The Rise of Generative AI & The New “Old” Challenges

Generative AI

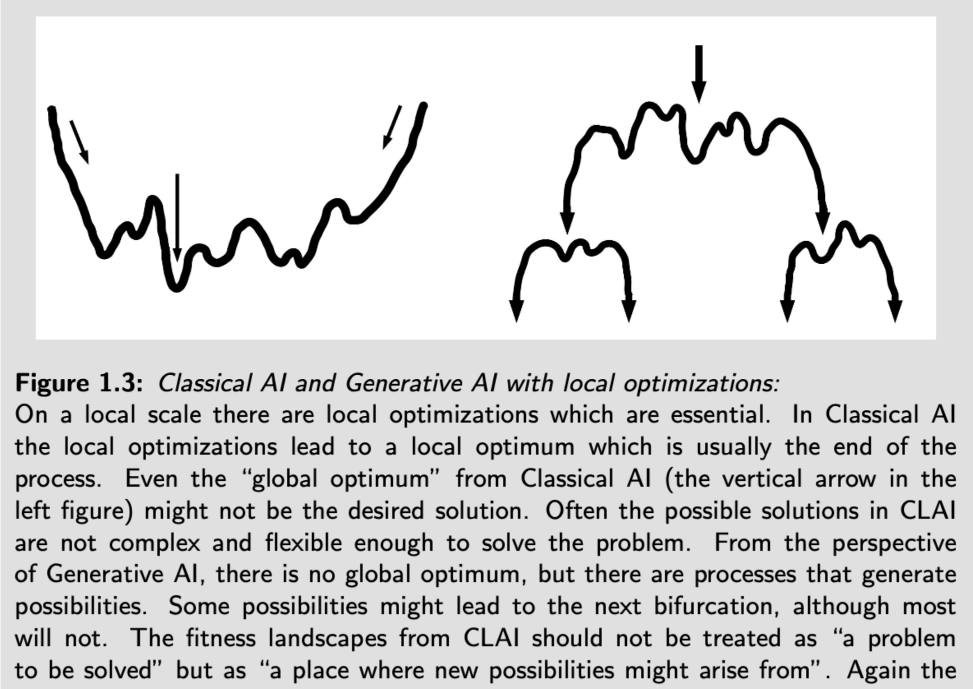

Figure 1.3: *Classical AI and Generative AI with local optimizations:*
On a local scale there are local optimizations which are essential. In Classical AI the local optimizations lead to a local optimum which is usually the end of the process. Even the "global optimum" from Classical AI (the vertical arrow in the left figure) might not be the desired solution. Often the possible solutions in CLAI are not complex and flexible enough to solve the problem. From the perspective of Generative AI, there is no global optimum, but there are processes that generate possibilities. Some possibilities might lead to the next bifurcation, although most will not. The fitness landscapes from CLAI should not be treated as "a problem to be solved" but as "a place where new possibilities might arise from". Again the

Source: "Generative AI: a neo-cybernetics analysis", Tijn van der Zant, University of Groningen, 2010

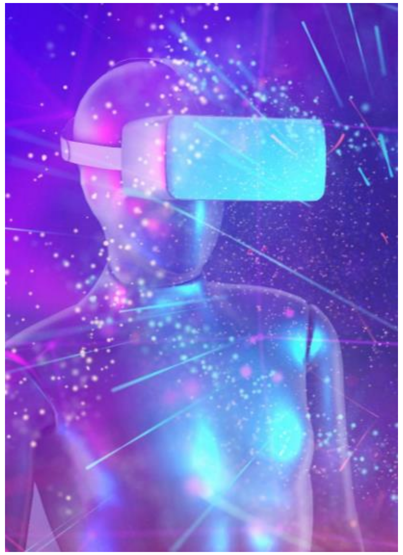

Generative Artificial Intelligence: The machine is provided with **generative mechanisms to actively explore possibilities**

Generative AI is expected to be unpredictable by design

The Amazingness of Generative AI

- Powerful and creative:
  - Identify whether painting is Mona Lisa (discriminative) vs paint Mona Lisa (generative)
  - When calling the automated system of your bank, the system offers predetermined options for your intent (“balance”, “new account”, …, “help me with something else”): discriminative
    - Generative AI allows for non-prespecified new intents
- Various types of inputs, sophisticated outputs
- Interactive via natural language; Easily accessible by all

Credit: Leonardo Da Vinci

New “Old” Challenges:

Challenges that arise from the idiosyncrasies of GAI & either are not being emphasized enough or there is not much to do about...

Explainability & Transparency

How realistic is it to expect these from such complex models?

Explainable to whom?

Intellectual Property Infringement

Most of the training in current GAI/LLMs has occurred by others with their trading data choices.

Stochasticity, hallucinations

These are really new “challenges” difficult to address: because there are what models where built around...

So...

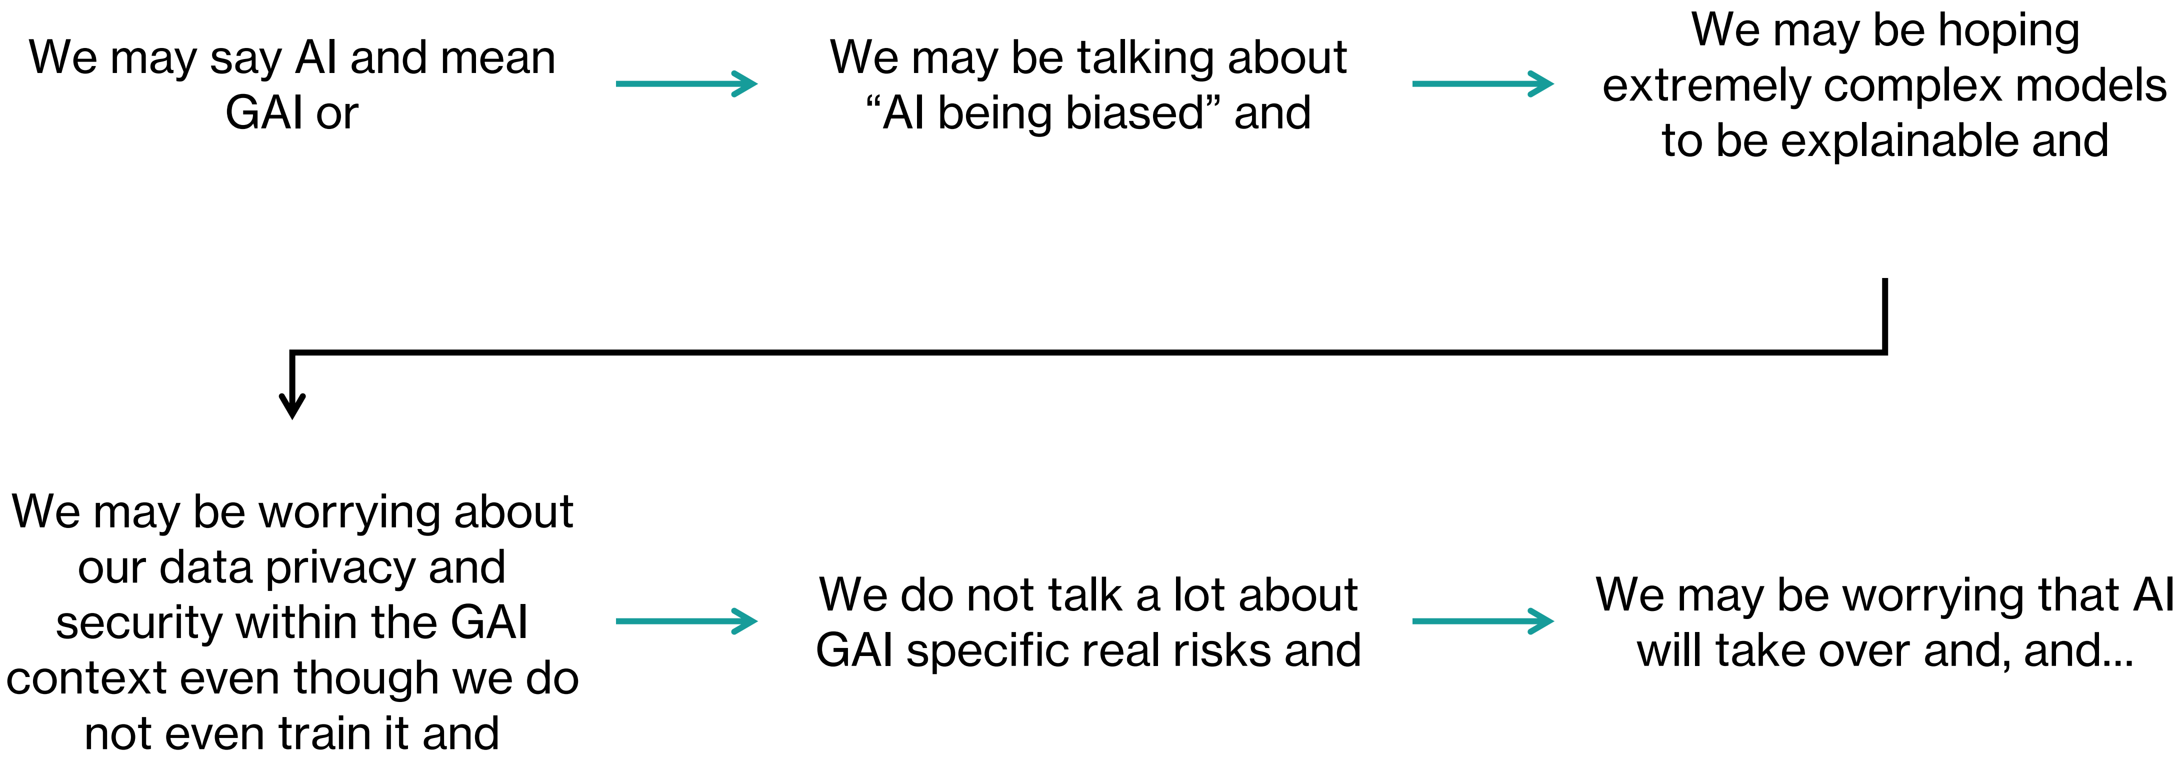

We may say AI and mean GAI or

We may be talking about "AI being biased" and

We may be hoping extremely complex models to be explainable and

We may be worrying about our data privacy and security within the GAI context even though we do not even train it and

We do not talk a lot about GAI specific real risks and

We may be worrying that AI will take over and, and...

So what?

Challenges Faced by Businesses

Understanding Regulations

A comprehensive understanding of regulatory frameworks is essential for businesses to effectively navigate compliance requirements.

Keeping Up with Regulations

Businesses must continuously adapt their operations to comply with ever-changing regulations, which can be a significant hurdle.

Maintaining Competitiveness

Navigating regulatory challenges is crucial for businesses to maintain their competitiveness in the market.

Understanding AI & GAI

What it is and what it is not; what specific risks it carries. What makes sense to be regulated as AI specific and what not.

Instead of A Conclusion

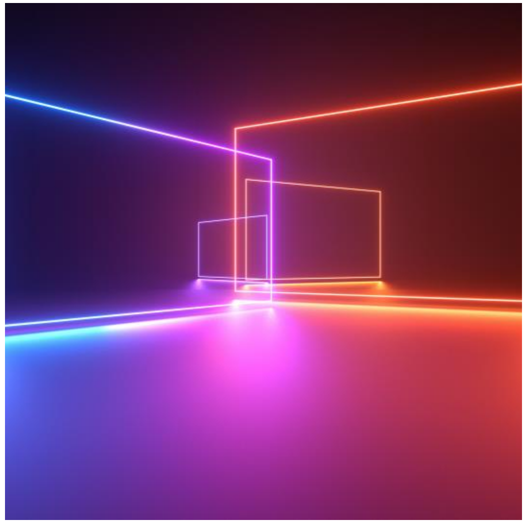

We need to Continue Working with Regulators

We cannot expect to regulate correctly something we do not understand correctly. We need to develop consistent understanding of risks we have been assuming well before AI brought them to the forefront.

Fostering collaboration between innovators and regulators is essential to ensure proper understanding of AI and its risks.

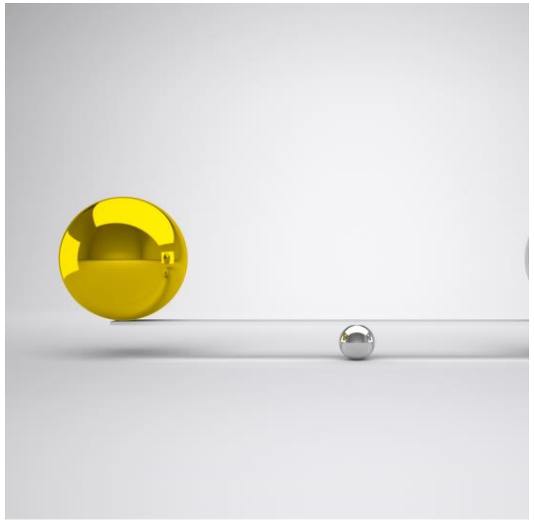

The Potential Impact of Wrong Regulation can be Disastrous

While regulation is necessary for safety, excessive or wrong regulation can stifle innovation, making it difficult for new technologies to emerge and thrive in the market.

Striking a balance between safety and innovation is essential to allow AI to grow while maintaining public trust and safety.