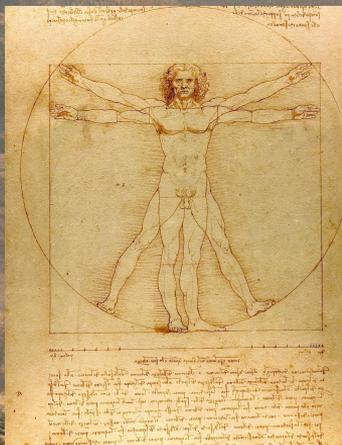
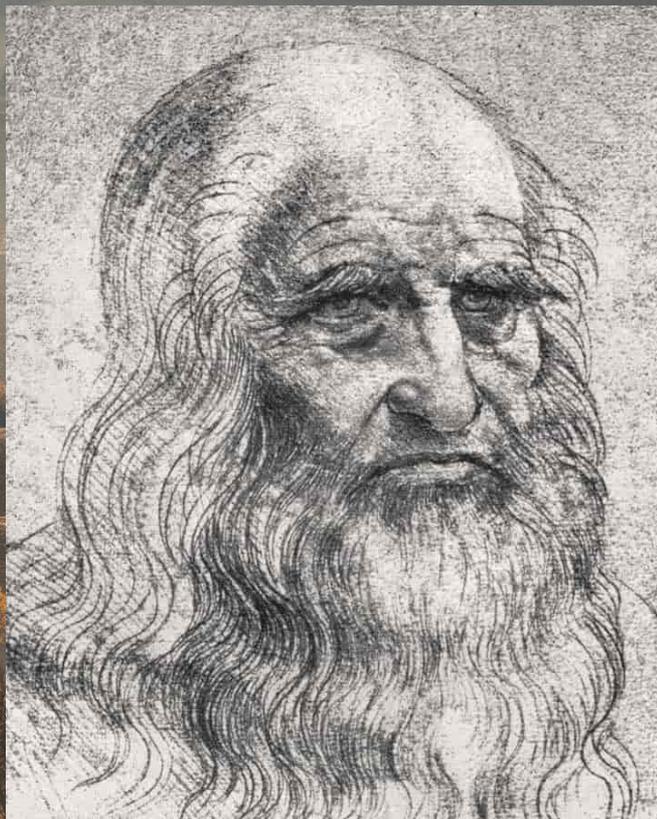
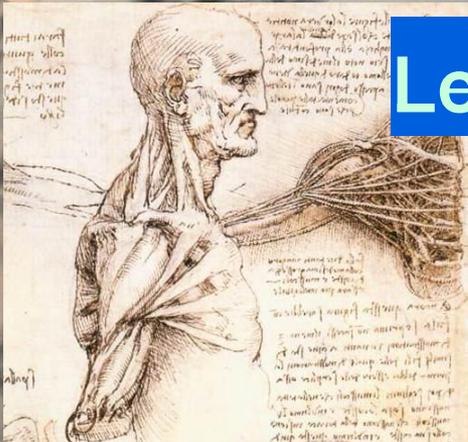


A wide-angle landscape photograph of rolling hills in Florence, Italy, captured in a hazy, atmospheric setting. The hills are covered in a mix of green and golden-brown vegetation, suggesting a late autumn or early spring scene. In the middle ground, two tall, slender cypress trees stand prominently on a hillside. The background is filled with more distant, misty hills, and a small cluster of buildings is visible on the horizon. The overall mood is serene and historical. The text "Florence, 1505" is overlaid in the center in a clean, white, sans-serif font.

Florence, 1505

Leonardo da Vinci (1452 - 1519)





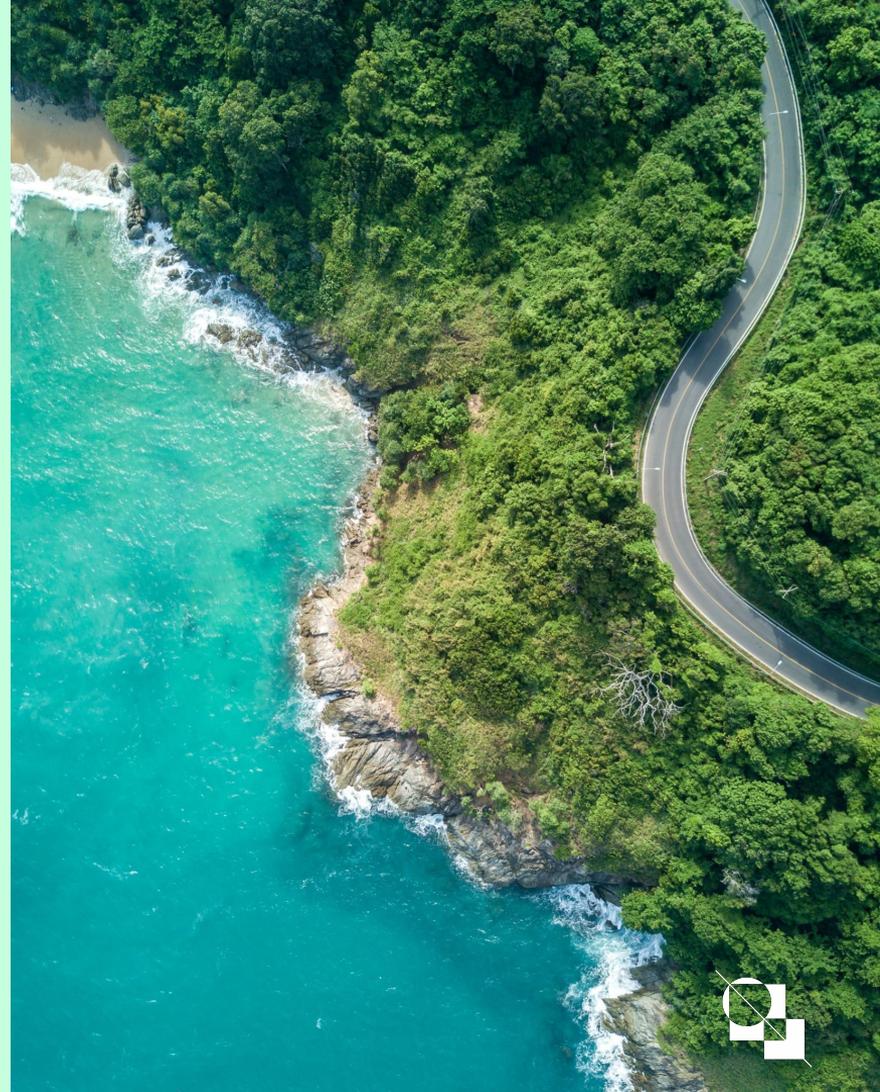
Mastering AI: Business in the generative age

March 2024

Contents

This presentation will cover

- > Understanding generative AI
- > AI's economic impact
- > Framing the art of the possible
- > Open Questions and Issues



About me: Walter Pasquarelli

- Expert advisor in Generative A.I. Strategy
- Led AI editorial and advisory programmes at The Economist. Work together with large tech firms and global central governments

Services:

- **Research and advisory** on AI impact, policy implications and strategy
- **Workshops and keynotes:** Moderator for events (World Economic Forum, IBM, The Economist)

LinkedIn



Understanding AI



Which one is
real?



Which one is
real?



How would you define artificial intelligence?



Mapping gen AI

Foundation Models

**Vision, Language, Music,
Audio**

Generative AI

**Branch of AI able to
generate media following
prompts**

Machine Learning

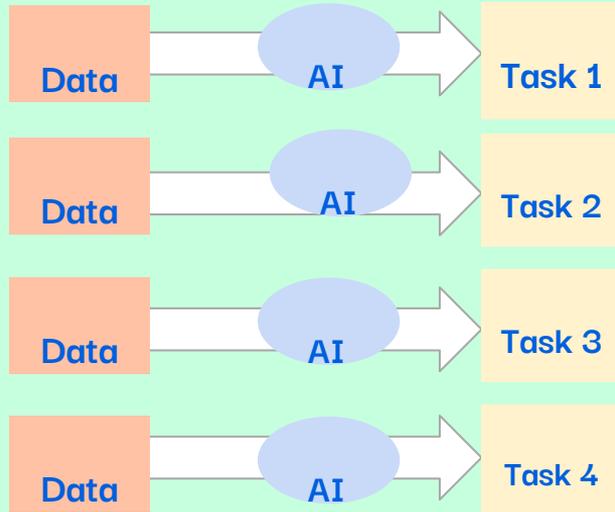
Machines able to improve
performance through
experience

Artificial Intelligence

Ability of artefacts to solve
problems and accomplish
tasks



Traditional AI research (Deep Learning)

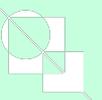


1. Supervised Learning

Example: Client demand forecast

2. Unsupervised Learning

Example: Customer segmentation



What happened?



SCALE

(Data, computing, models)

New Models

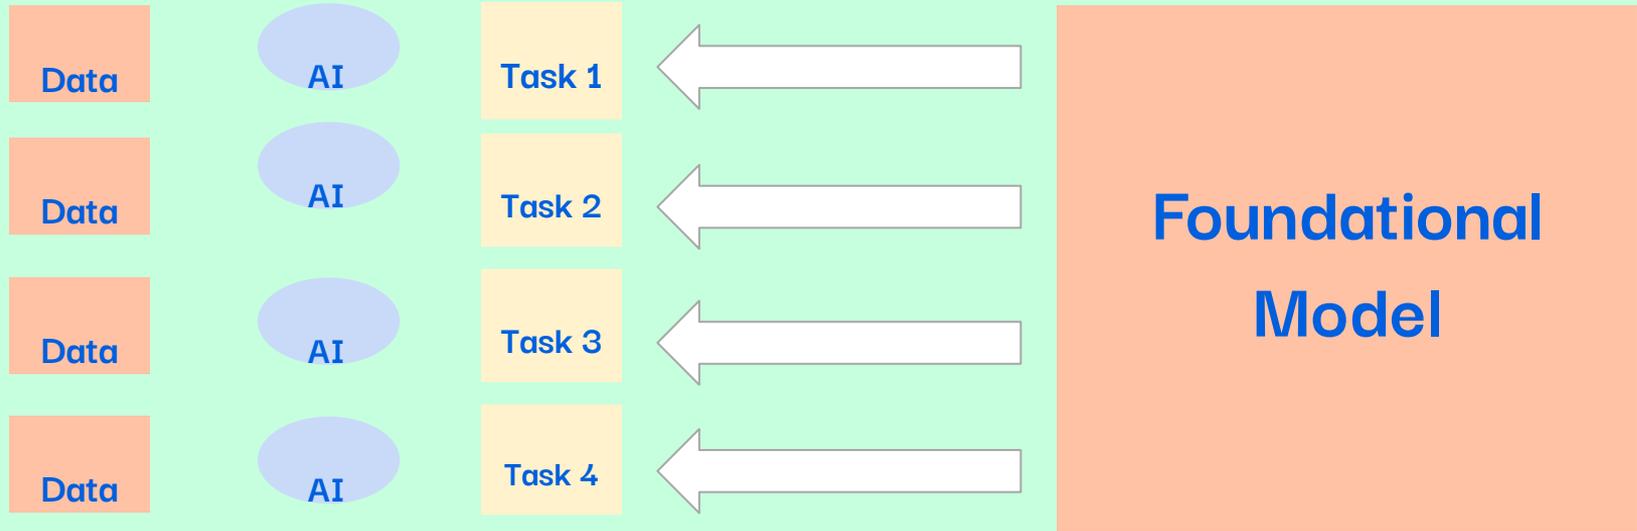
(Transformers)

Pre-trained

(In-context
learning)



Foundation Model Process



Towards a new form of legal tech startups

Gen AI
startup

Gen AI
startup

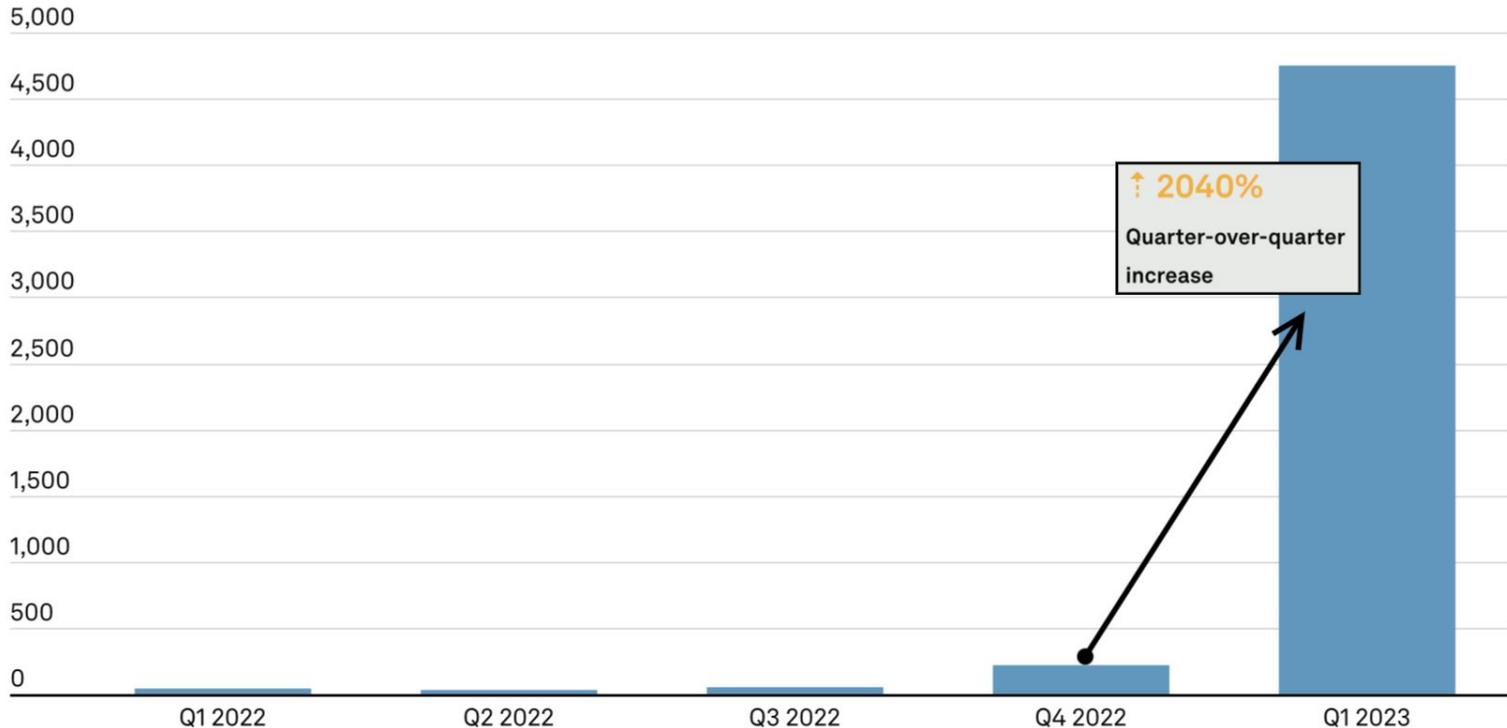
Gen AI
startup

Foundation Model

(LLama, GPT, etc.)



Quarterly Volumes of Generative AI Reports



Source: S&P Global Market Intelligence. Data referenced in this graphic aggregates Investment Research reports from more than 1,000 brokers and research providers globally that reference the trend in generative artificial intelligence or related permutations of searched keywords published between January 1, 2022 and March 31, 2023.



Text-to-text



tell me a story



Once upon a time, in a far-off kingdom, there was a brave knight named Sir Roland. He had served the kingdom for many years and had won numerous battles against enemies of the kingdom. Sir Roland was known for his courage, wisdom, and his unwavering loyalty to the king.

One day, the king summoned Sir Roland to his court and asked him to undertake a dangerous mission. The mission was to rescue the princess, who had been kidnapped by a fearsome dragon. The dragon was known to be invincible, and no one had ever returned alive after confronting it.



Text-to-image



Text-to-video



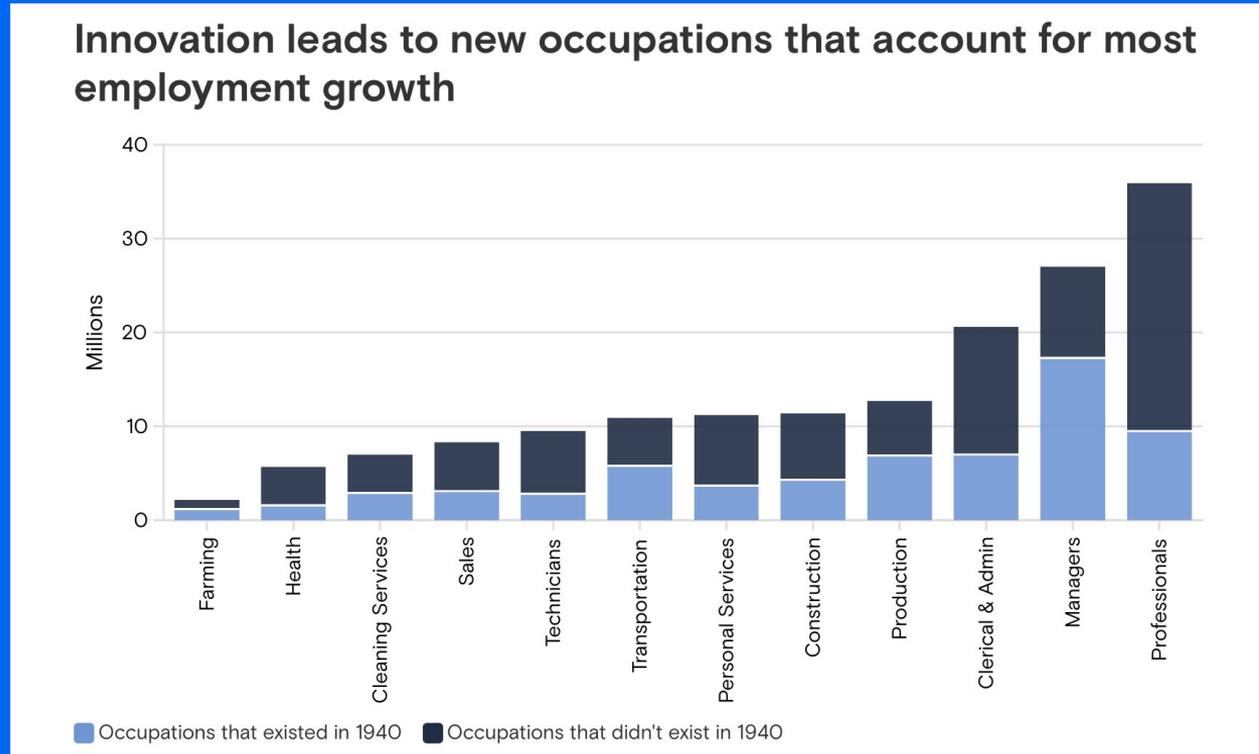


Future of AI: towards AI agents

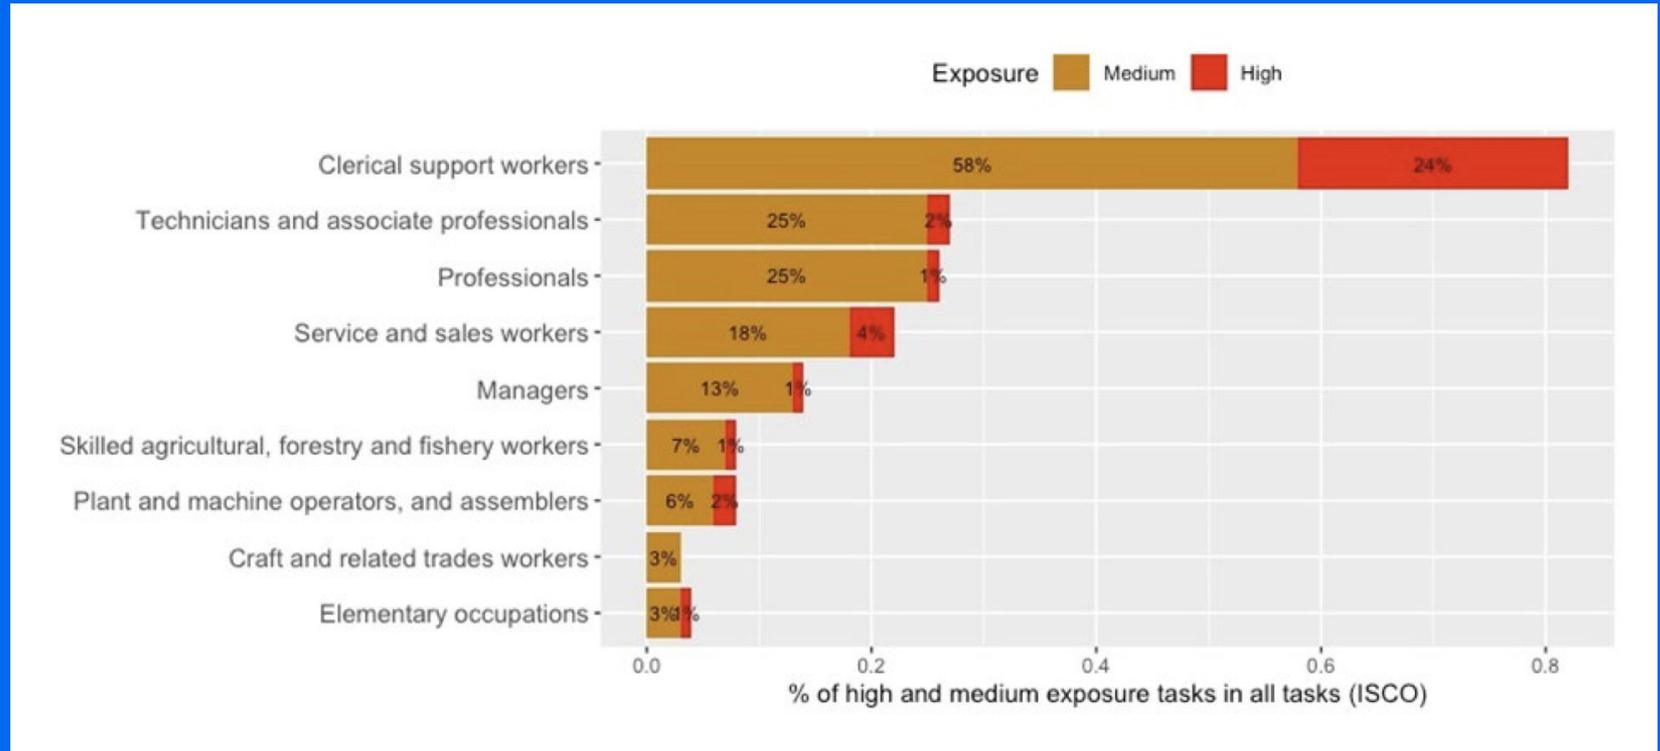
- **AI Assistants**
 - Tools that can fulfil a number of different tasks and navigate the web to fulfil those
- **Future of gen AI?**
 - Eminent voices in tech including Bill Gates and Tim Berners Lee, argue that these will be integrated into most tech devices

Economic impact of AI

Augmentation not automation



Tasks exposed to automation are overestimated



Increasing worker productivity and job satisfaction

- In a study by Noy et al (2023) at MIT, ChatGPT improved writing task productivity, reducing time taken by 0.8 SDs and enhancing output quality by 0.4 SDs.
- It also reduced inequality and shifted tasks towards idea-generation and editing.
- ChatGPT increased job satisfaction, self-efficacy, and generated both concerns and excitement about automation.
- Similarly, in customer support, a generative AI-based conversational assistant increased productivity by 14% on average, particularly benefiting novice and low-skilled workers (Brynjolfsson et al, 2023).

Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence

Shakked Noy Whitney Zhang
MIT MIT

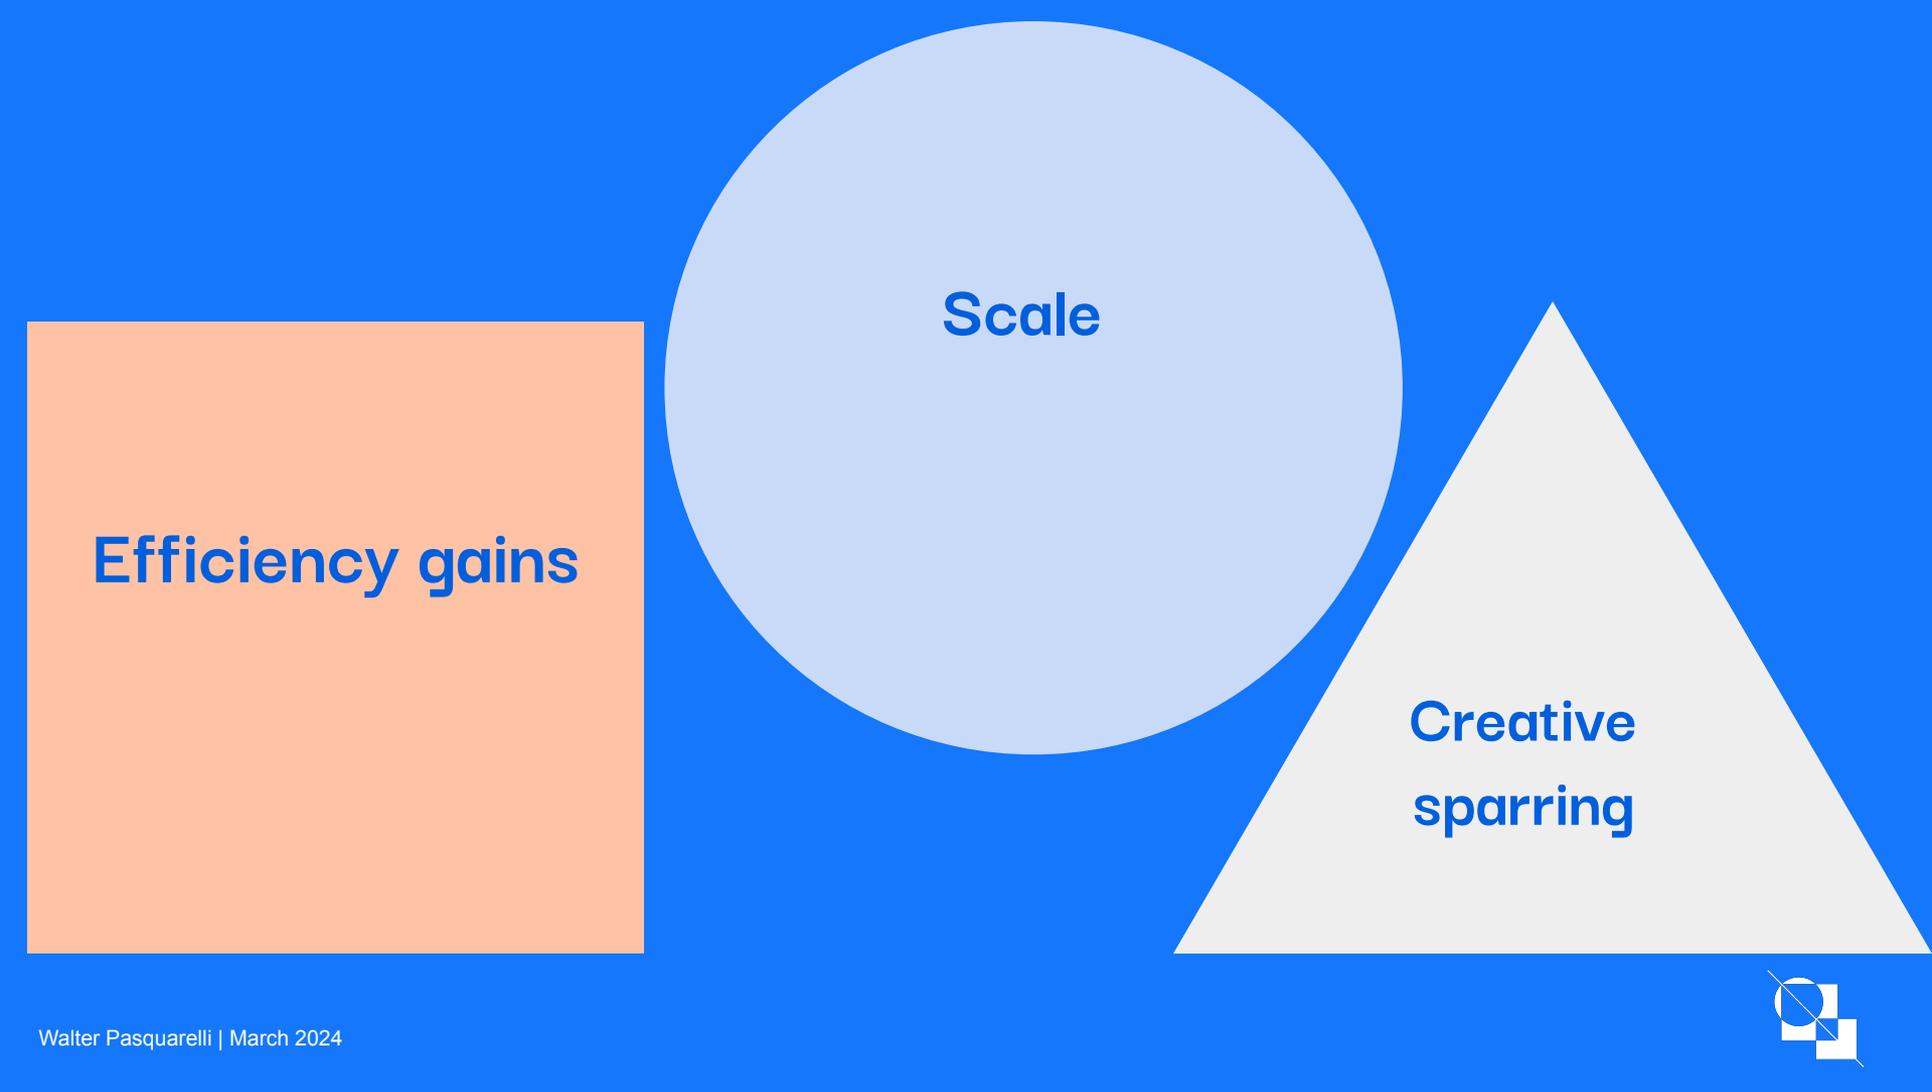
March 10, 2023
Working Paper (not peer reviewed)

Abstract

We examine the productivity effects of a generative artificial intelligence technology—the assistive chatbot ChatGPT—in the context of mid-level professional writing tasks. In a preregistered online experiment, we assign occupation-specific, incentivized writing tasks to 444 college-educated professionals, and randomly expose half of them to ChatGPT. Our results show that ChatGPT substantially raises average productivity: time taken decreases by 0.8 SDs and output quality rises by 0.4 SDs. Inequality between workers decreases, as ChatGPT compresses the productivity distribution by benefiting low-ability workers more. ChatGPT mostly substitutes for worker effort rather than complementing worker skills, and restructures tasks towards idea-generation and editing and away from rough-drafting. Exposure to ChatGPT increases job satisfaction and self-efficacy and heightens both concern and excitement about automation technologies.

We gratefully acknowledge financial support from an Emergent Ventures grant, the George and Obie Shultz Fund, and the National Science Foundation Graduate Research Fellowship under Grant No. 1745302. The research described in this article was approved by the MIT Committee on the Use of Humans as Experimental Subjects, and was preregistered at the AEA RCT Registry (AEA RCT-0010882). We thank Daron Acemoglu, Nikhil Agarwal, David Autor, Lucas Barros, Talia Benheim, Amy Finkelstein, John Horton, Simon Jäger, Ailidh Leslie, Jackson Mejia, Ilan Noy, Liora Noy, Emily Partridge, Charlie Rafkin, Aakaash Rao, Nina Roussille, Chris Roth, Frank Schillbach, Benjamin Schoefer, Lexi Schubert, Advik Shreekumar, Shine Wu, and participants at the MIT Labor Lunch for helpful comments and conversations.





Scale

Efficiency gains

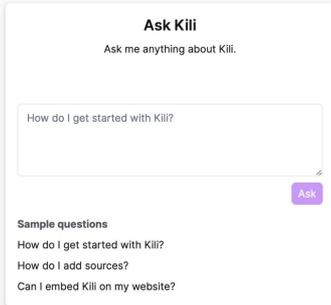
**Creative
sparring**



SCALE: Kili assists users in better understanding legal services

Create an assistant in seconds

Kili uses information related to your product or service to answer questions. Integrate your assistant within your product or implement it internally to assist your team.



Ask Kili
Ask me anything about Kili.

How do I get started with Kili?

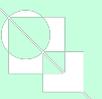
Ask

Sample questions
How do I get started with Kili?
How do I add sources?
Can I embed Kili on my website?

Kili provides a gen AI—powered customers sales assistant that uses information related to products and services to answer user questions. It is different from ChatGPT in that Kili uses private and public knowledge to ensure responses are accurate and specific—whereas chatGPT uses mainly public information. Kili helps customers:

- Engage visitors in real-time, answer their questions, and provide a personalised experience.
- Satisfaction rose to 78% compared to 54% with a human operator
- Equips internal teams to resolve customer questions faster and more efficiently. (Brynjolfsson et al, 2023)
- Using conversational AI to analyze retailer data and consumer sentiment across brands and channels

Data required: Data about products and services is imported and used to tune the foundational model used for Kili.

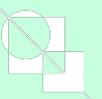


CREATIVE SPARRING: Twilio Optimises knowledge management with gen AI tools



- Twilio relies on ServiceNow for its ticketing system, handling a substantial volume of weekly tickets.
- Their knowledge base serves as an extensive resource hub, offering troubleshooting guides, product documentation, and support articles.
- With Aisera's AI Service Desk, Twilio achieved an impressive 65% auto-resolution rate for incoming support requests.

→ This same technology is used by banks such as Morgan Stanley, for getting together investment reports and analysis



Strategic implications

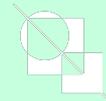
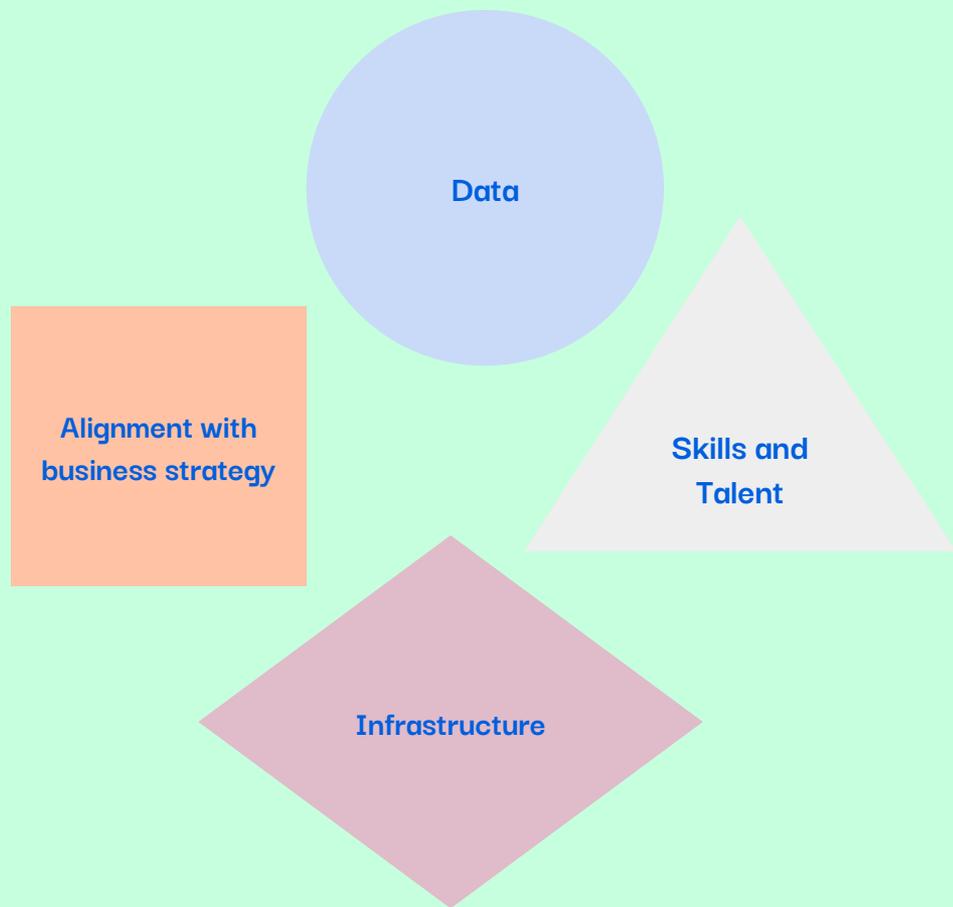
- **Shift in Job Roles:**
 - Decreased need for roles focused on manual, routine work and an
 - Increased demand for roles that leverage human judgment and strategic thinking.
- **Demand for New Skills:**
 - Multi-disciplinary approach where expertise needs to be complemented with a strong grasp of technology and its applications in legal contexts.
- **Impact on Education and Training (LET):**
 - Education institutions must adapt to prepare students for a future where AI plays a central role in their sector.
- **Organizational and Business Model Changes:**
 - Changes in the power structure, business models, and recruitment strategy



Framing the art of the possible



Elements of an AI ready firm



LLMs don't store facts.
They store **probabilities**



Prompt Engineering

“The task of developing prompts that guide an LLM to perform specialised tasks”

AKA SIMPLE QUESTIONS AND TASKS TO CHATBOTS

Advantages of prompt engineering

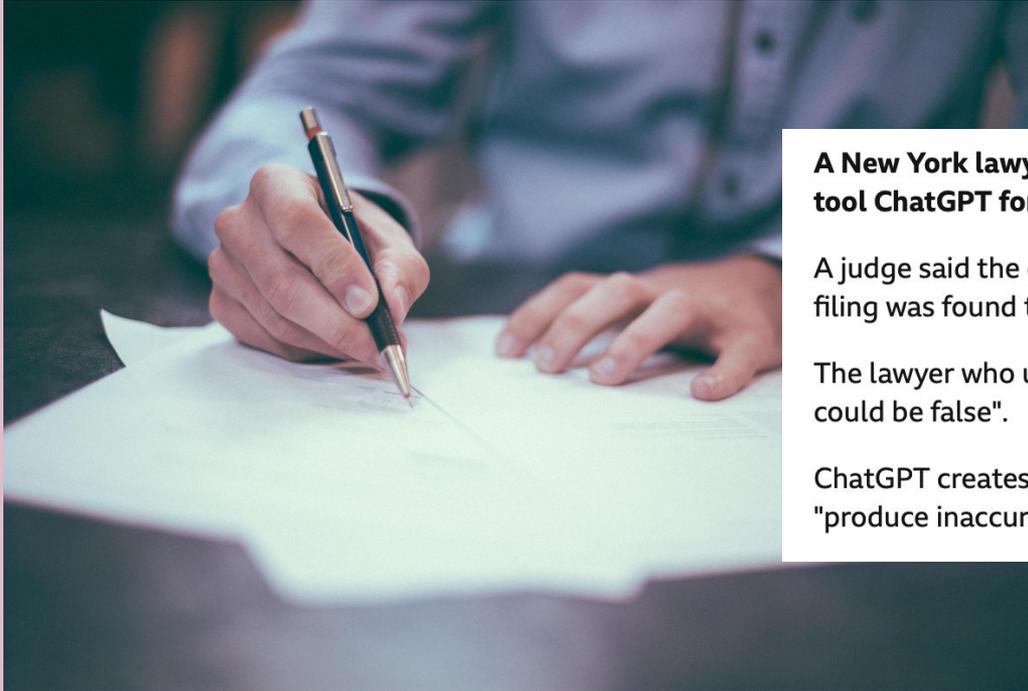
- **Quick and easy**
- **A no-code option for non-technical people**

Disadvantages of prompting include:

- **Risk of hallucination**
- **Non-specialised**
- **Typically “average” responses with edge cases harder to spot**



Fabricating truth: “*LawyerGPT*”



A New York lawyer is facing a court hearing of his own after his firm used AI tool ChatGPT for legal research.

A judge said the court was faced with an "unprecedented circumstance" after a filing was found to reference example legal cases that did not exist.

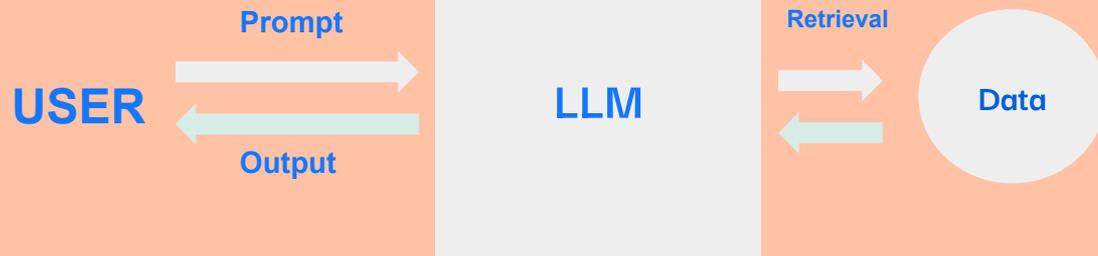
The lawyer who used the tool told the court he was "unaware that its content could be false".

ChatGPT creates original text on request, but comes with warnings it can "produce inaccurate information".

Source: BBC

Embedding

(Retrieval
augmented
generation)



Integrating your dataset within a foundation model to make decisions based on that.

RAG is a great tool for:

- **Question and answer agents**
- **Retrieving relevant information**
- **Reducing the chance of hallucination**

Disadvantages of RAG include:

- **Dataset needs to be clean and complete or otherwise answers may be incomplete.**



Case Study: Morgan Stanley Explores GPT-4 for Financial Advisor Support



- While not yet in full production, Morgan Stanley is advancing its experimentation with GPT-4 technology.
- They've identified over 100,000 documents relevant to their 16,000 financial advisors, addressing investment recommendations, business queries, and procedural questions.
- GPT-4 has been fine-tuned using generative AI on this vast corpus to provide answers to financial advisors' inquiries.





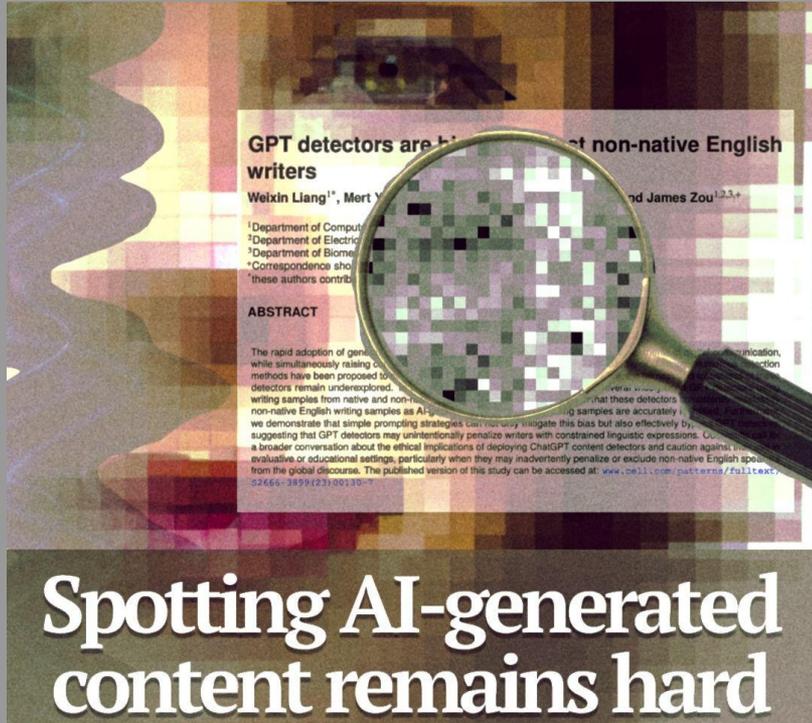
A word of caution

How can we distinguish what is
real vs what is *not*?

Scam emails and phishing



Remedies for authenticating synthetic content



Detection tools:

Text : 26% accuracy (Source: OpenAI)

Image and video: highly unreliable

Provenance:

C2PA standard

Embedded cryptographic hashes

Education Campaigns

Data governance

Settings

✕

- ⚙️ General
- 📄 Data controls**

Chat history & training

Save new chats on this browser to your history and allow them to be used to improve our models. Unsaved chats will be deleted from our systems within 30 days. This setting does not sync across browsers or devices. [Learn more](#)

Shared links Manage

Export data Export

Delete account Delete

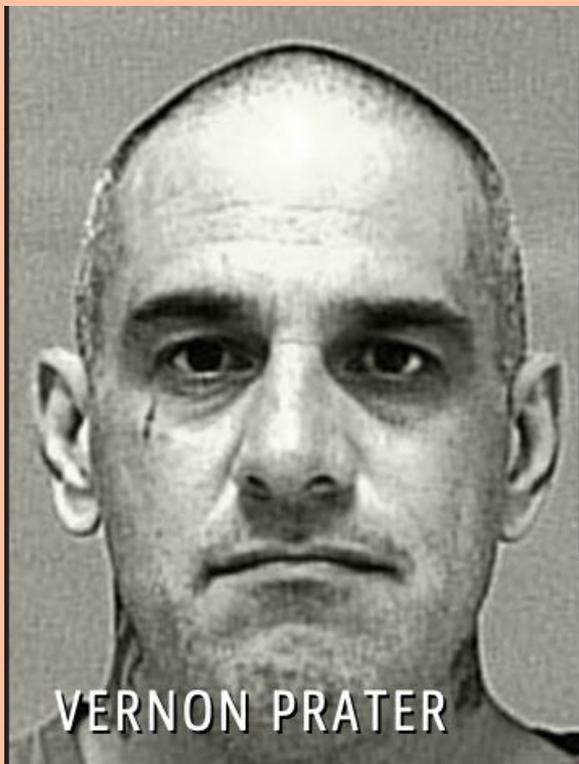
Machine Bias

There's software used across the country to predict future criminals. And it's biased against blacks.

by Julia Angwin, Jeff Larson, Surya Mattu and Lauren Kirchner, ProPublica

May 23, 2016

- Criminologists have long aimed to predict recidivism rates
- Northpointe developed a tool to predict recidivism rates of offenders based on a set of data points
- ProPublica investigation found that it was biased against blacks



VERNON PRATER

LOW RISK

3



BRISHA BORDEN

HIGH RISK

8

VERNON PRATER

Prior Offenses

2 armed robberies, 1
attempted armed
robbery

Subsequent Offenses

1 grand theft

LOW RISK

3

BRISHA BORDEN

Prior Offenses

4 juvenile
misdemeanors

Subsequent Offenses

None

HIGH RISK

8

	WHITE	AFRICAN AMERICAN
Labeled Higher Risk, But Didn't Re-Offend	23.5%	44.9%
Labeled Lower Risk, Yet Did Re-Offend	47.7%	28.0%

Remedies for trustworthy decisions



Auditing and transparency

Recourse to challenging algorithmic decisions

Assessing data used to train systems

Auditing and transparency

Recourse to challenging decisions

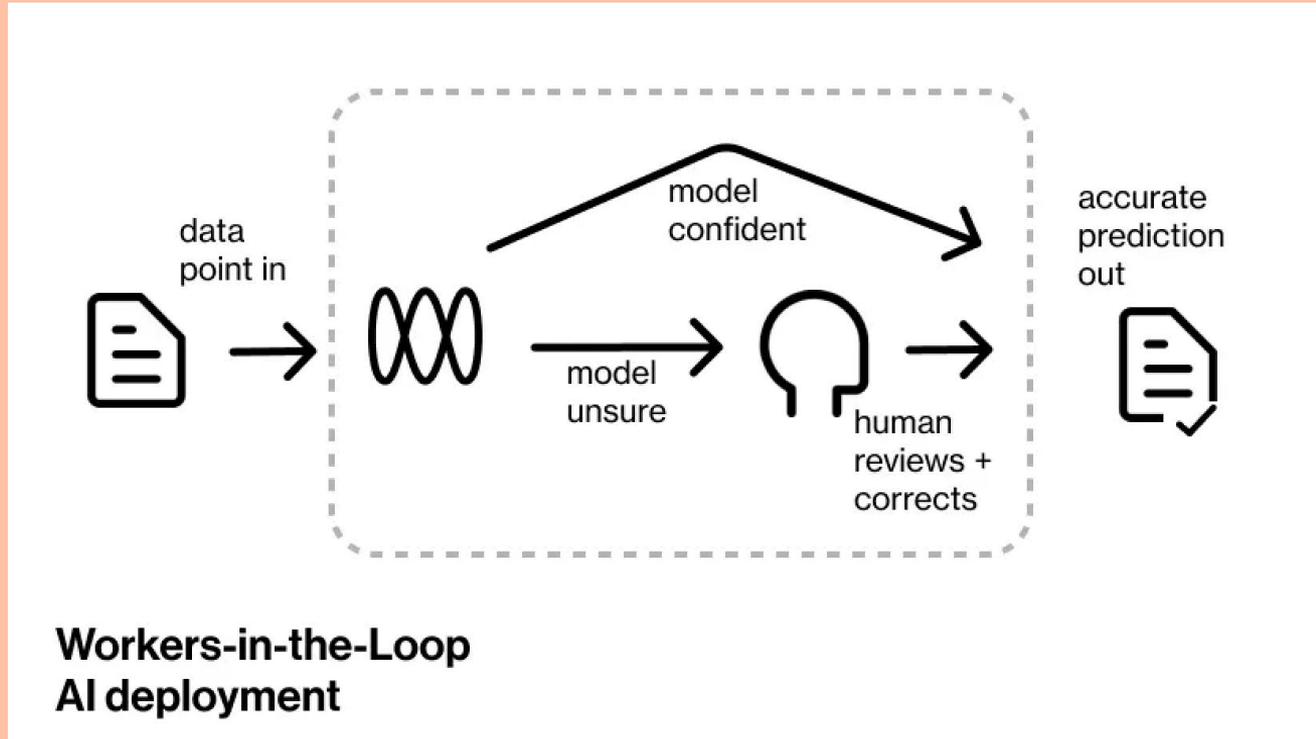
Continuous monitoring

Control

Human in the loop

Avoid high stakes decisions

Human in the loop can improve precision of outcomes



- 1. Demystify generative AI and emerging techn**
- 2. Identify the use cases for your firm**
- 3. Review your capabilities and identify constraints**
- 4. Start a small pilot, before scaling**
- 5. Weave in ethics and trust from the trust**



AI is a journey



Questions?

