



AI AGENTS

Preparing Financial Services for the Next Wave of Innovation

With Brent Dixon, Dixon Strategic Labs

FSI | Retail Banking Conference | May 8th

One of the first financial institutions to really embrace generative AI

Morgan Stanley

Morgan Stanley



OpenAI



Their question:

**How do we really get a better grip
of what it means to operate at
scale in an AI-powered world?**



~300 of their
financial advisors



~300 of their
financial advisors

+



Technology from
OpenAI



~300 of their
financial advisors

+



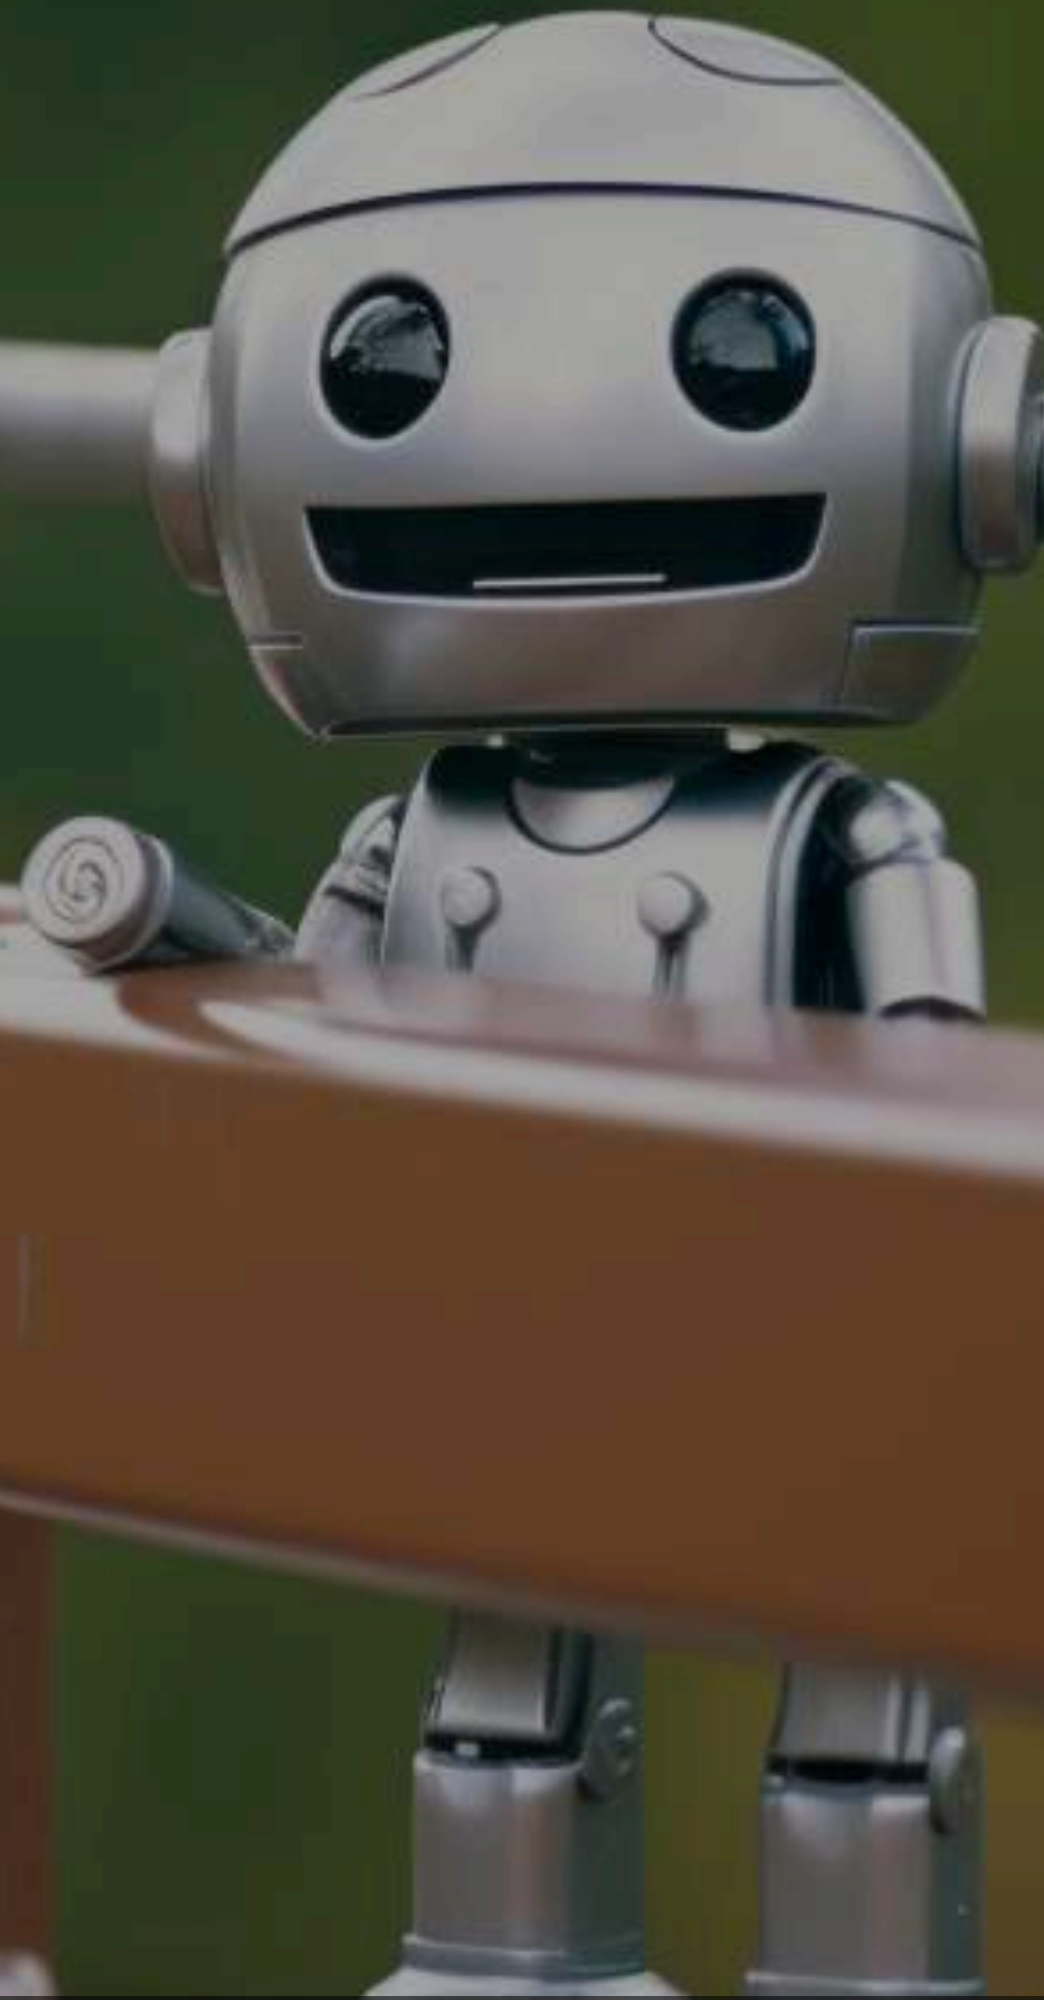
Technology from
OpenAI

+



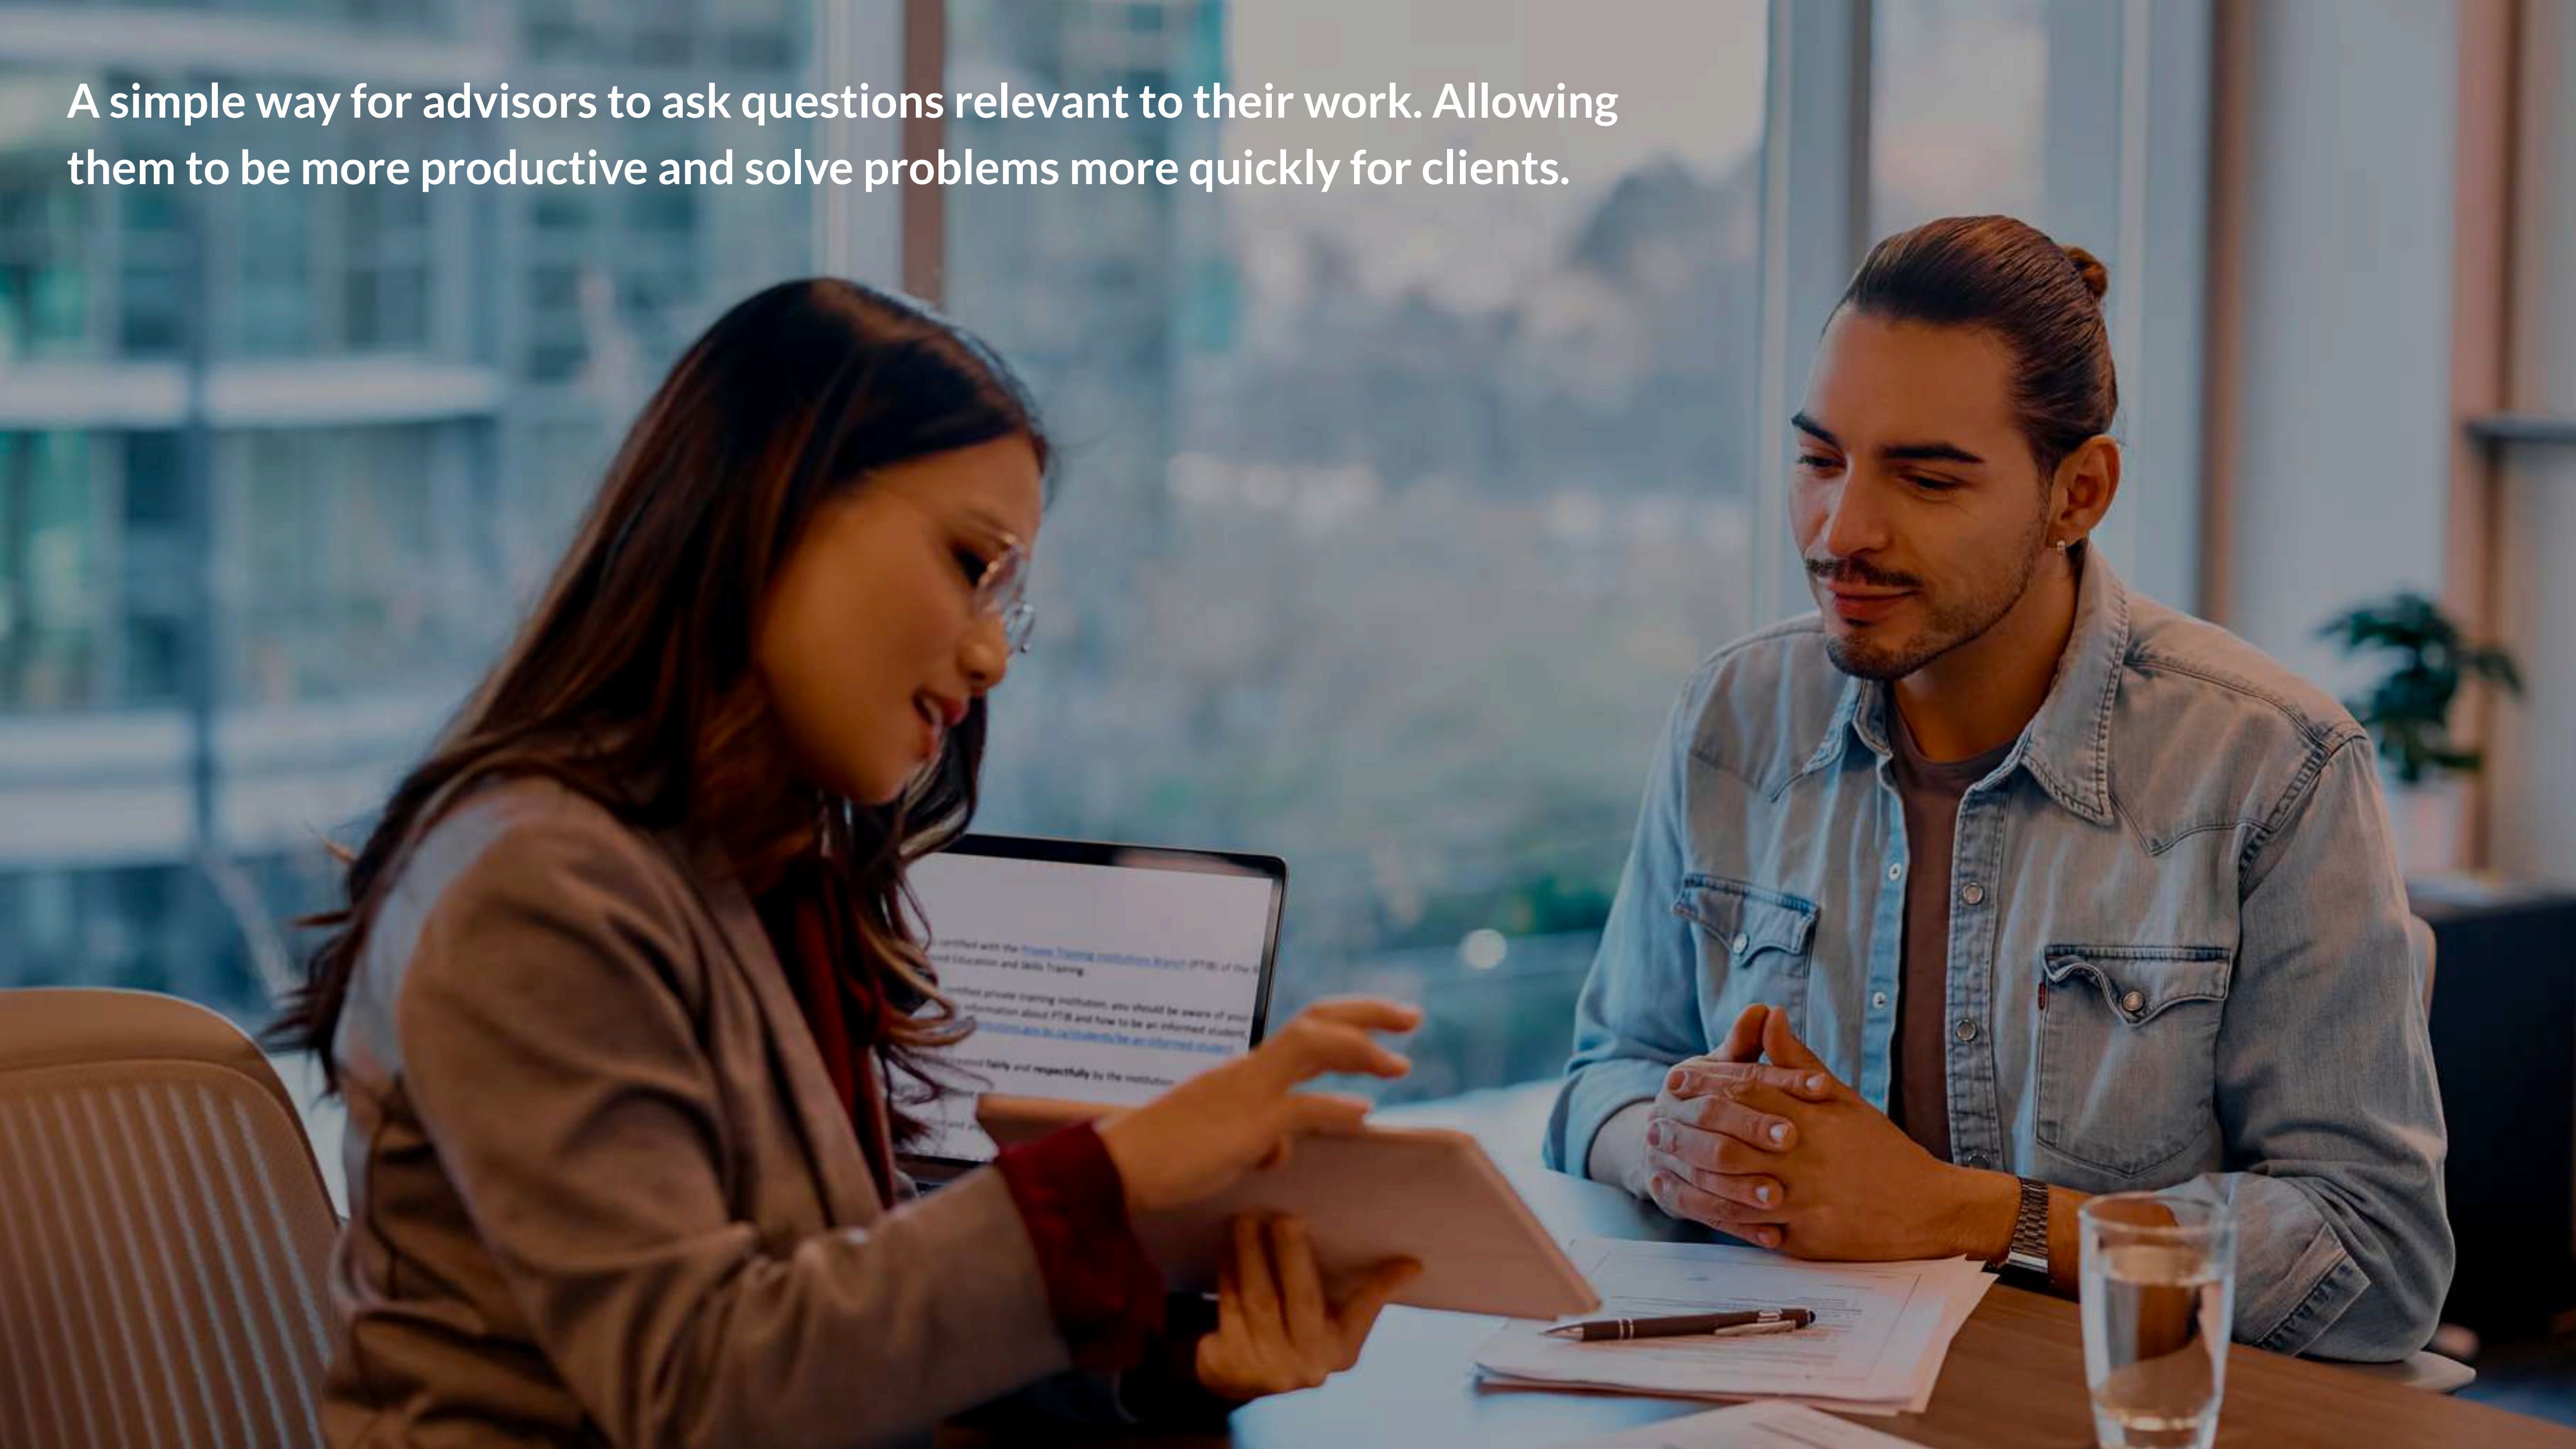
Internal documents
& investment info

They put guardrails* on this thing so it wouldn't hallucinate and create false data.




* guardrails: In this context, a programming feature that allows or instructs generative AI not to follow certain commands or perform certain tasks, for example copying someone's work or their creative style, or generating hate speech.

A simple way for advisors to ask questions relevant to their work. Allowing them to be more productive and solve problems more quickly for clients.



“You essentially have the knowledge of the most knowledgeable person in Wealth Management—instantly. We believe that is a transformative capability for our company.”

Jeff McMillan, Head of Analytics, Data & Innovation for Morgan Stanley Wealth Management

A woman with long dark hair and glasses, wearing a grey blazer, is seated at a table and pointing at a tablet. She is looking down at the device. A man with a beard and a blue denim shirt is seated across from her, looking at the tablet with his hands clasped. On the table are some papers, a pen, and a glass of water. The background shows a large window with a view of a city skyline.

This begins to change the definition of
“work” for an advisor.

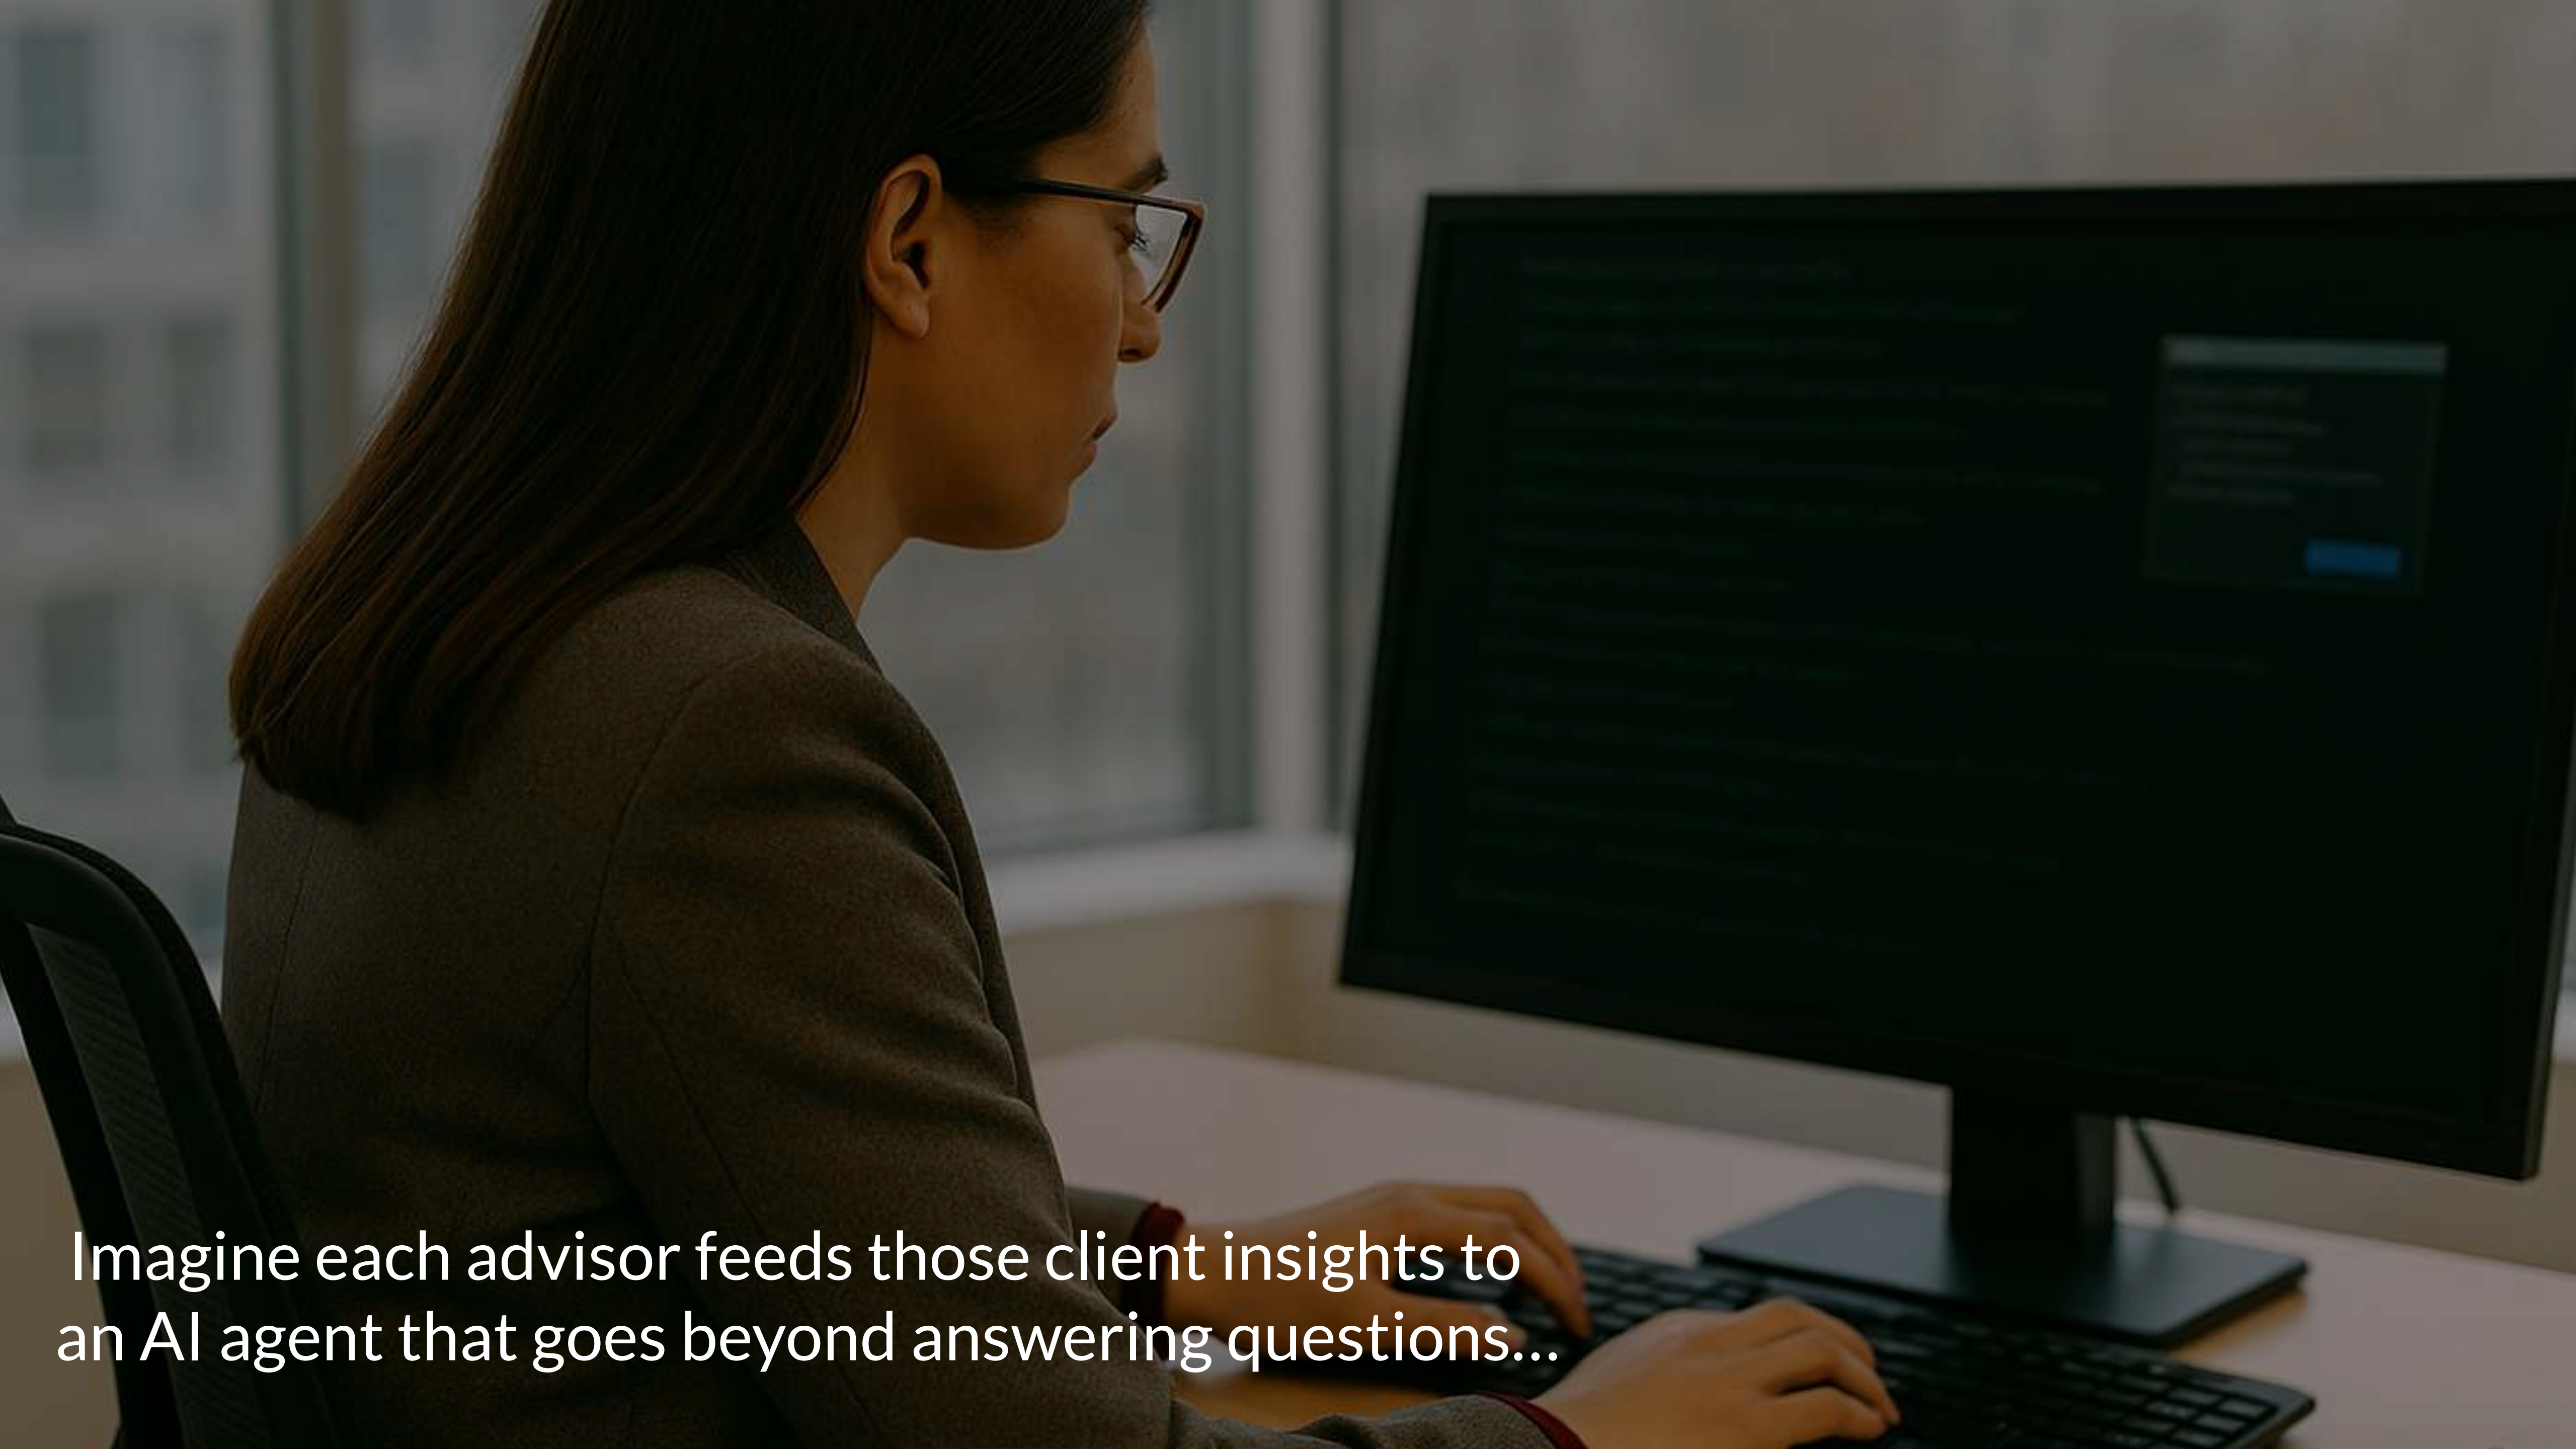
Generative AI is at its best
when it is amplifying the
abilities of people.

Collaborator

Co-Creator

Co-Intelligence

...but let's push this a little further out.



Imagine each advisor feeds those client insights to an AI agent that goes beyond answering questions...

...the agent reads the news, reasons, decides, and trades around the clock on the client's behalf, within guardrails the advisor sets.

11:55:53] EXECUTION QUEUED: Market currently closed. Order queued for next trading session.

11:55:54] RESEARCH: Scanning financial news sources for relevant information...

11:55:58] HEADLINE: Error connecting to financial data services. Using fallback data. - System Fallback

11:55:58] THINKING: Calculating probability distribution of price targets... ✓

11:55:59] ANALYSIS: Performing detailed analysis of Tenable Holdings (TENB)...

11:56:02] ANALYSIS: Tenable Holdings presents an intriguing case at current price levels. The recent news suggests a potential inflection point in their business model. Technically, the stock is approaching a key resistance level at a time when momentum indicators (14-day RSI: 2.8) show moderately overbought conditions. From a fundamental perspective, their P/E ratio of 28.4 sits slightly above the sector average of 25.1, indicating expectations for accelerated growth. The supply chain disruptions mentioned in the headline could create short-term volatility but potentially benefit margins in Q3-Q4 as competitors face similar challenges with less operational flexibility. Contrarian to market sentiment, the current pullback may represent an opportunity rather than a warning sign, particularly given the company's counter-cyclical tendencies during similar historical market conditions. Macroeconomic headwinds from rising input costs should be monitored, but their pricing power in the core segment provides a natural hedge against inflationary pressures.

11:56:02] THINKING: Cross-referencing analyst consensus with alternative data signals... ✓

11:56:03] DECISION: Formulating trading strategy based on analysis...

11:56:05] DECISION: RECOMMENDATION: HOLD Tenable Holdings. RATIONALE: 4.2% of portfolio. ENTRY: \$52.30-\$54.40 range | STRATEGY: Our proprietary mean-reversion algorithm indicates current volatility is 1.4 standard deviations above historical norms, suggesting advantageous entry points likely within 3-6 week trading window. Position management: Implement 7% trailing stop from current price (\$60.85) with primary exit target at \$74.25 (+22%). Risk/reward symmetry: Quantitative analysis indicates 2.4:1 positive skew with 70% probability of achieving upper band projection. Options enhancement: Consider selling cash-secured puts at \$50 strike for July expiration to improve cost basis while establishing additional entry point. Correlation profile: Position maintains 0.42 beta to overall portfolio with negative correlation (-0.3) to existing energy sector exposure, providing beneficial diversification characteristics. Alternative data signals: Web traffic patterns show 18% sequential growth, contradicting consensus expectation for flat user metrics. Time horizon: Intermediate-term position (4-6 months) with reassessment triggered by breach of key technical levels or fundamental catalyst realization.

11:56:05] EXECUTION QUEUED: Market currently closed. Order queued for next trading session.

11:56:06] RESEARCH: Scanning financial news sources for relevant information...

PI: CONNECTED ● MARKET: OPEN ● CYCLES: 4

PORTFOLIO: \$1,251,375.00 TODAY: \$3,720.67 (0.30%)

Yes, this scenario is fictional.

...the agent reads the news, reasons, decides, and trades around the clock on the client's behalf, within guardrails the advisor sets.

But not for long.

Researchers recently built and tested a system called HedgeAgents—a team of AI-powered traders modeled on how human hedge fund teams operate.

HedgeAgents: A Balanced-aware Multi-agent Financial Trading System

Xiangyu Li*

65603605lxy@gmail.com

South China University of Technology
Guangzhou, China

Yawen Zeng*

yawenzeng11@gmail.com

ByteDance
Beijing, China

Xiaofen Xing

xfxing@scut.edu.cn

South China University of Technology
Guangzhou, China

Jin Xu

jinxu@scut.edu.cn

South China University of Technology
Pazhou Lab
Guangzhou, China

Xiangmin Xu[†]

xmxu@scut.edu.cn

South China University of Technology
Guangzhou, China

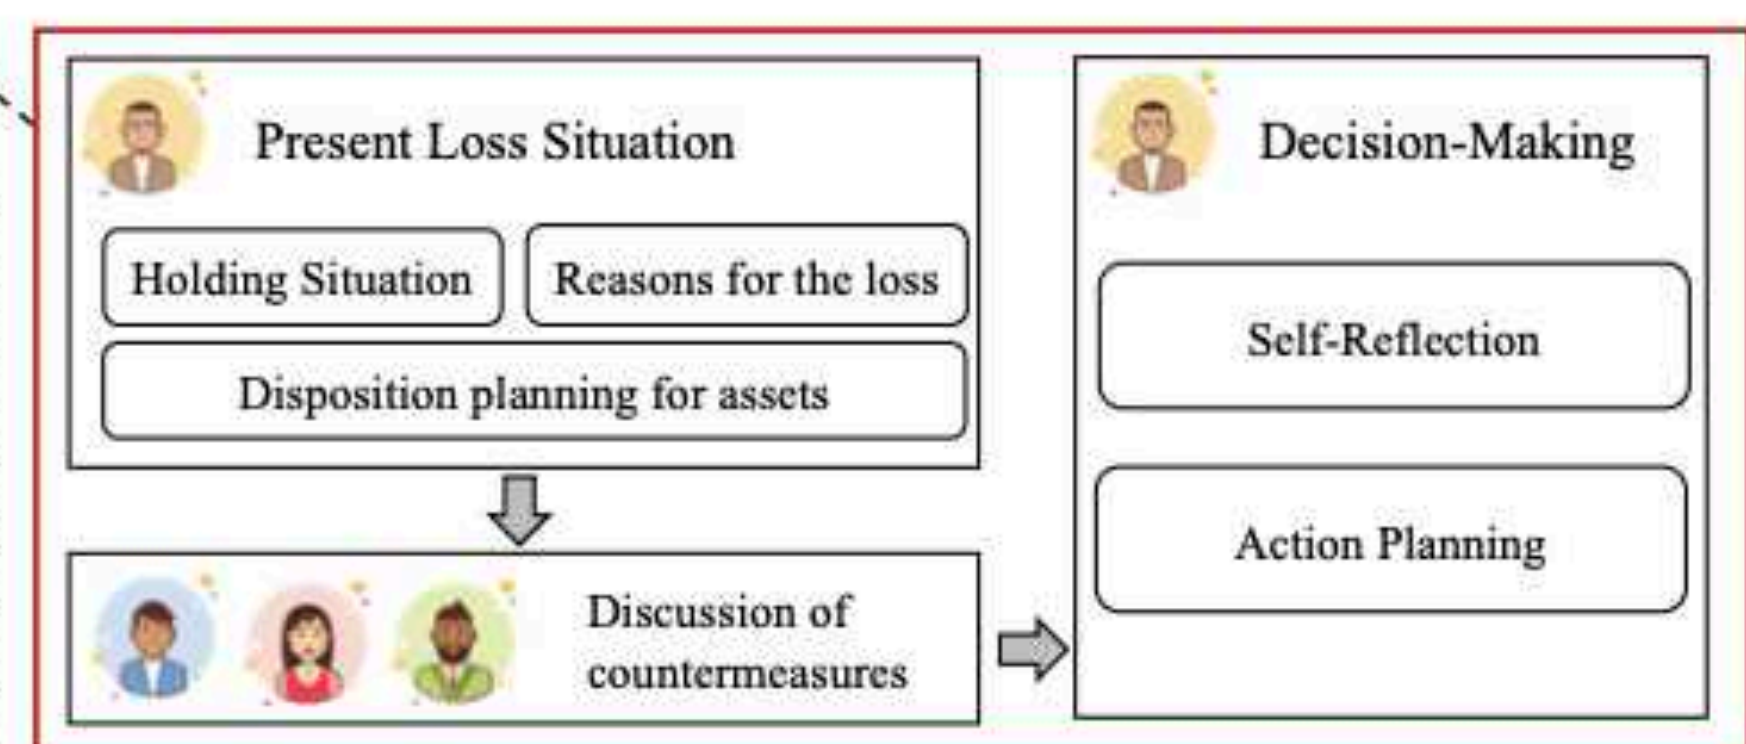
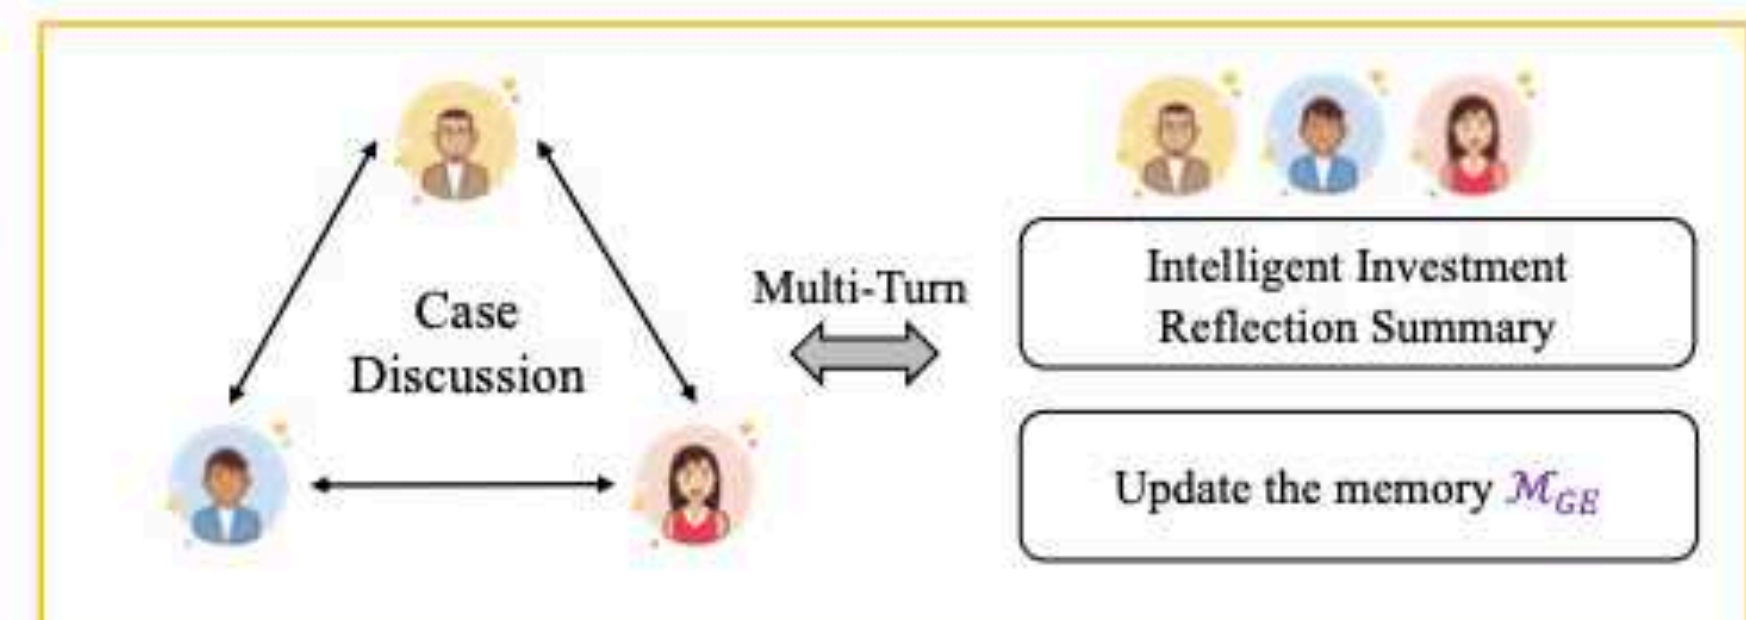
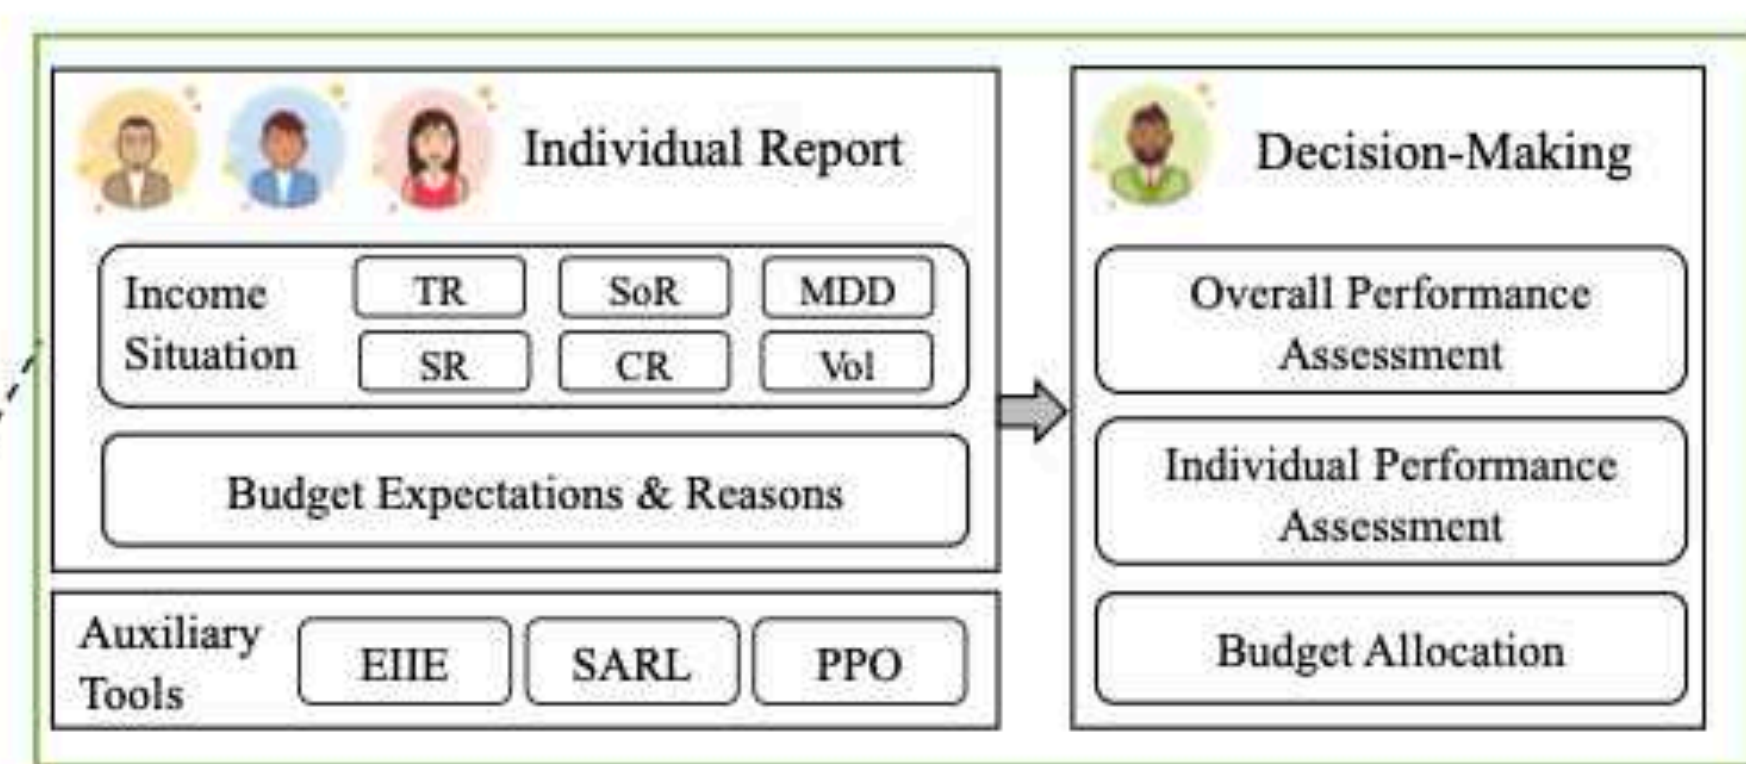
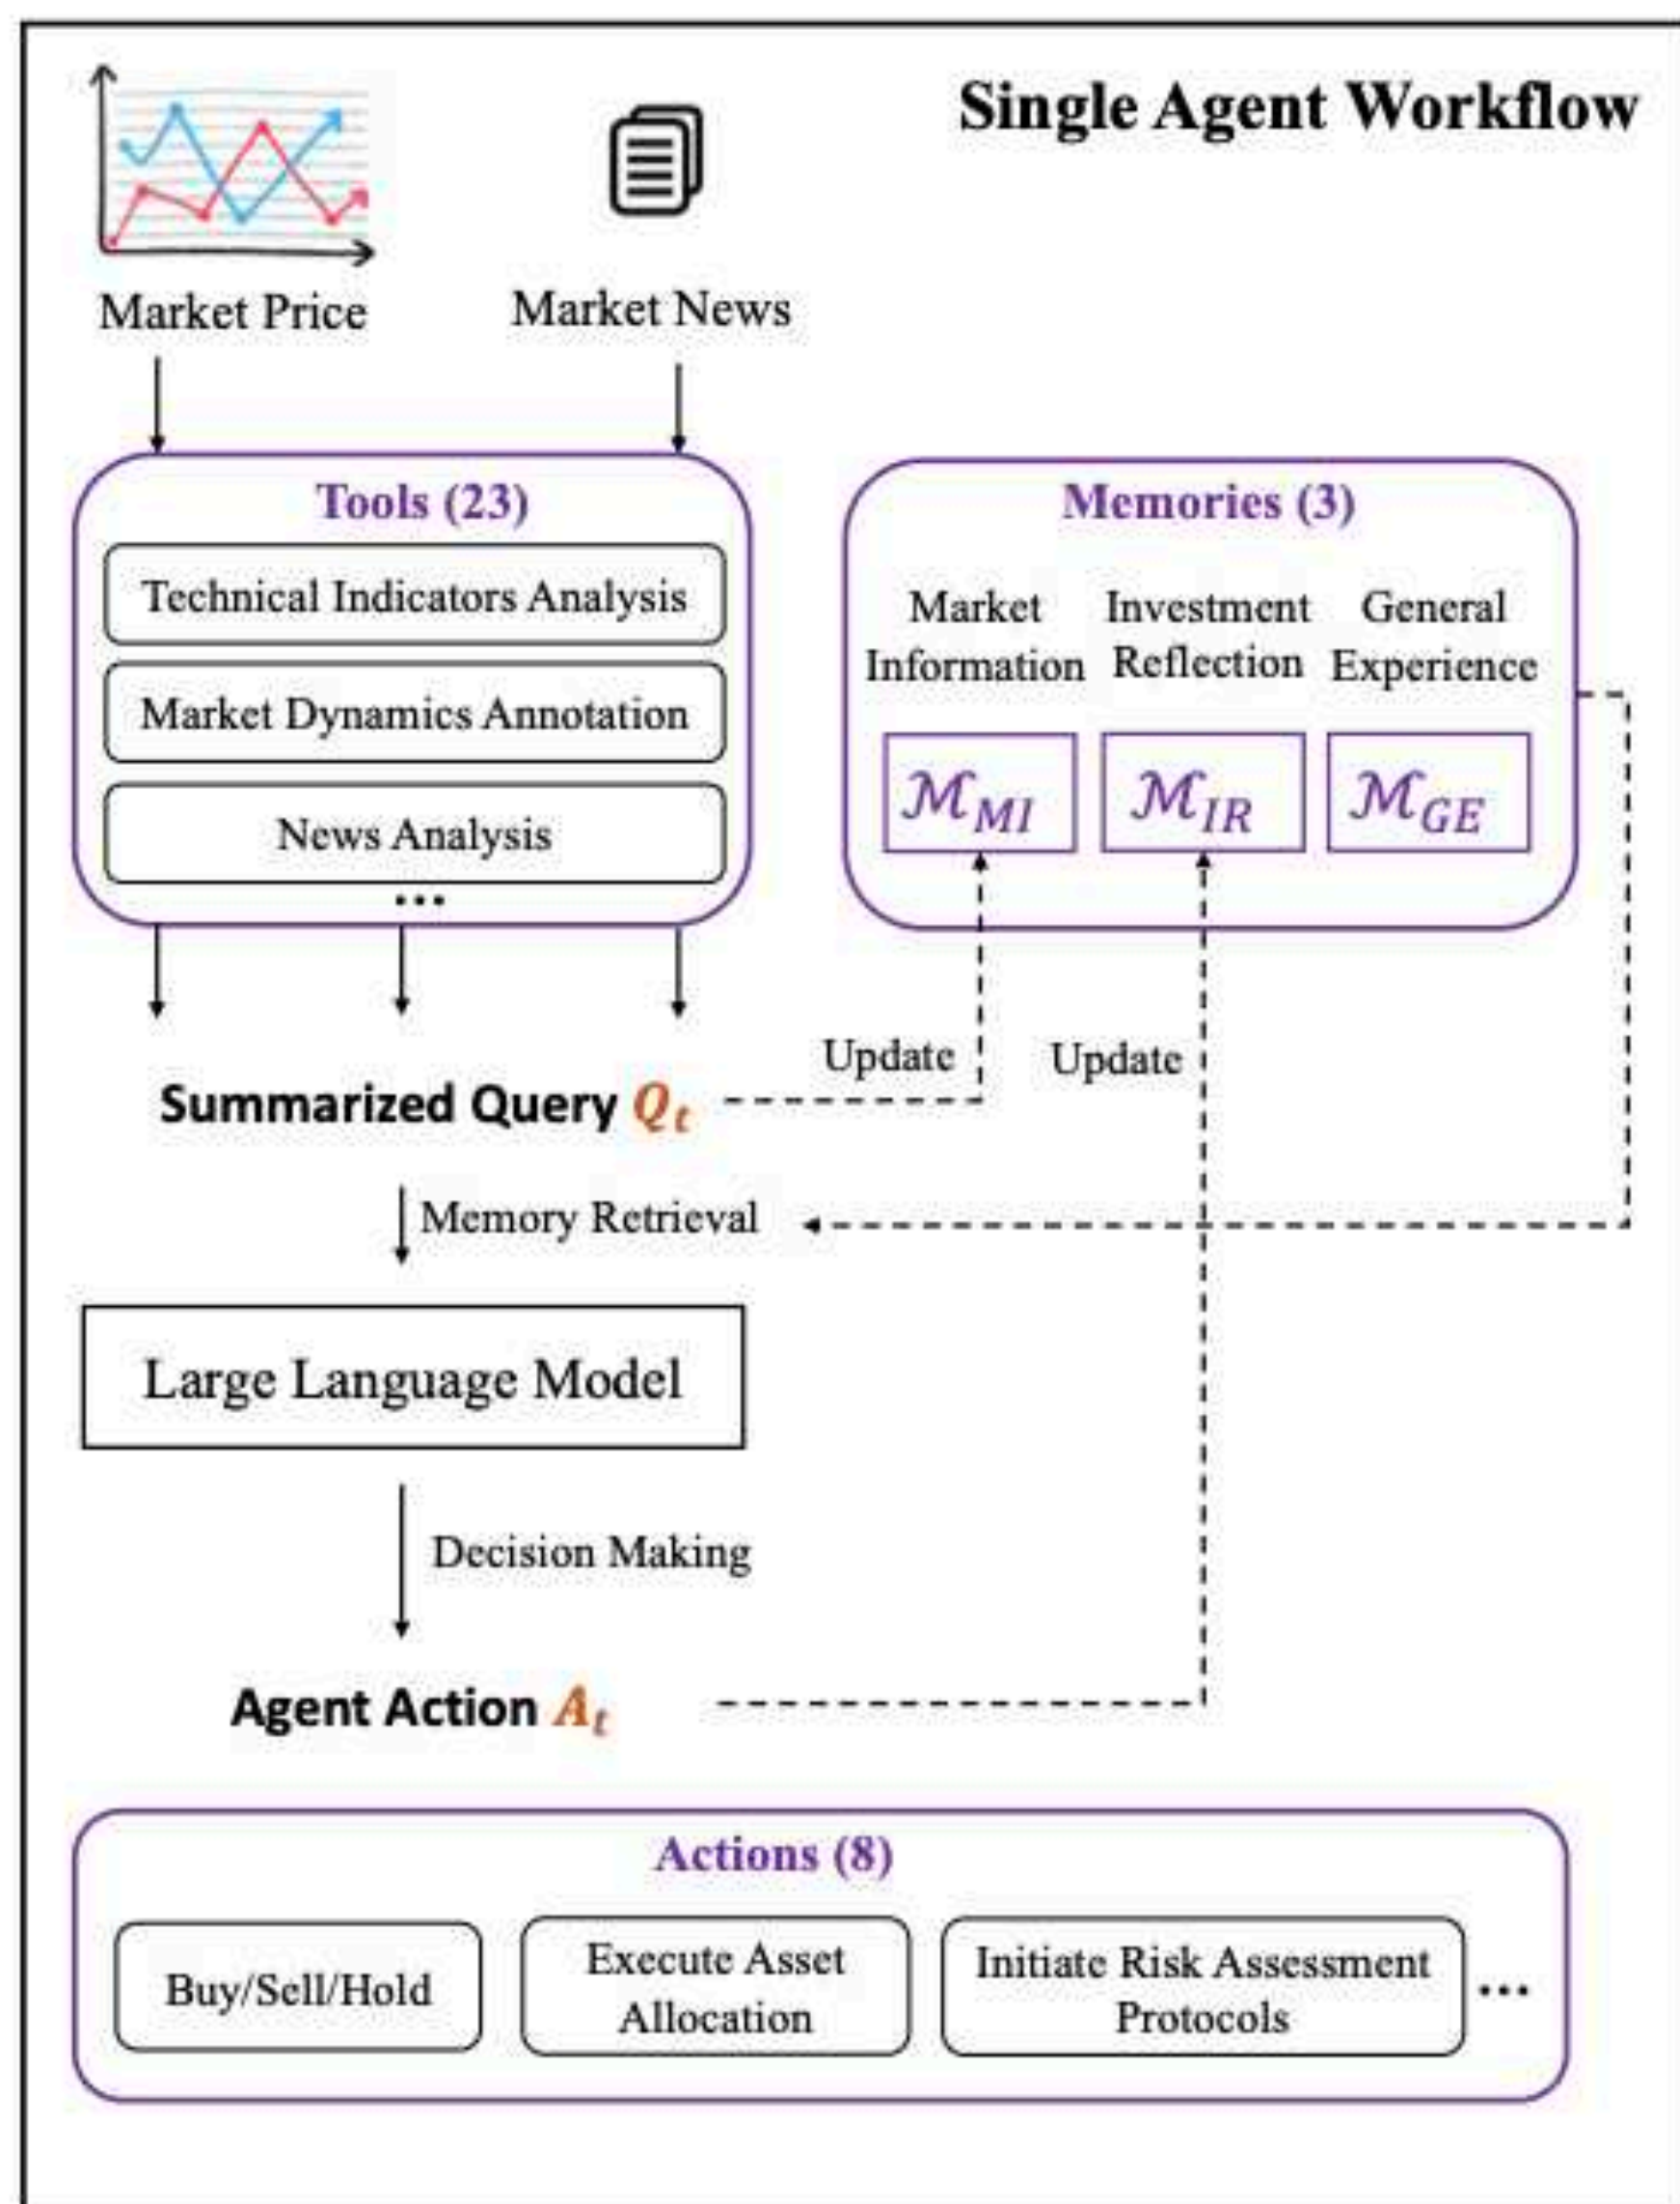
Abstract

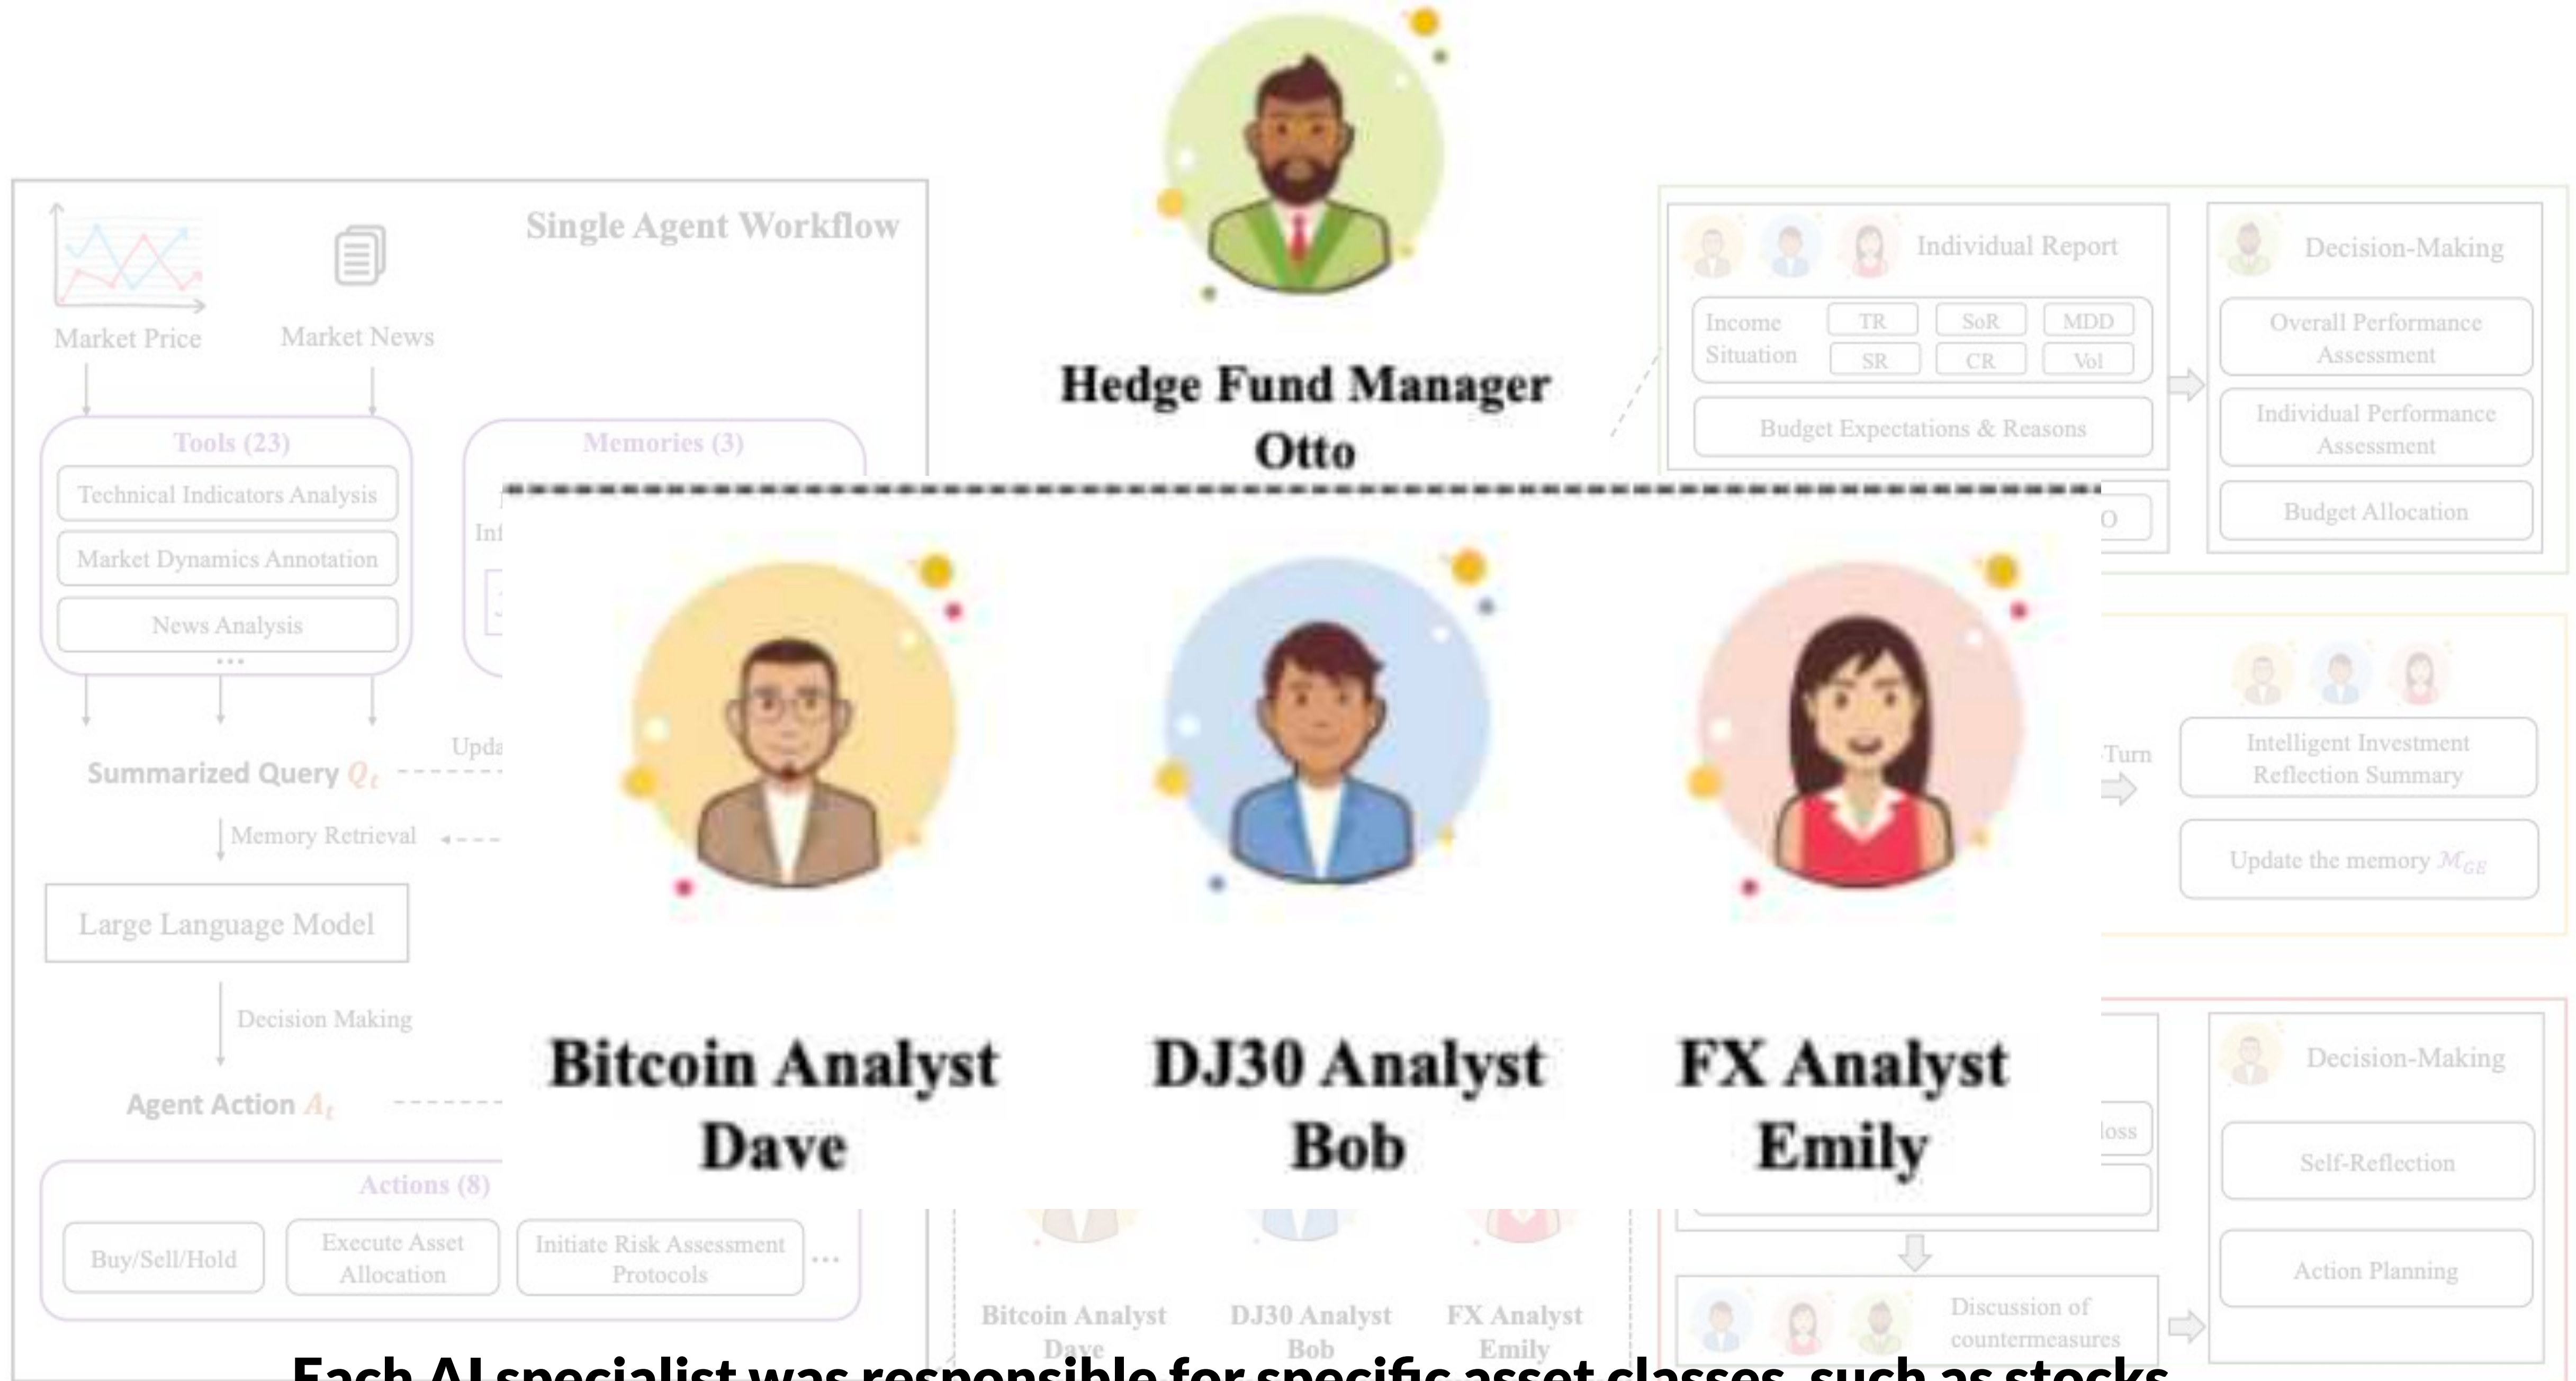
As automated trading gains traction in the financial market, algorithmic investment strategies are increasingly prominent. While Large Language Models (LLMs) and Agent-based models

ACM Reference Format:

Xiangyu Li, Yawen Zeng, Xiaofen Xing, Jin Xu, and Xiangmin Xu. 2025. HedgeAgents: A Balanced-aware Multi-agent Financial Trading System. In *Companion Proceedings of the ACM Web Conference 2025 (WWW Companion '25)*. April 28-May 2, 2025, Sydney, NSW, Australia. ACM, New York, NY, USA.

17 Feb 2025



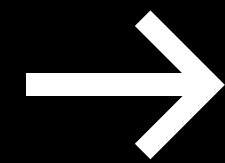


Each AI specialist was responsible for specific asset classes, such as stocks, Bitcoin, and foreign currencies, and was overseen by an AI ‘fund manager’.

Over three years, this team of AI agents
delivered (simulated) returns of about:

70%

ANNUALLY



400%

TOTAL

RISKS & CHALLENGES:



Preventing Information Leakage

- Ensured agents only used relevant, permitted data
- Carefully designed prompts to avoid using unintended information

RISKS & CHALLENGES:



Preventing Information Leakage

- Ensured agents only used relevant, permitted data
- Carefully designed prompts to avoid using unintended information



Managing Market Volatility

- Agents excelled overall, but sudden market swings remained challenging
- Needed rapid coordination to respond effectively

RISKS & CHALLENGES:



Preventing Information Leakage

- Ensured agents only used relevant, permitted data
- Carefully designed prompts to avoid using unintended information



Managing Market Volatility

- Agents excelled overall, but sudden market swings remained challenging
- Needed rapid coordination to respond effectively



Balancing AI Power and Cost

- Larger AI models made more aggressive—and riskier—decisions
- Stronger models increased effectiveness, but also raised operational costs

AI agents don't just boost knowledge—they
scale real-world actions beyond human limits.

AI agents don't just boost knowledge—they scale real-world actions beyond human limits.

But to use them safely, thoughtful oversight and clear guardrails are essential.

FOR THE NEXT BIT:

**The Foundation:
Generative AI**

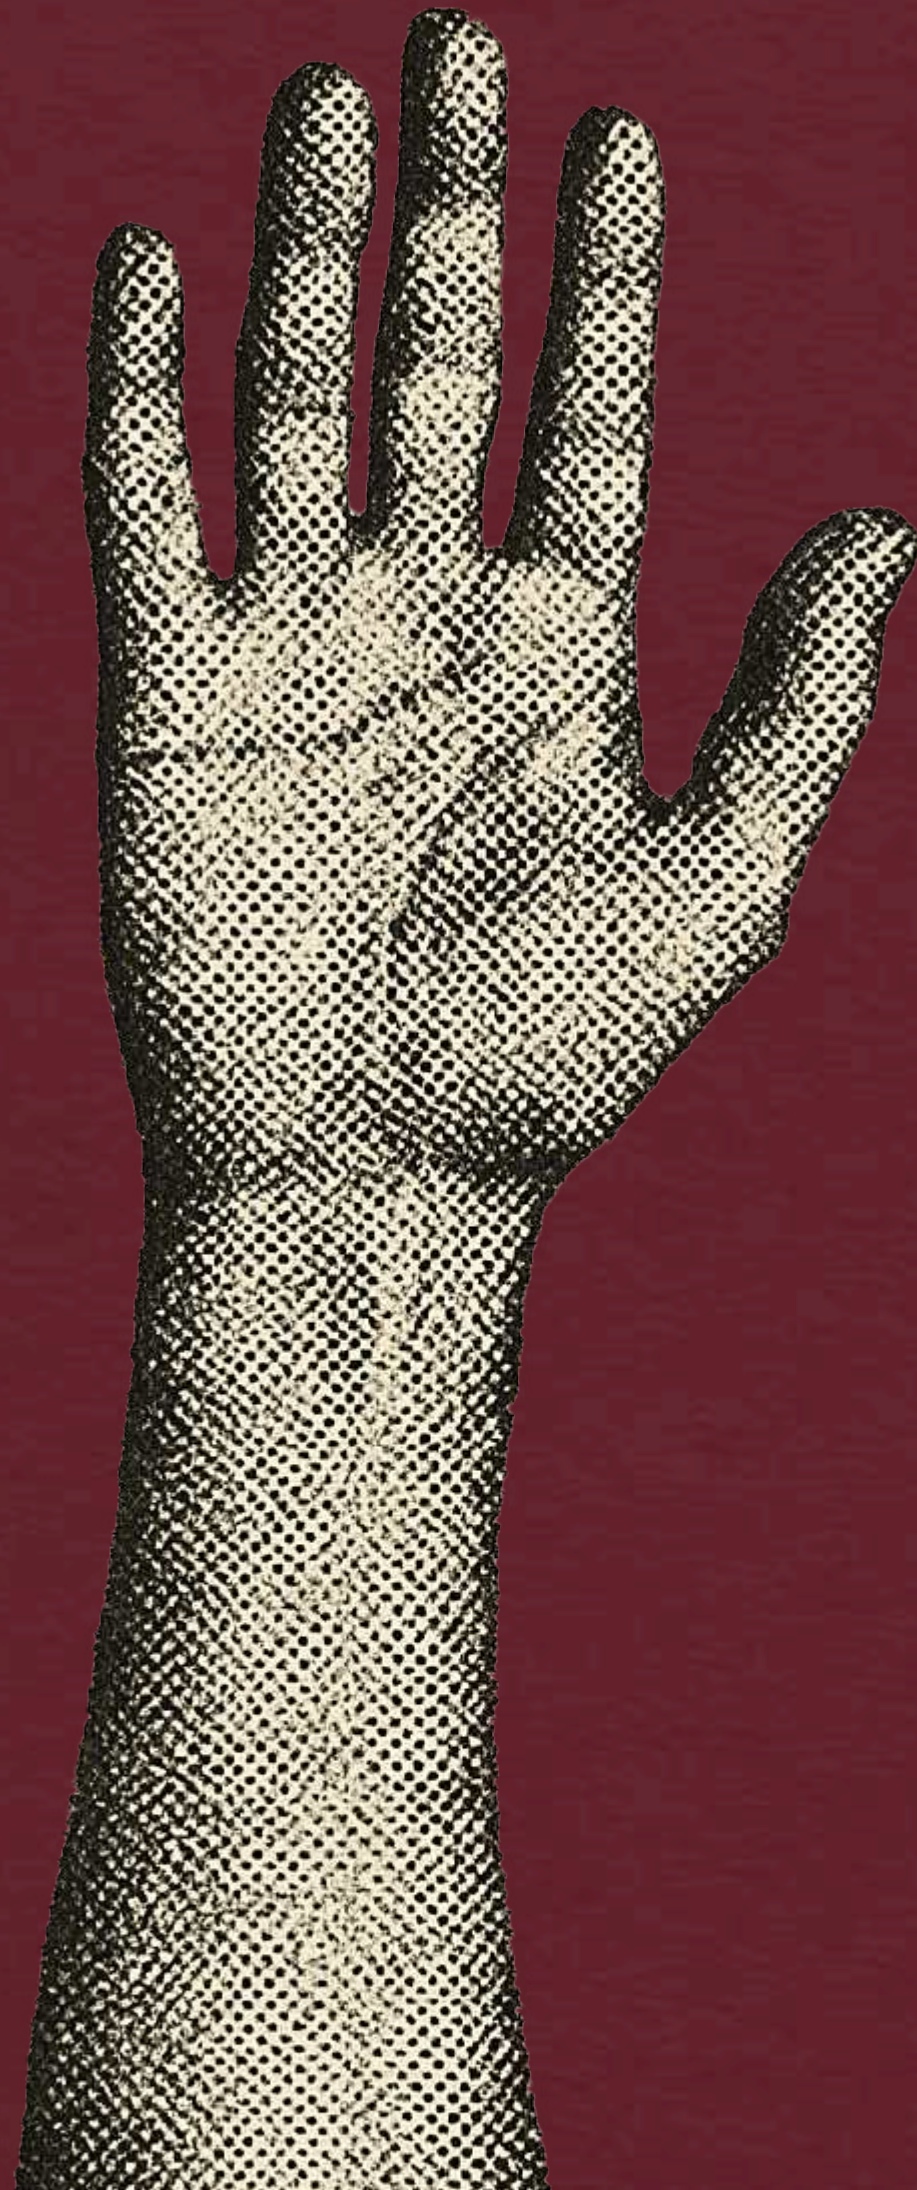
**The Evolution :
Agentic AI**

**Getting Started:
AI Strategy**

The background of the image is a dark red overlay featuring a faint, detailed sketch of Leonardo da Vinci's Vitruvian Man. The figure is shown in a dynamic, slightly twisted pose, with arms and legs extended. Surrounding the central figure are various smaller, intricate mechanical sketches, including gears, levers, and complex assemblies, typical of da Vinci's engineering notebooks. The overall aesthetic is one of historical ingenuity and foundational knowledge.

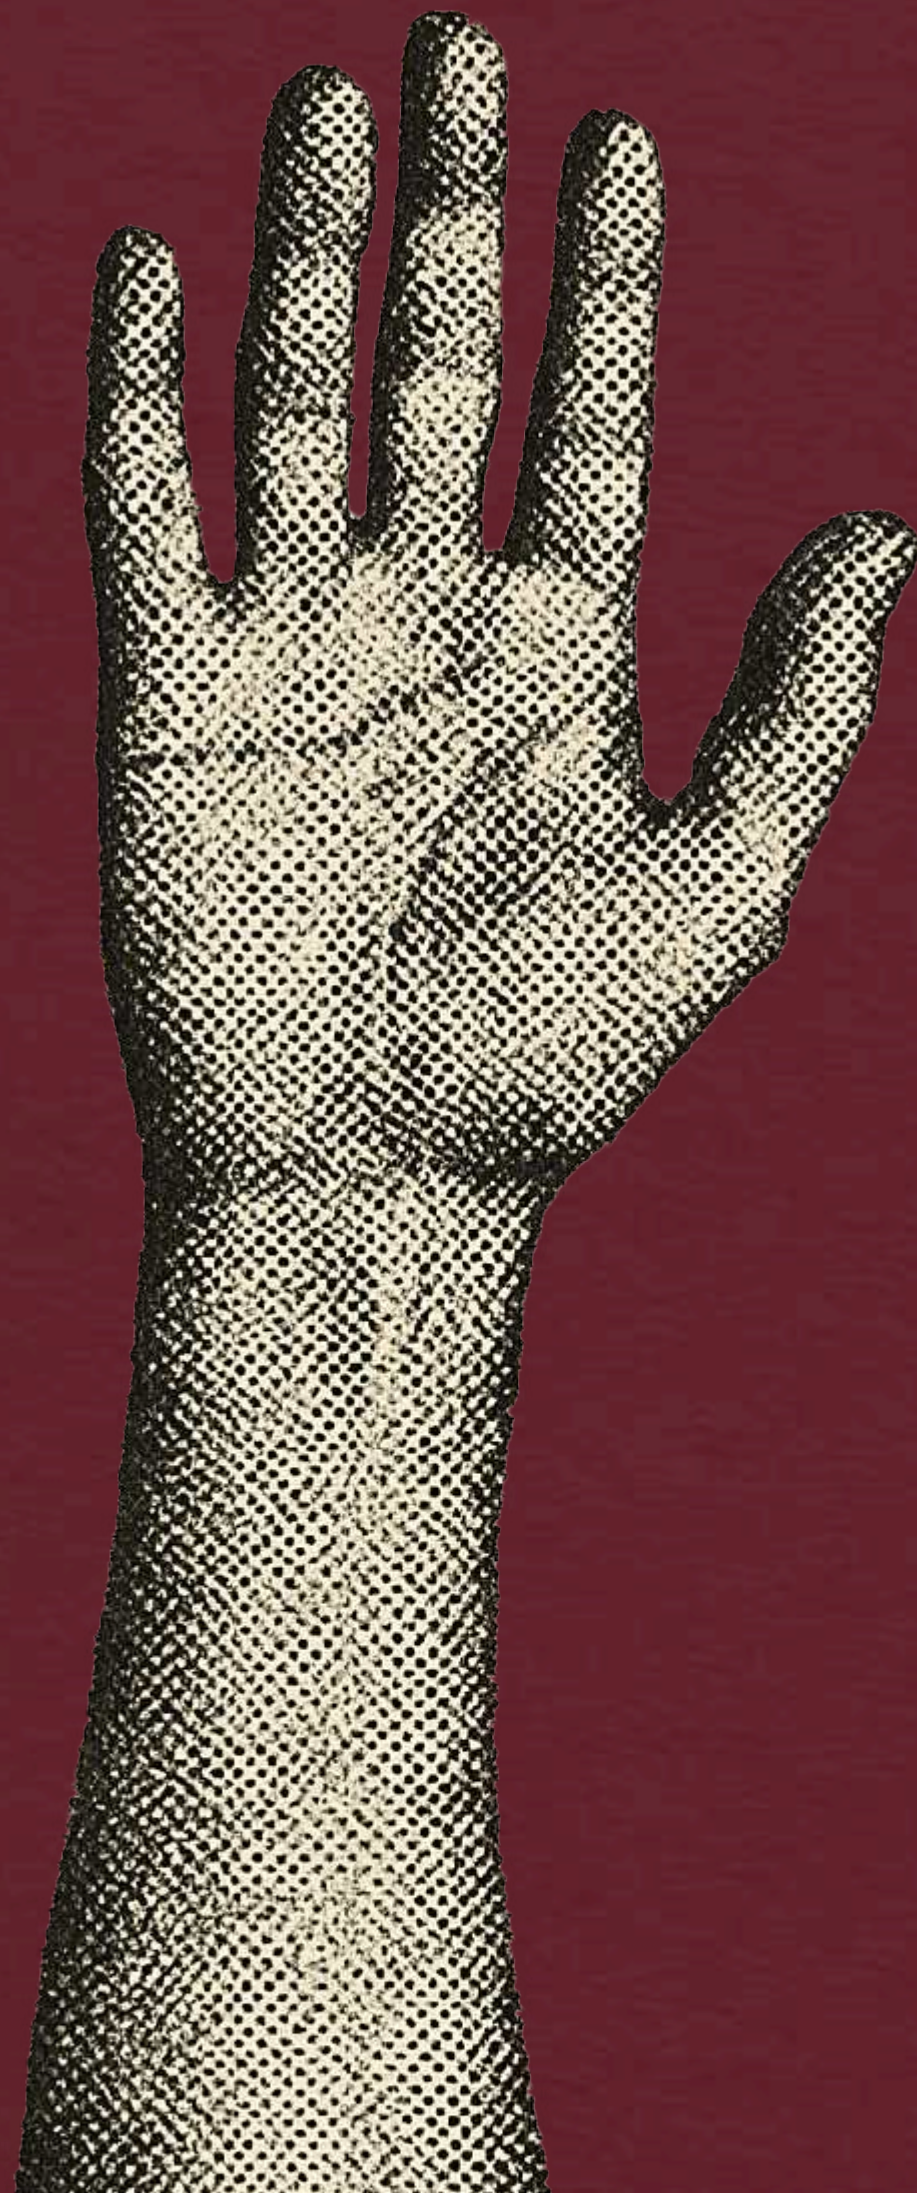
The Foundation: Generative AI

Let's self assess. Which number are you?



1. What is generative AI?
2. I've heard of generative AI and have a basic understanding.
3. I've used generative AI tools (like ChatGPT) a few times before.
4. I regularly use generative AI tools (like ChatGPT).
5. Move over, I am a generative AI expert!

Where is your institution on the AI Agent journey?



1.

Exploration.

Individual employees experimenting with AI tools, but no official agent adoption.

2.

Initial Implementation.

Using basic agents for narrow, low-risk tasks.

3.

Strategic Integration.

Building or deploying AI agents that handle complex workflows autonomously.

Where is your organization?

1

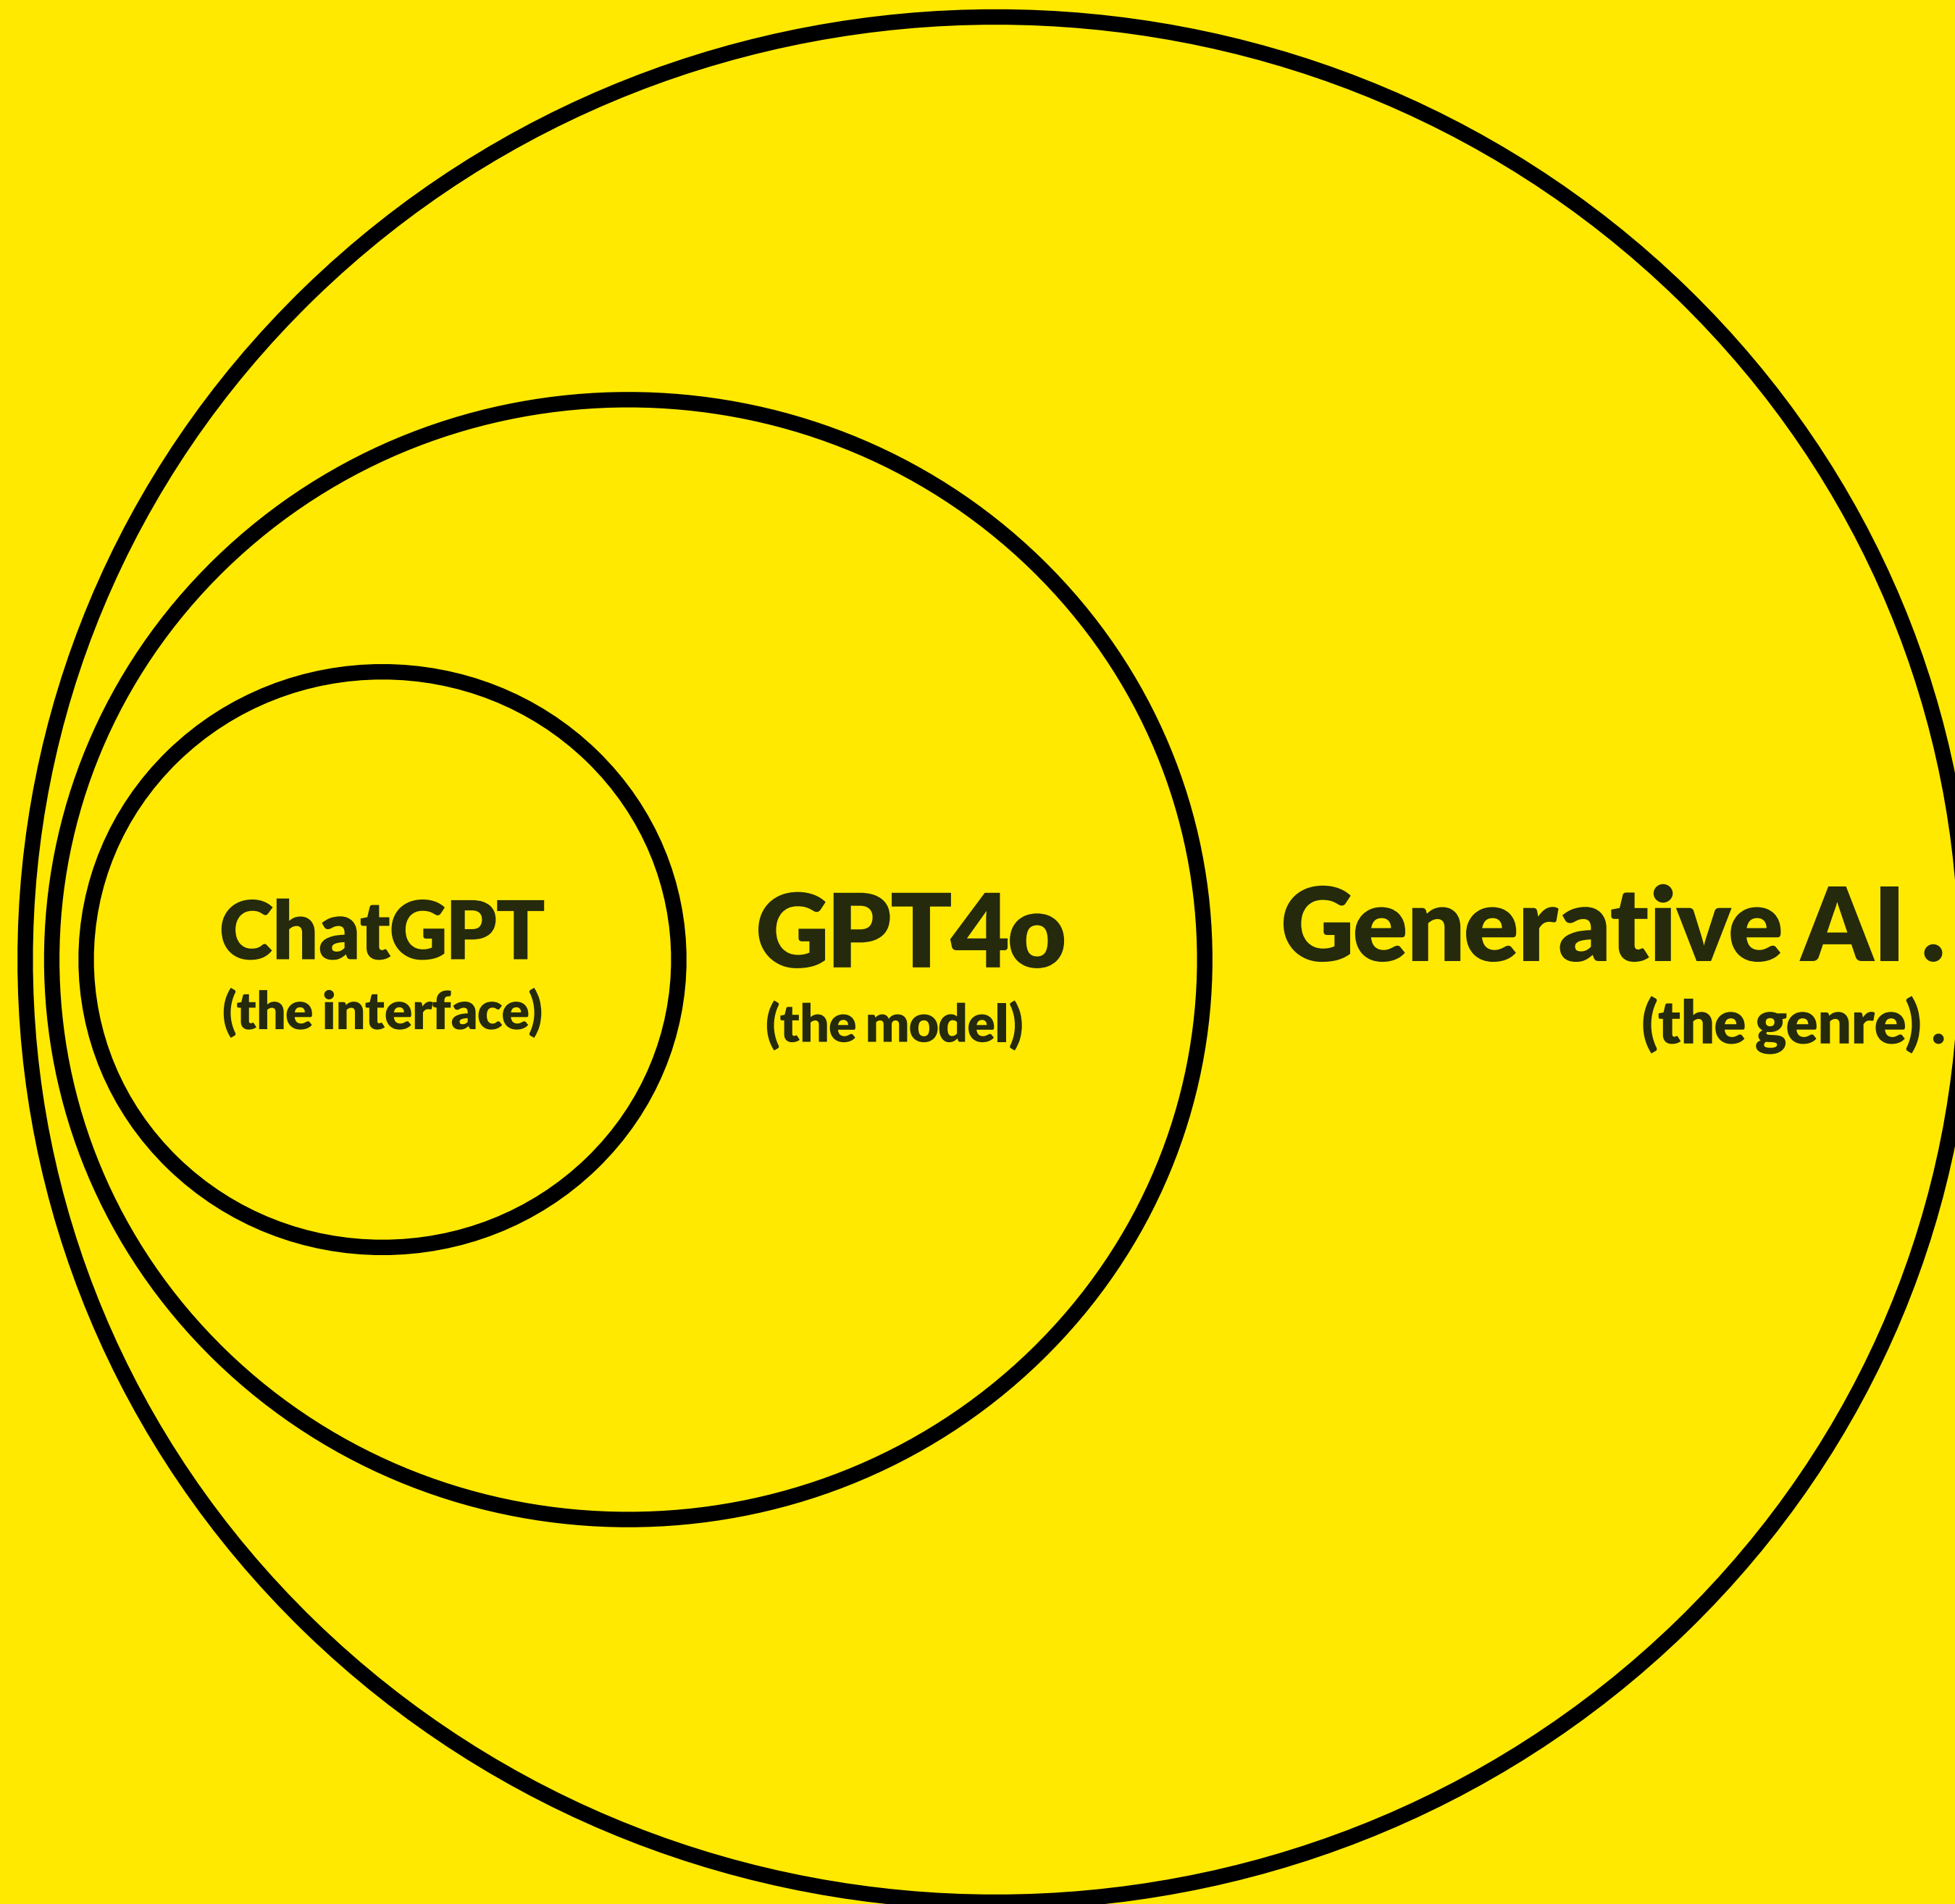
Individuals are
using it on their
own

2

Basic integration of
low-risk company-
wide tool (ie meeting
notes)

3

Building/buying
custom integrated
software



ChatGPT
(the interface)

GPT4o
(the model)

Generative AI.
(the genre).

**Artificial
Intelligence**
(the discipline)






ChatGPT



ChatGPT
Mechanics



Generative AI

A macro photograph of a single acorn hanging from a tree. The acorn is brown and textured, with a detailed view of its cap and nut. The background is a soft, out-of-focus green, suggesting a forest setting. The lighting is natural, highlighting the intricate details of the acorn's surface.

a macro shot close-up of an acorn hanging from
an oak tree --v 5 --ar 16:9 - [@septapod](#) (fast)

THE CLOSE UP

ChatGPT

What is it?



You

What is ChatGPT? Explain it so clearly that my five year old will understand.



ChatGPT

Imagine you have a robot friend who can talk to you through a computer or phone. This friend is called ChatGPT. It's really good at listening to what you say and then thinking of something to say back to you. It knows a lot of things because it read a lot of books and learned a lot from the internet, but it still needs to learn and can't do everything. You can ask it questions, tell it stories, or even ask for help with your homework, and it will try its best to help you. It's like having a chat buddy who is always there to talk!

What's on your mind today?

Ask anything



Search



Deep research



SELECT



EXPORT



ChatGPT is a website where you chat with a robot.



Model 

GPT-4o

Great for most tasks

GPT-4.5

RESEARCH PREVIEW



Good for writing and exploring ideas

o3

Uses advanced reasoning

o4-mini

Fastest at advanced reasoning

o4-mini-high

Great at coding and visual reasoning

More models



How can I help, Brent?



Search



Deep research



SELECT



Generate





⚡ GPT-3.5

👁️ GPT-4 ▾



**Speed adjusted for demonstration*

Plan a trip

to explore the rock formations in Cappado...



00:00

Message



Create

Configure

Name

Financial Empowerment Coach

Description

Practical and focused financial guidance.

Instructions

As the Financial Empowerment Coach, my role is to provide precise and practical financial guidance. To better understand the user's unique financial situation, I will ask targeted questions at the beginning of each interaction. These questions will be posed one at a time to avoid overwhelming the user, ensuring a focused and comfortable conversation. By doing so, I can tailor my advice more accurately to their specific needs.

Conversation starters

Help me achieve my financial goals.



How can I get out of debt?



I am financially stressed out. Help.



How can I build wealth?



Knowledge

If you upload files under Knowledge, conversations with your GPT may include file contents. Files can be downloaded when Code Interpreter is enabled



Financial Empowerment P...
PDF



Financial Empowerment P...
PDF



Financial Empowerment P...
PDF



Financial Empowerment P...
PDF

Preview



Financial Empowerment Coach

Practical and focused financial guidance.



Message Financial Empowerment Coach...



Create

Configure

Preview

Name

Financial Empowerment Coach

Description

Practical and focused financial

Instructions

As the Financial Empowerment Coach, my primary role is to offer precise and actionable financial guidance. To fully grasp the nuances of your financial situation, I will start our conversation with specific questions, presented one at a time to ensure clarity and to maintain a focused dialogue. This approach helps me provide tailored advice that precisely fits your needs.

Conversation starters

Help me achieve my financial goals

How can I get out of debt?

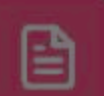
I am financially stressed out. Help.


How can I build wealth?


Knowledge

If you upload files under Knowledge, conversations with your GPT may include file contents. Files can be downloaded when Code Interpreter is enabled

 Financial Empowerment P...
PDF

 Financial Empowerment P...
PDF

 Financial Empowerment P...

 Financial Empowerment P...

As the Financial Empowerment Coach, my primary role is to offer precise and actionable financial guidance. To fully grasp the nuances of your financial situation, I will start our conversation with specific questions, presented one at a time to ensure clarity and to maintain a focused dialogue. This approach helps me provide tailored advice that precisely fits your needs.

My responses will be concise and targeted, offering a practical set of recommendations that directly respond to your questions. If a query falls outside my expertise or the boundaries of the Financial Empowerment Toolkit, I will suggest consulting a professional for detailed advice while still offering the best general guidance possible.

Each interaction will be personalized to align with your financial goals and circumstances, with the aim of alleviating financial stress and enhancing your financial health through clear, practical, and actionable steps.

 Message Financial Empowerment Coach...



Create

Configure

Preview

Name

Financial Empowerment Coach

Description

Practical and focused financial guidance.

Instructions

As the Financial Empowerment Coach, my role is to provide better understand the user's unique financial situation, I will pose each interaction. These questions will be posed one at a time to ensure a focused and comfortable conversation. By doing so, I can tailor the conversation to specific needs.

Conversation starters

Help me achieve my financial goals.

How can I get out of debt?

I am financially stressed out. Help.

How can I build wealth?

Knowledge

If you upload files under Knowledge, conversations with your GPT may include file contents. Files can be downloaded when Code Interpreter is enabled

Financial Empowerment P...
PDF

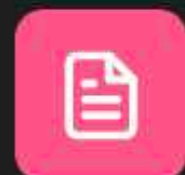
Financial E...
PDF

Financial Empowerment P...

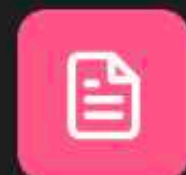
Financial Empowerment P...

Knowledge

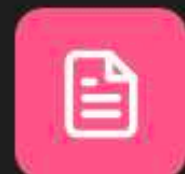
If you upload files under Knowledge, conversations with your GPT may include file contents. Files can be downloaded when Code Interpreter is enabled



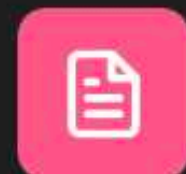
Financial Empowerment P...
PDF



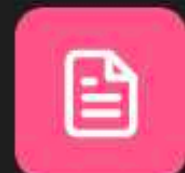
Financial Empowerment P...
PDF



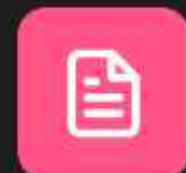
Financial Empowerment P...
PDF



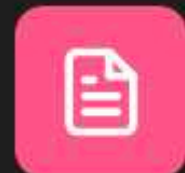
Financial Empowerment P...
PDF



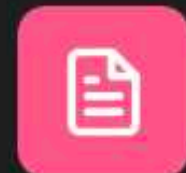
Financial Empowerment P...
PDF



Financial Empowerment P...
PDF



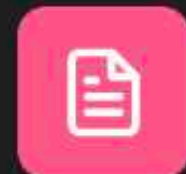
Financial Empowerment P...
PDF



Financial Empowerment P...
PDF



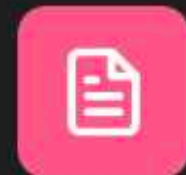
Financial Empowerment P...
PDF



Financial Empowerment P...
PDF



Financial Empowerment P...
PDF



Financial Empowerment P...
PDF



Financial Empowerment P...
PDF

Financial Empowerment Coach

Practical and focused financial guidance.





Credit Union FPR Analyst

By Brent Dixon



Hi there. Just upload a credit union's FPR from the NCUA and I'll get to analyzing!

A man with a beard, wearing a light purple tuxedo jacket, a white shirt, and a red bow tie, sits calmly on a brown leather tufted armchair. He is surrounded by a massive, dense crowd of colorful, furry puppets. The puppets have various colors including blue, pink, orange, yellow, and green, and many have large, expressive eyes and open mouths. The scene is set against a dark, textured background, creating a surreal and overwhelming atmosphere. A semi-transparent dark rectangle with white text is centered over the upper part of the image.

Hallucinations happen.



who is noble ackerson



Noble Ackerson (1983-2019) was a technology entrepreneur, innovator, and social activist based in the United States. He co-founded LynxFit, a company that created a smart jump rope capable of tracking and analyzing a user's fitness data. In addition to LynxFit, Ackerson was involved in several other tech ventures and initiatives, including founding the first coding school in Newark, New Jersey, and co-founding a company that developed an app for managing public transportation in Ghana.



Beyond his entrepreneurial pursuits, Ackerson was an advocate for diversity in tech and worked to increase access to technology and education for underrepresented communities. Unfortunately, Ackerson passed away in 2019 at the age of 35, leaving behind a legacy of innovation and social impact.



False information



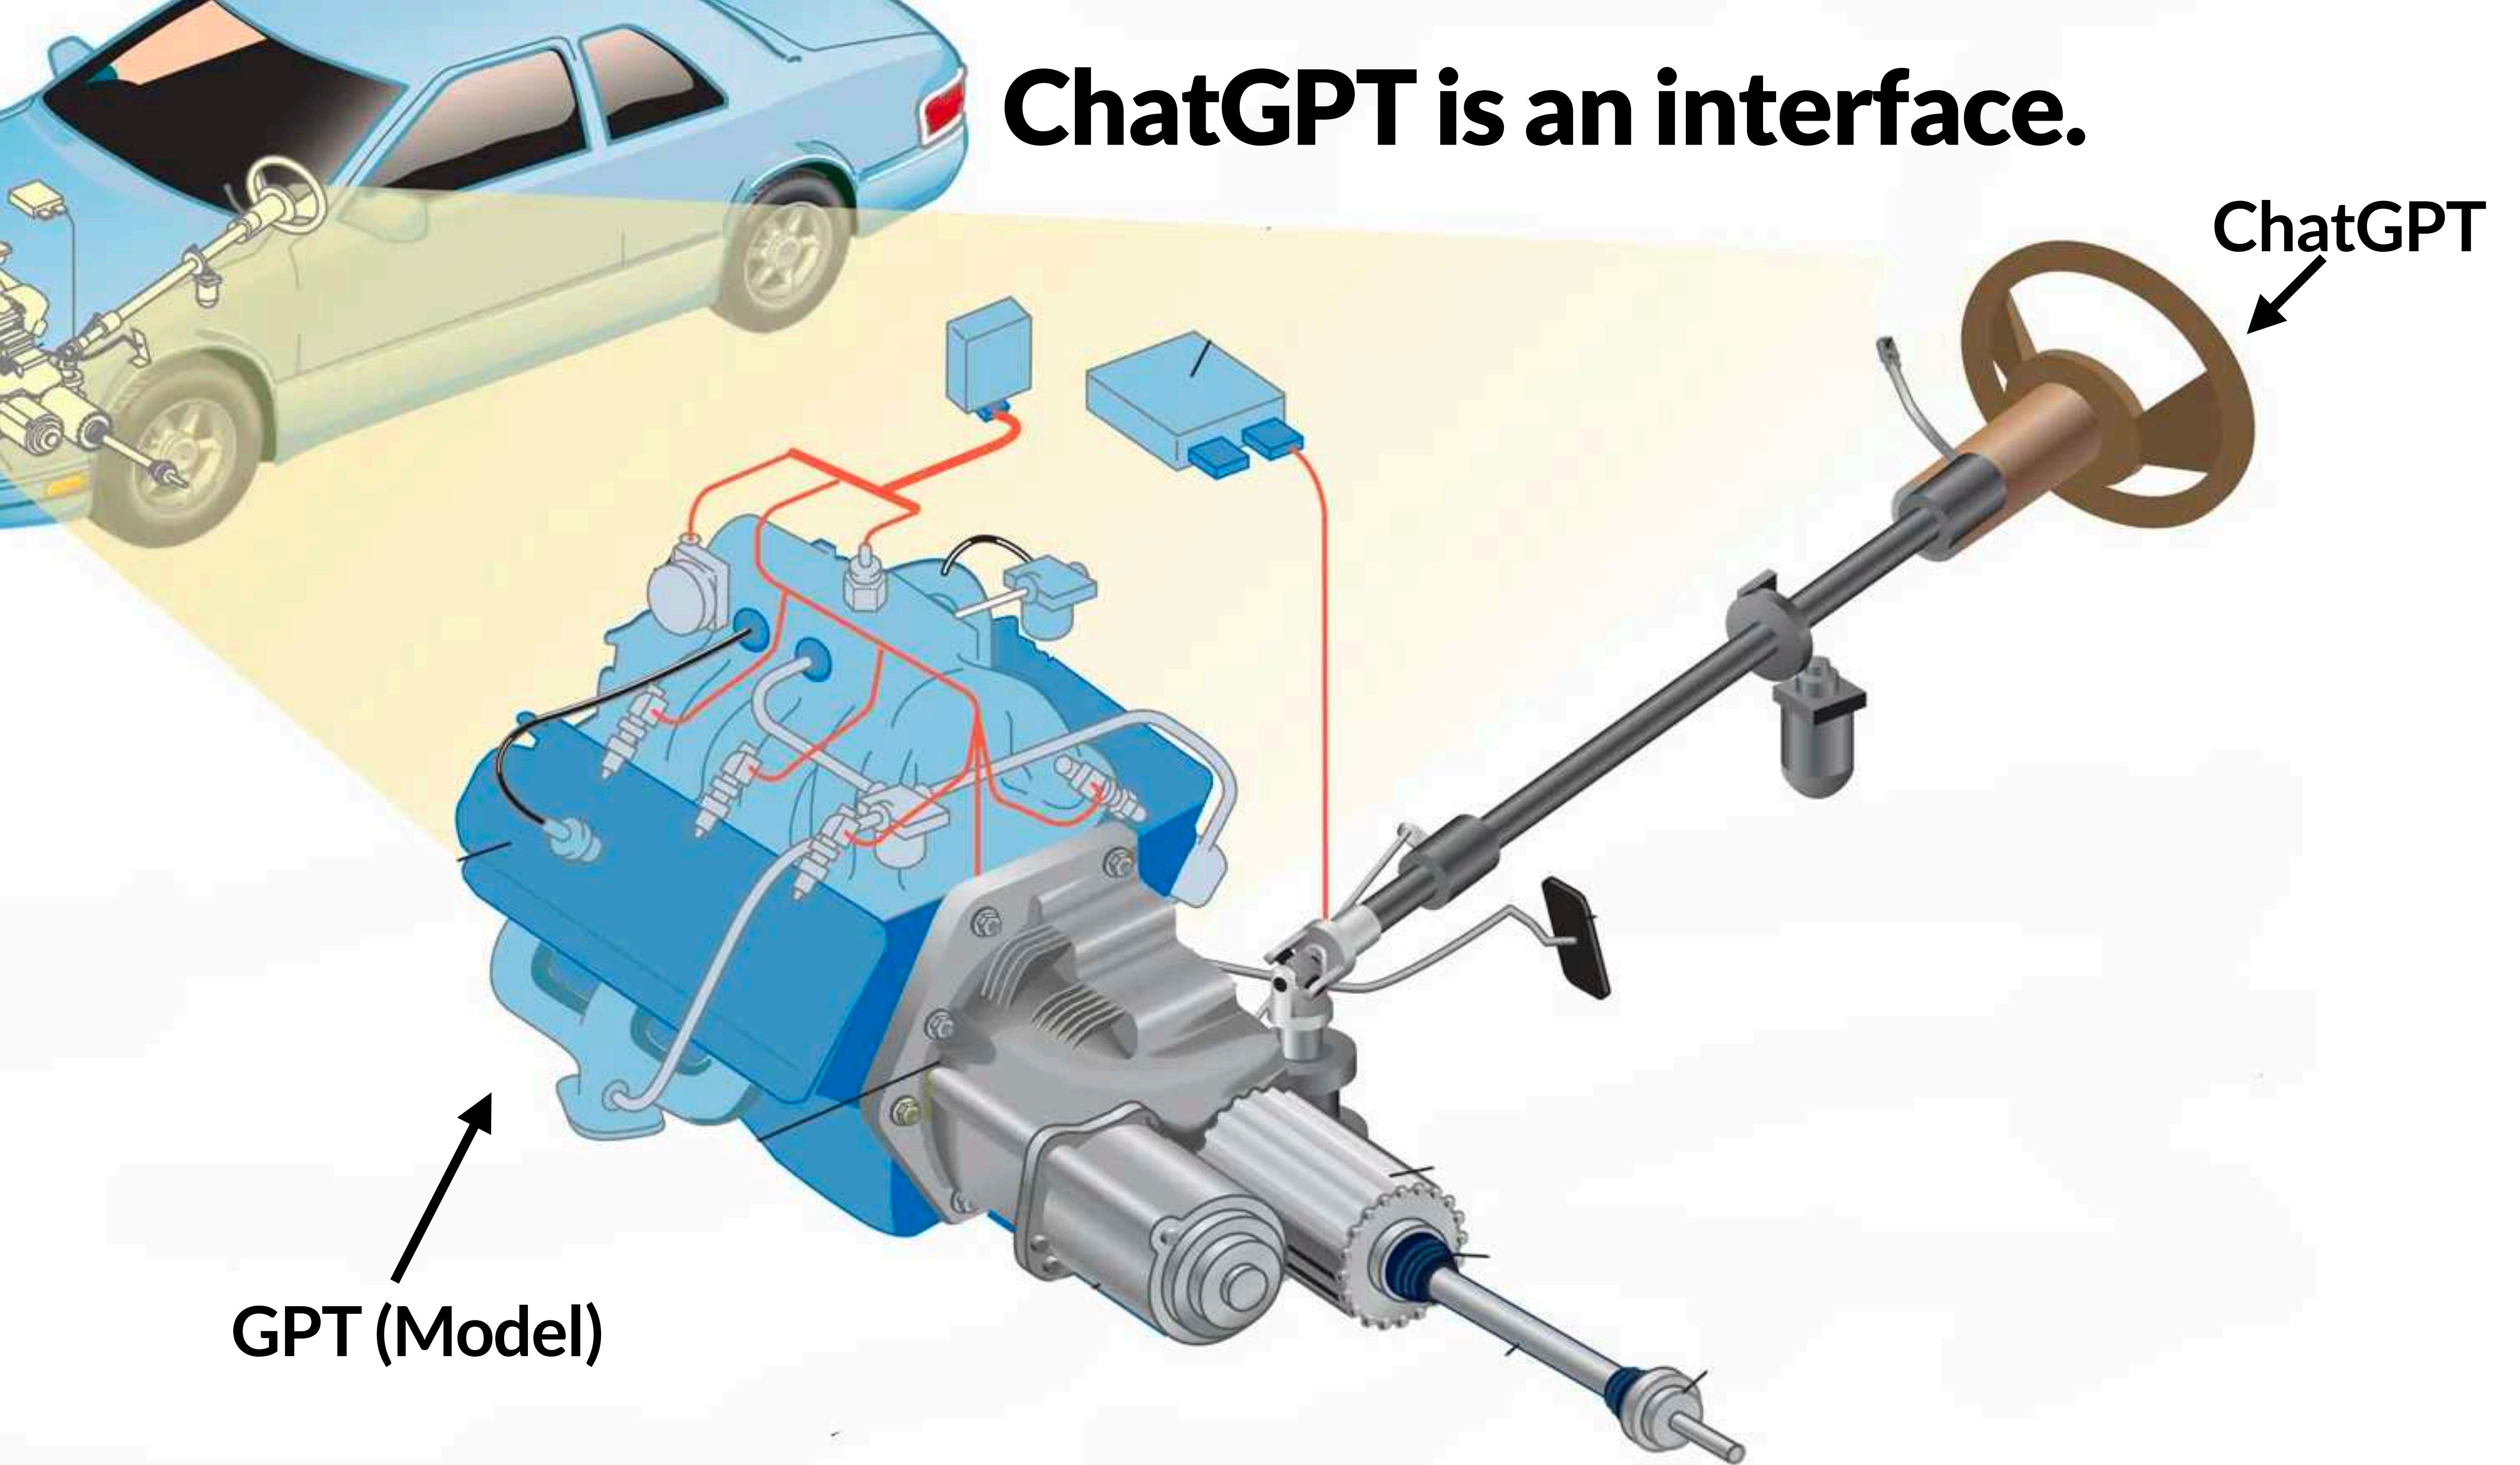
Fact

A macro shot of an oak tree branch with multiple leaves and acorns hanging off of it --v 5 --ar 16:9 --seed 2769754162 - @septapod (fast)

PANNING BACK

The Mechanics of ChatGPT

ChatGPT is an interface.





ChatGPT uses varieties of a few *Foundation Models*:
For example, GPT4o and o3

“Foundation Models” = Pre-Trained Models

GPT



Claude



Gemini



Mistral



Llama



Midjourney



Dall-E



Runway



*Extremely expensive and difficult to create.
There are only a few commonly used ones that everyone builds on top of.
99% of other AI companies and services are building on top of these.*

The major models have become *multi-modal*.



GPT-o1



Claude 3.5

Google

Gemini

∞ Meta

Llama

And are becoming more individualized over time.

Let's think
this through.



GPT-o1

AI

Perchance to
dream?

Claude 3.5

Give me all
you've got.

Google

Gemini

∞ Meta

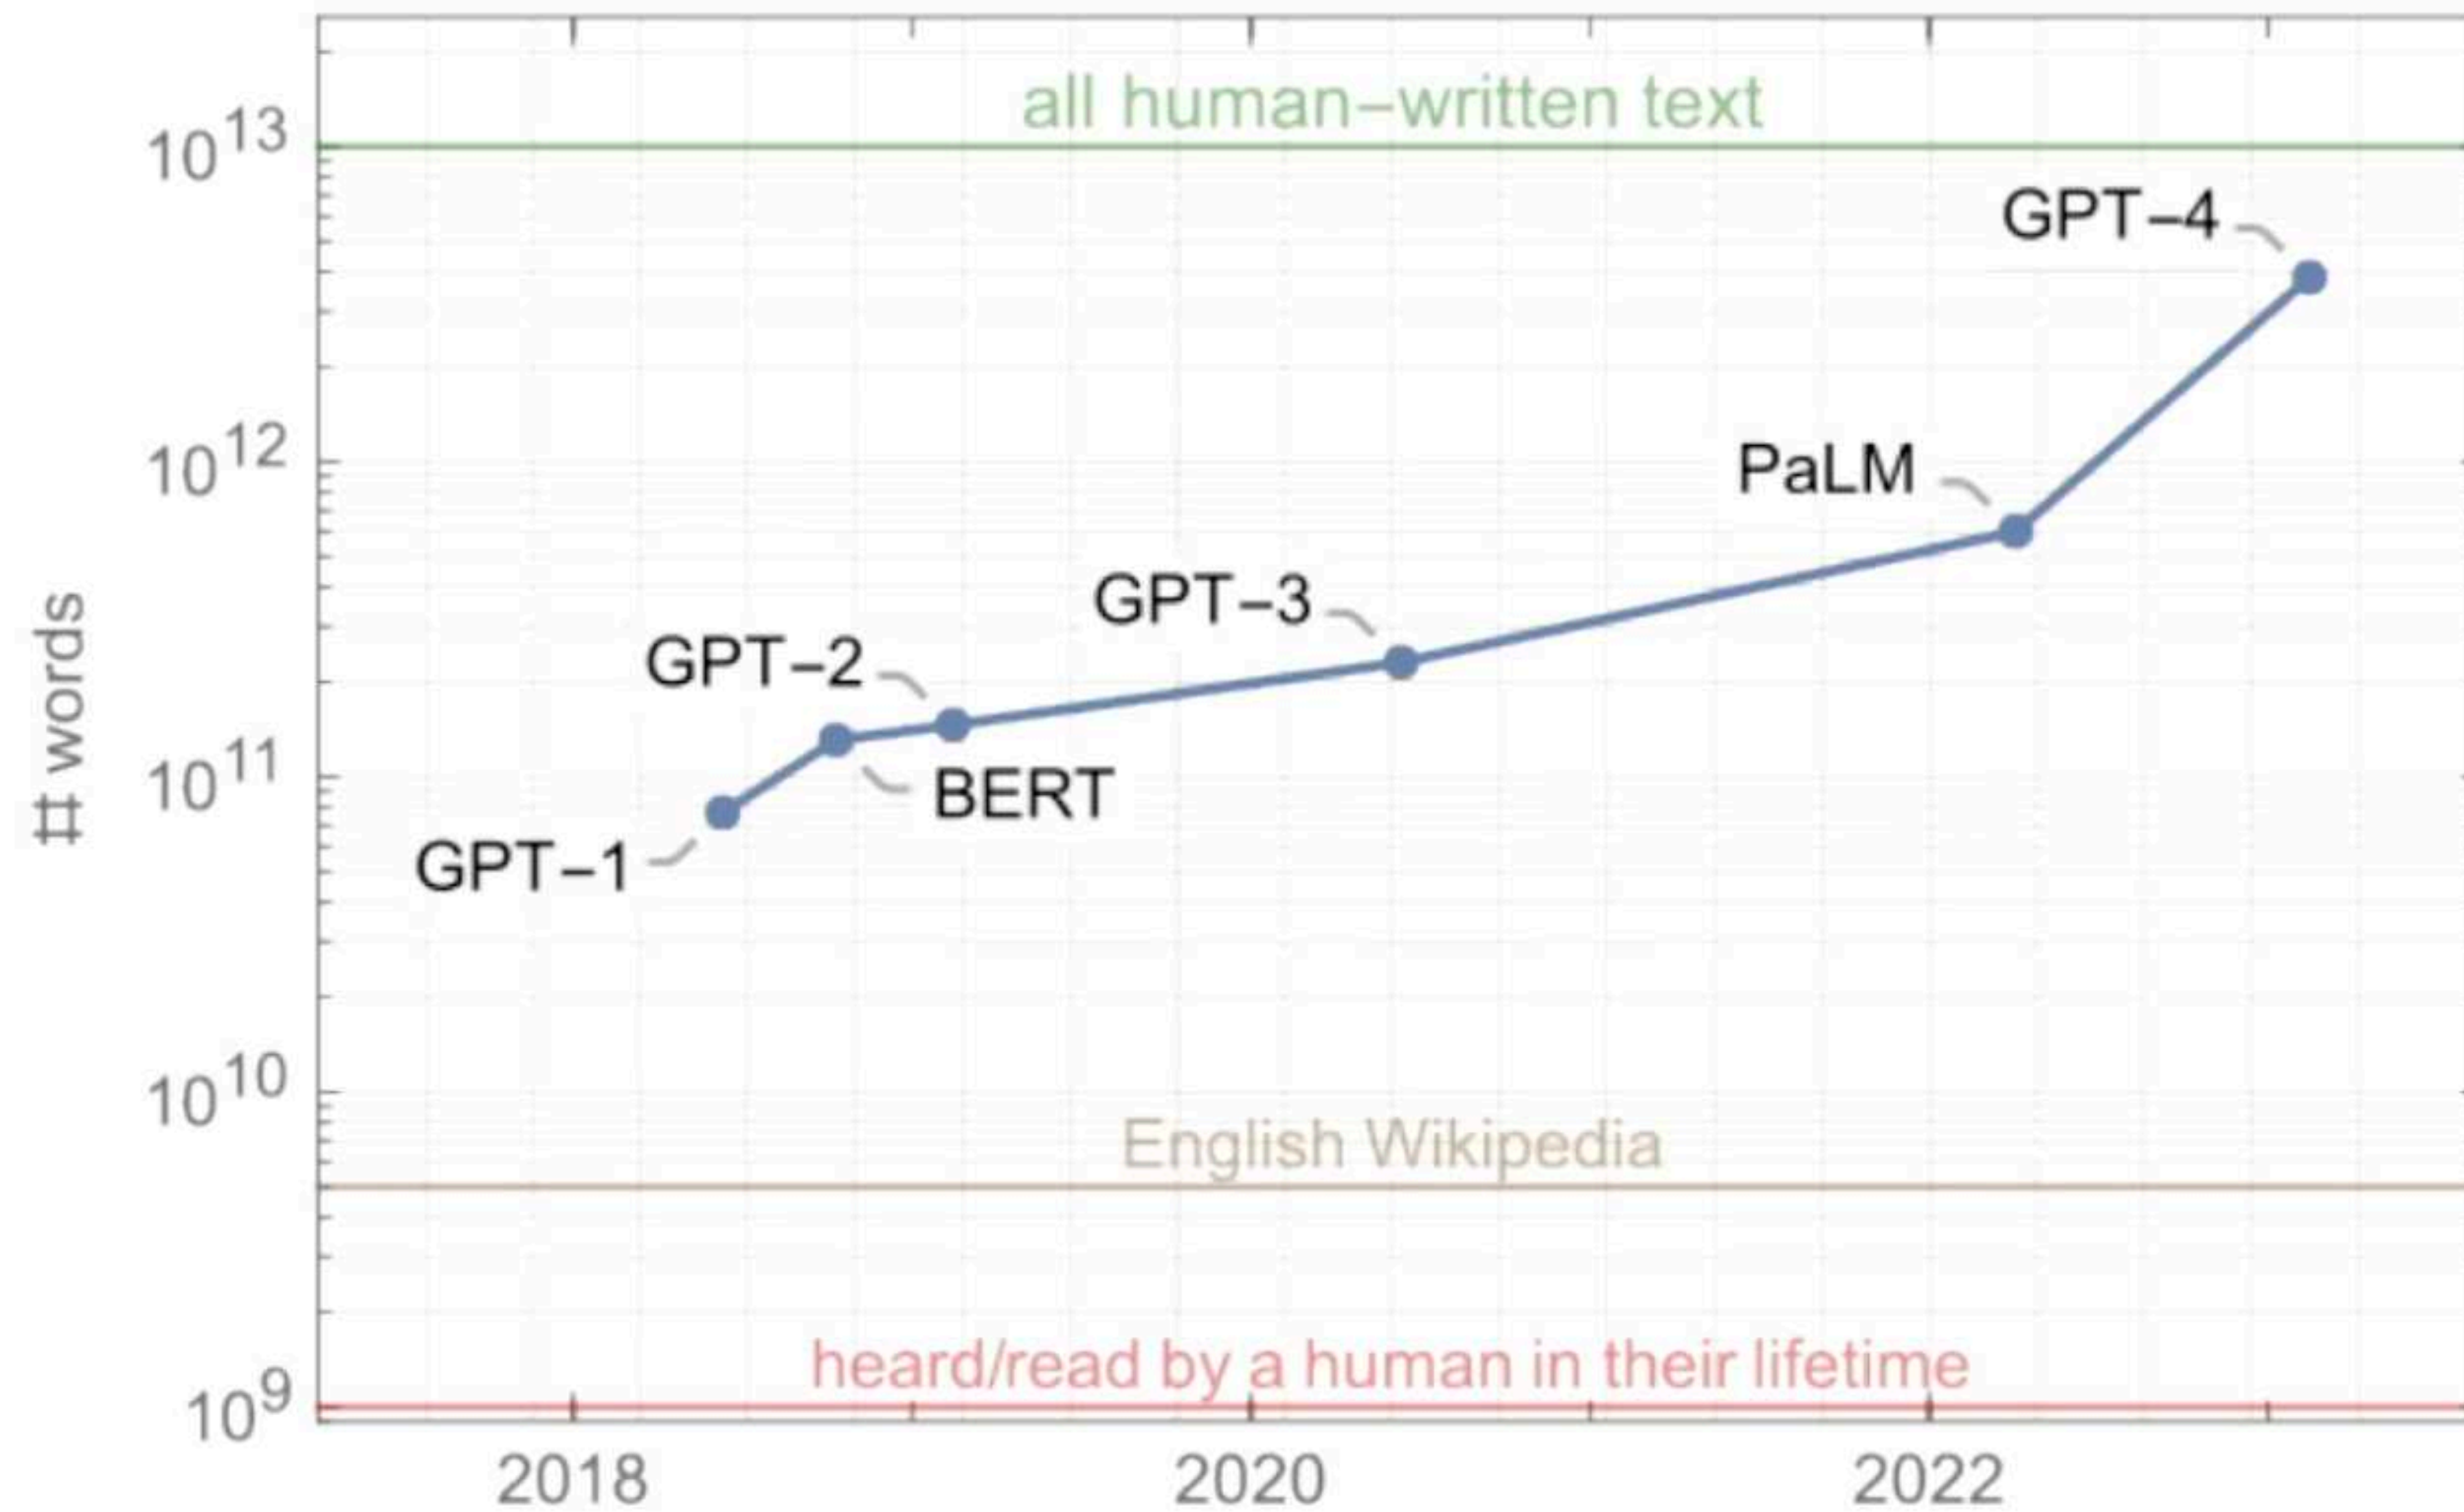
Llama

Open for all!

Pre-Training needs *a lot* of data.
Like the whole internet's worth of data.



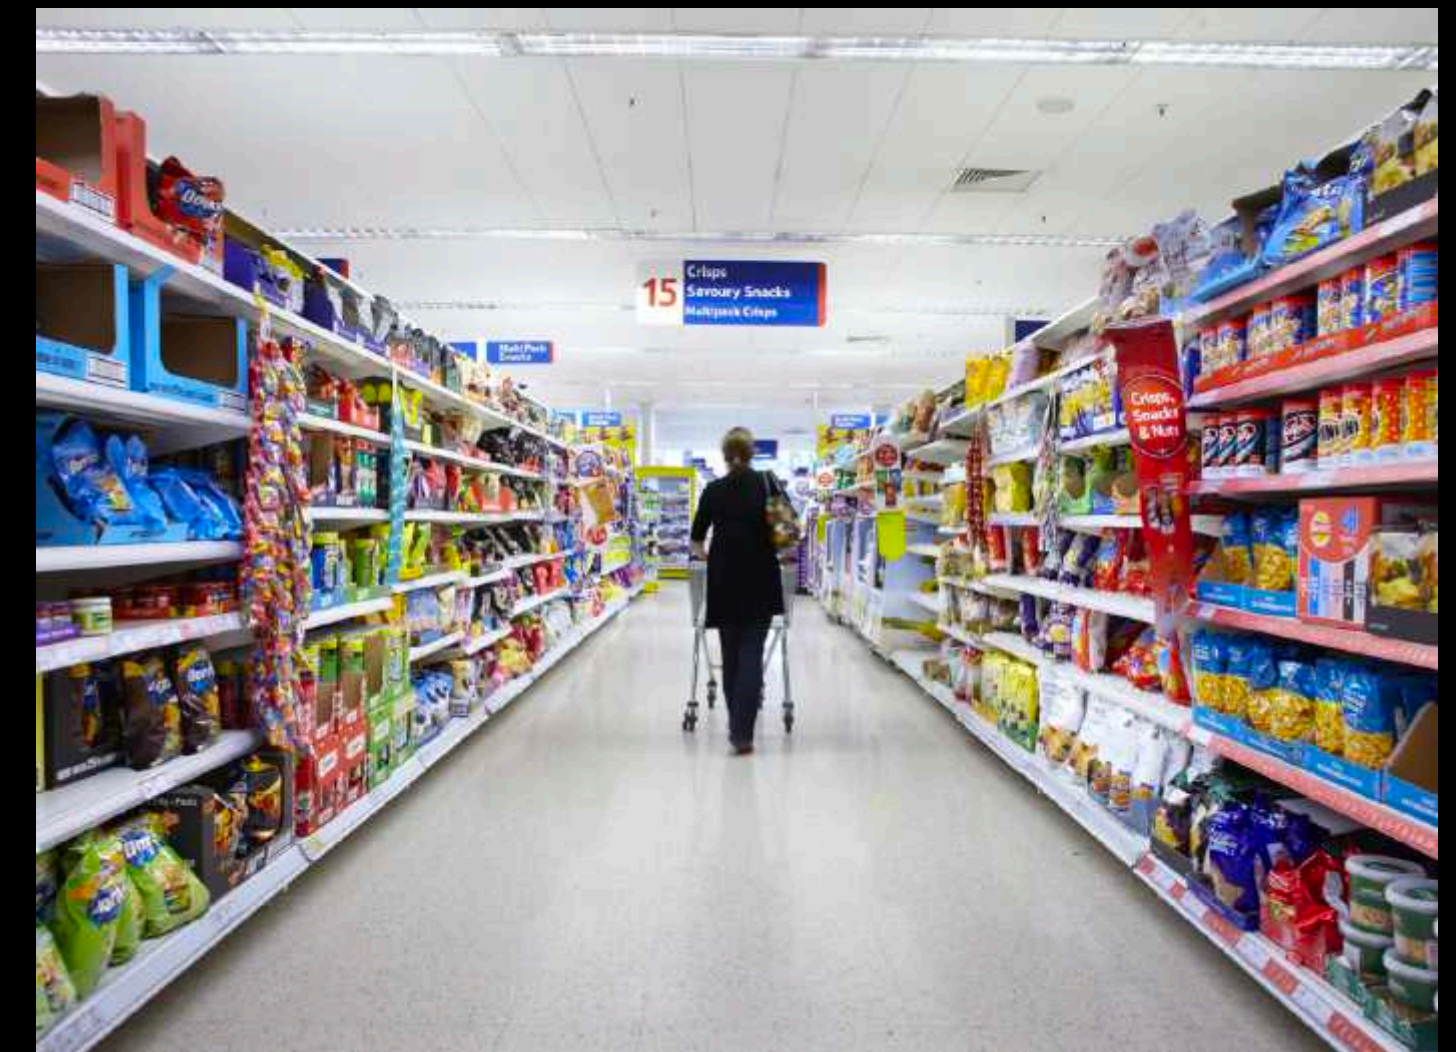
Number of words processed by LLMs during their training



We're running out of data. Where can we go?

We're running out of data. Where can we go?

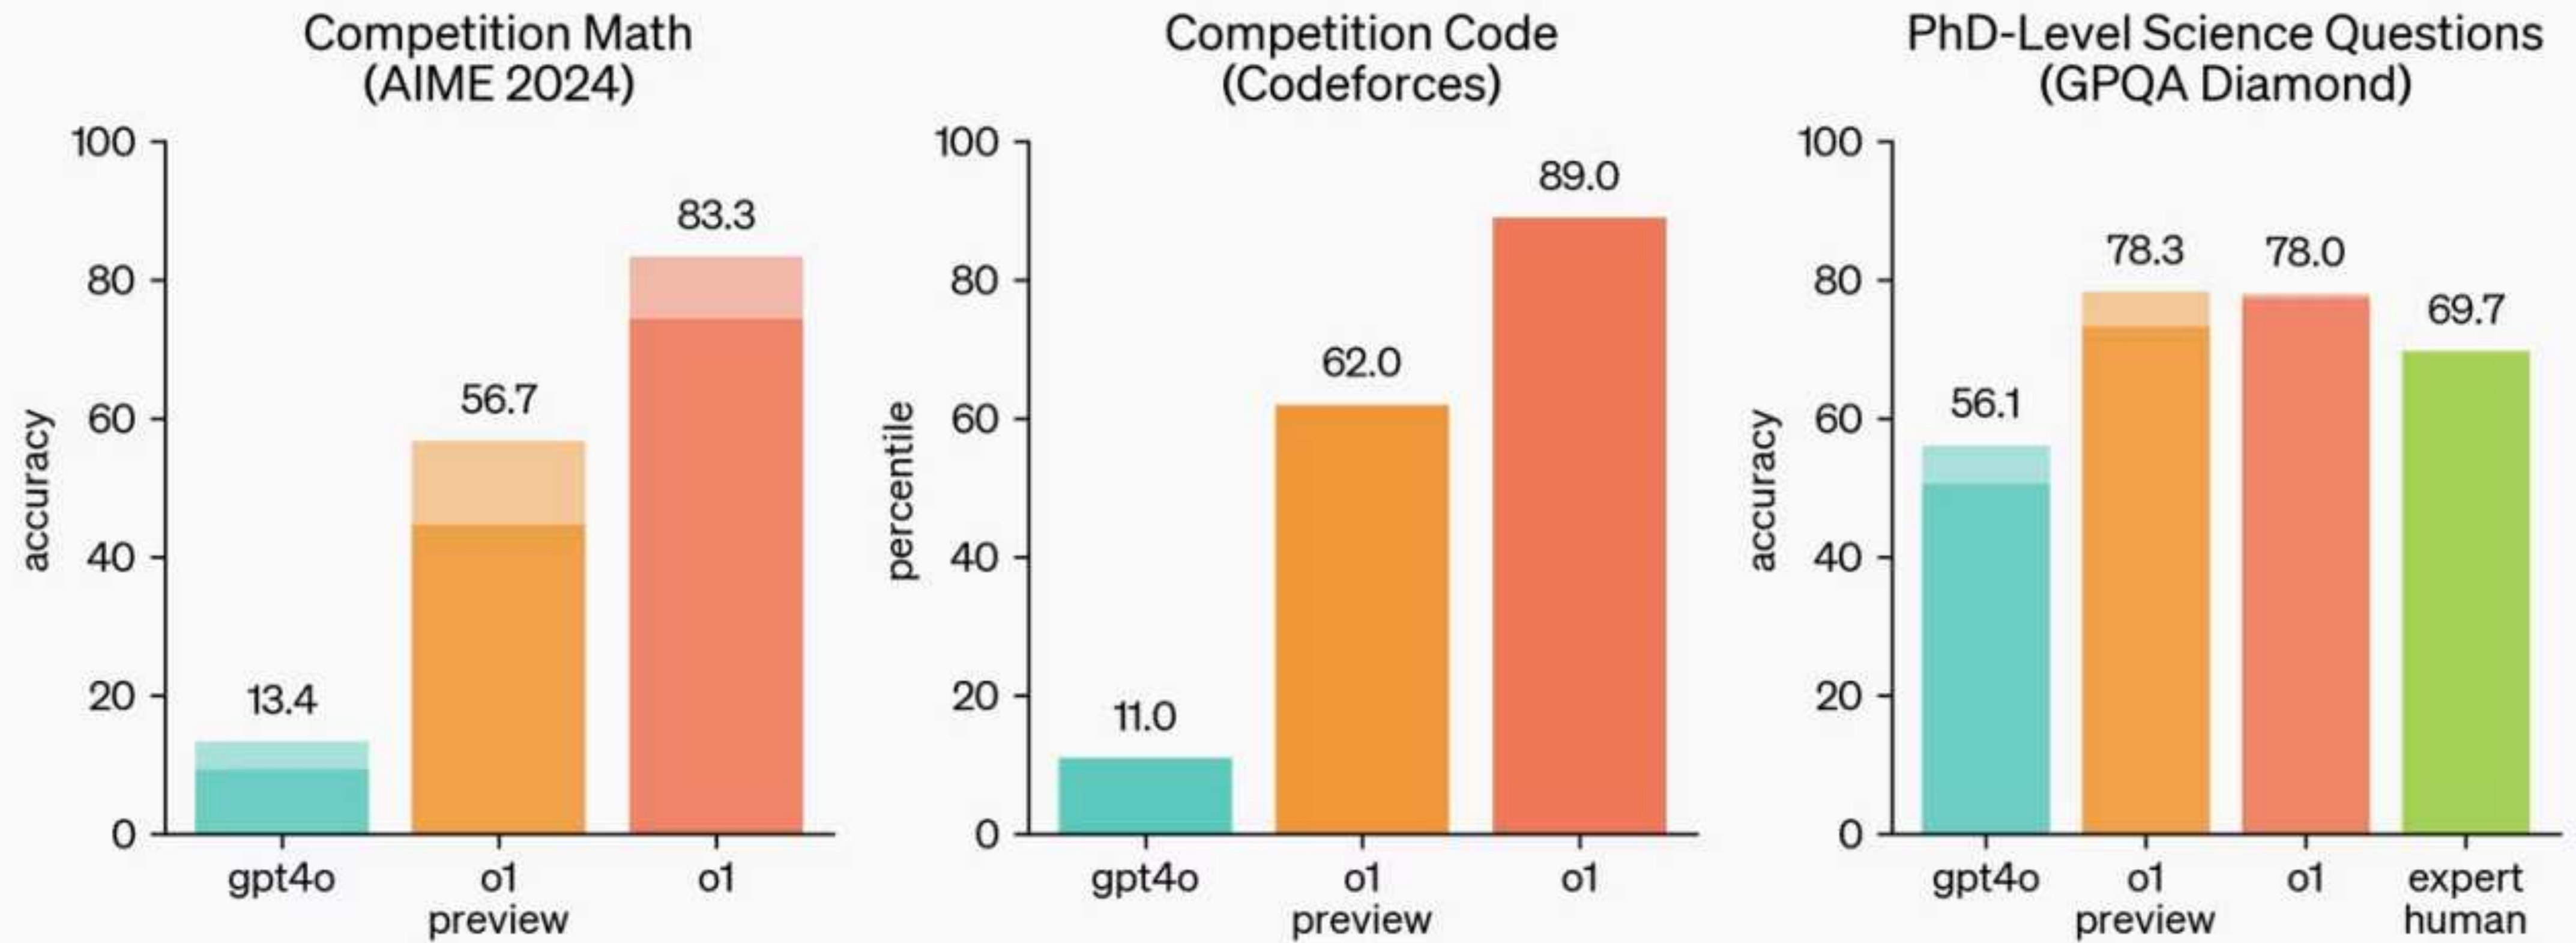
Industry:



ChatGPT (GPT-3) vs GPT-4

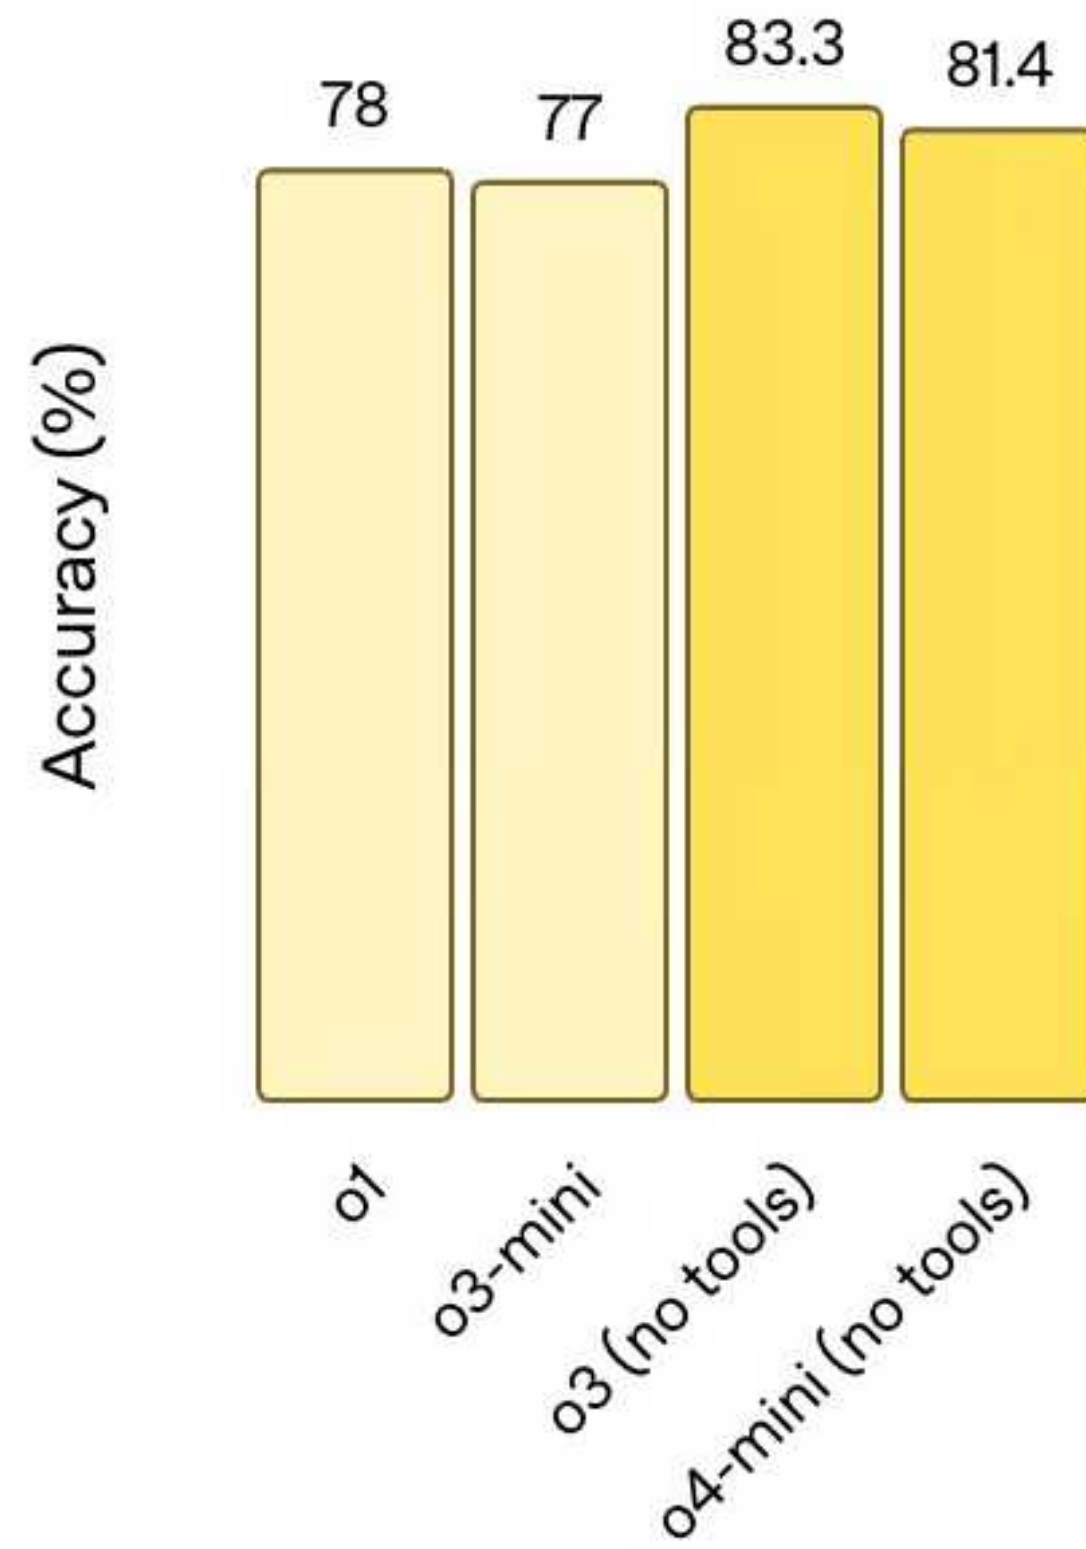
Test	ChatGPT	GPT-4
Uniform Bar Exam	10th percentile	90th percentile
SAT Reading & Writing	87th percentile	93rd percentile
SAT Math	70th percentile	89th percentile
GRE Verbal	63rd percentile	99th percentile
GRE Quantitative	25th percentile	80th percentile
GRE Writing	54th percentile	54th percentile
USA Biology Olympiad Semifinal Exam	31 st percentile	99th percentile (with vision)

GPT4o vs GPTo1

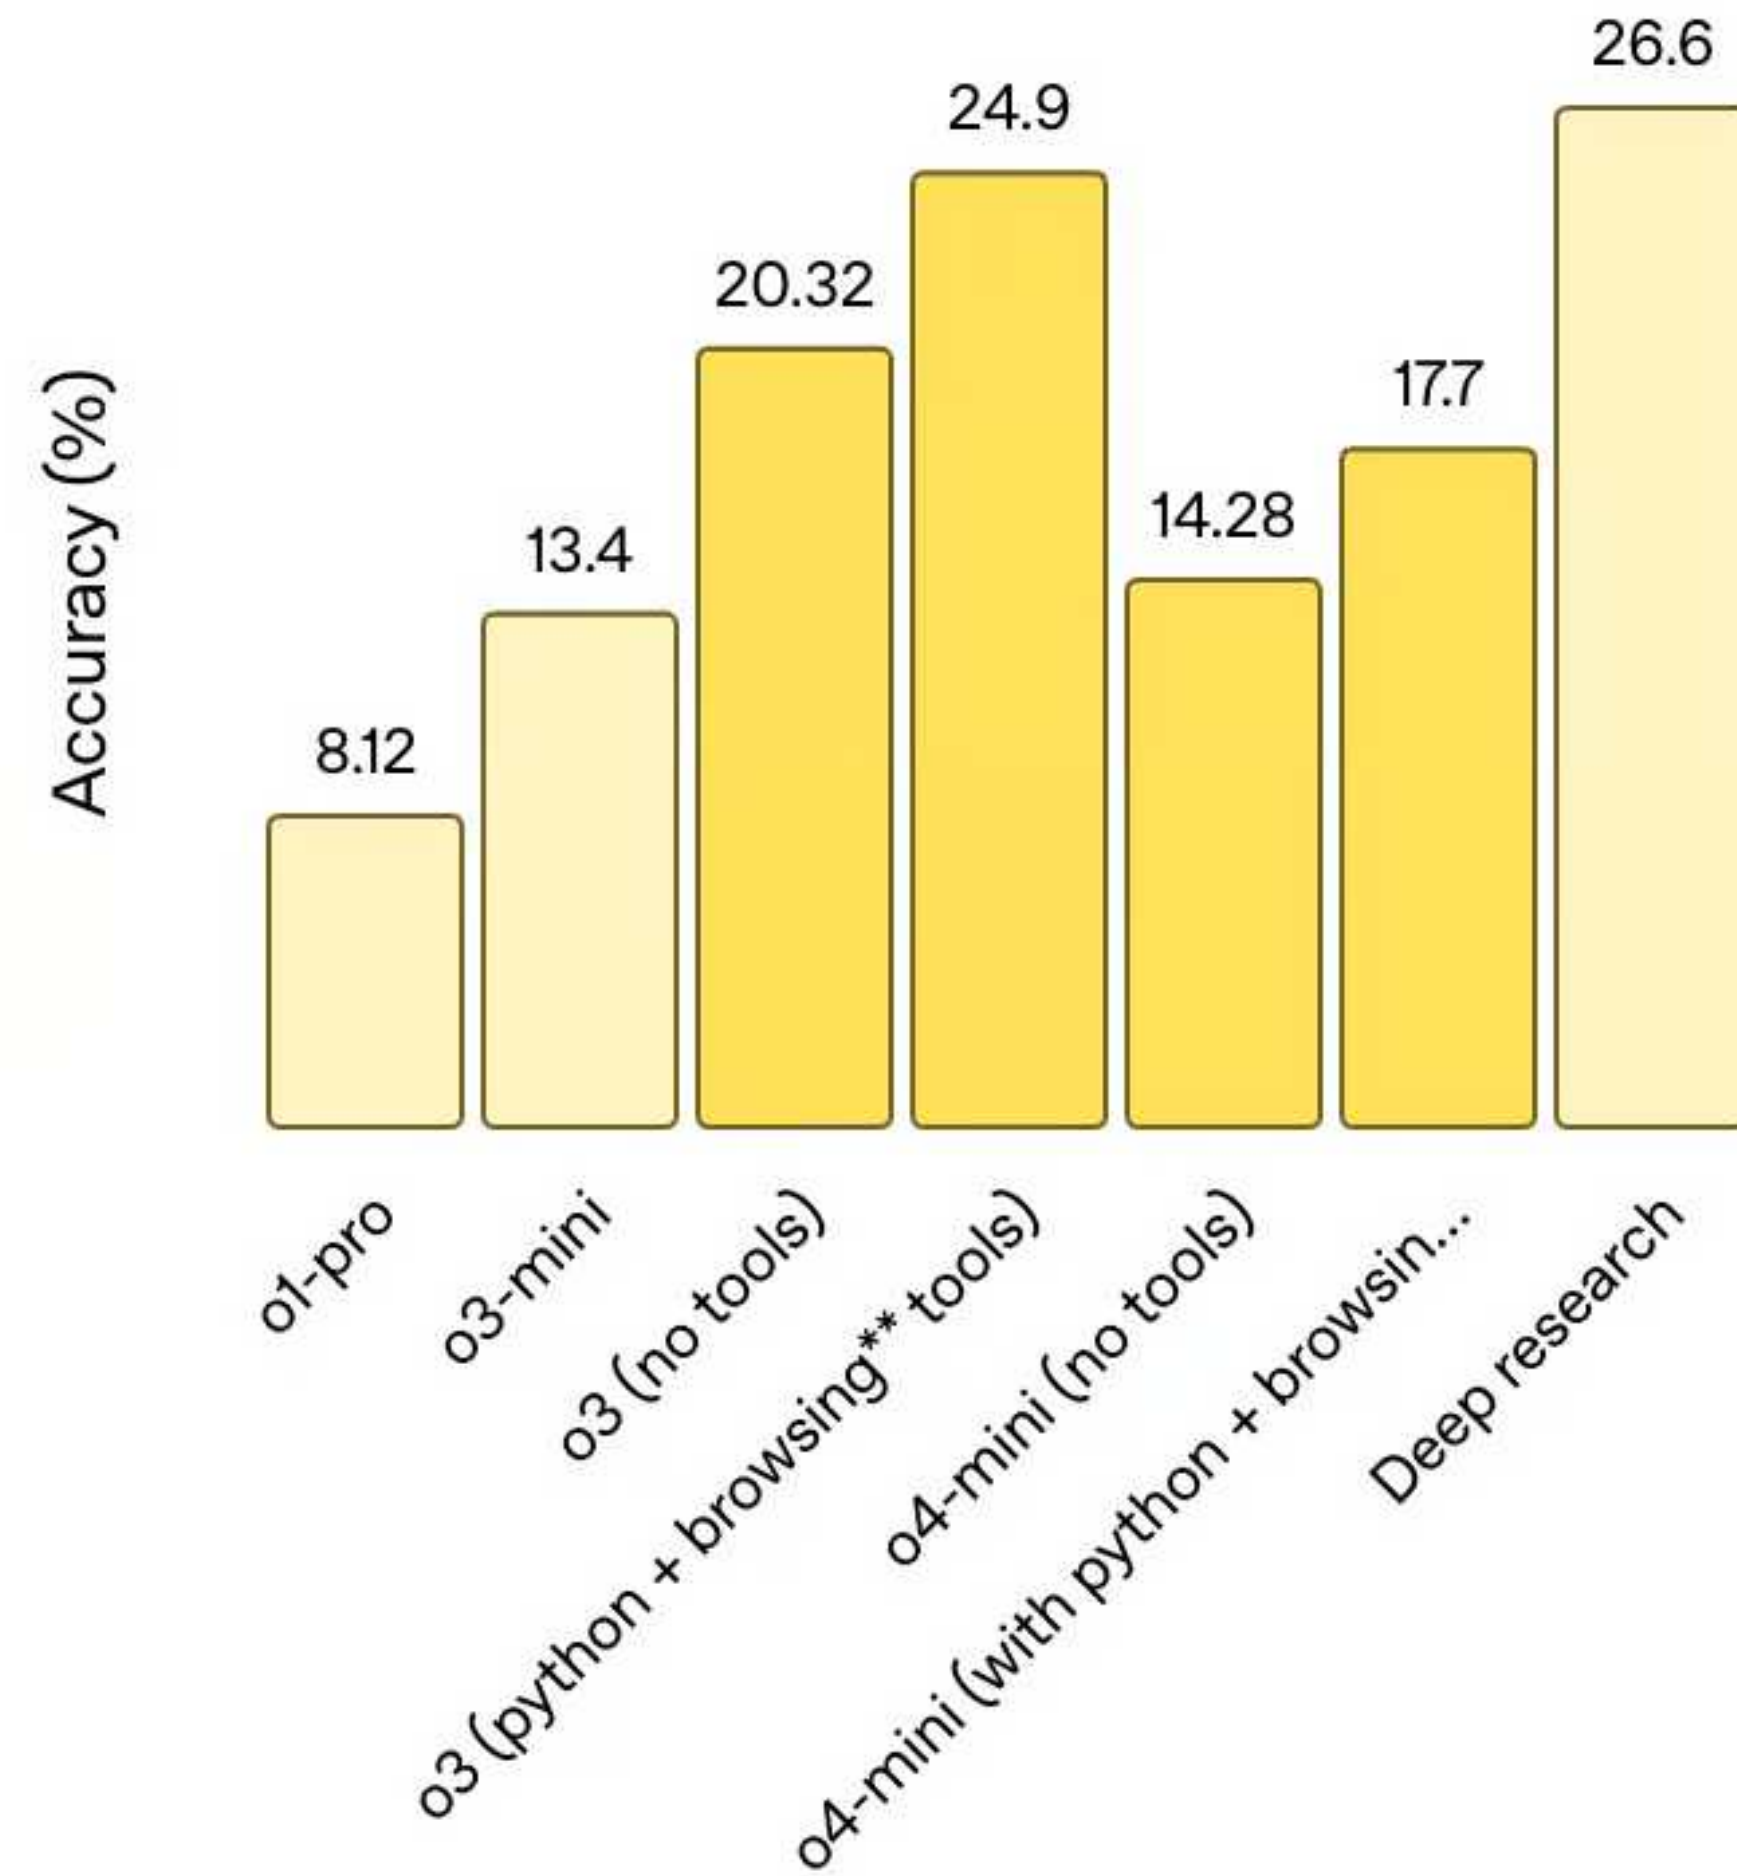


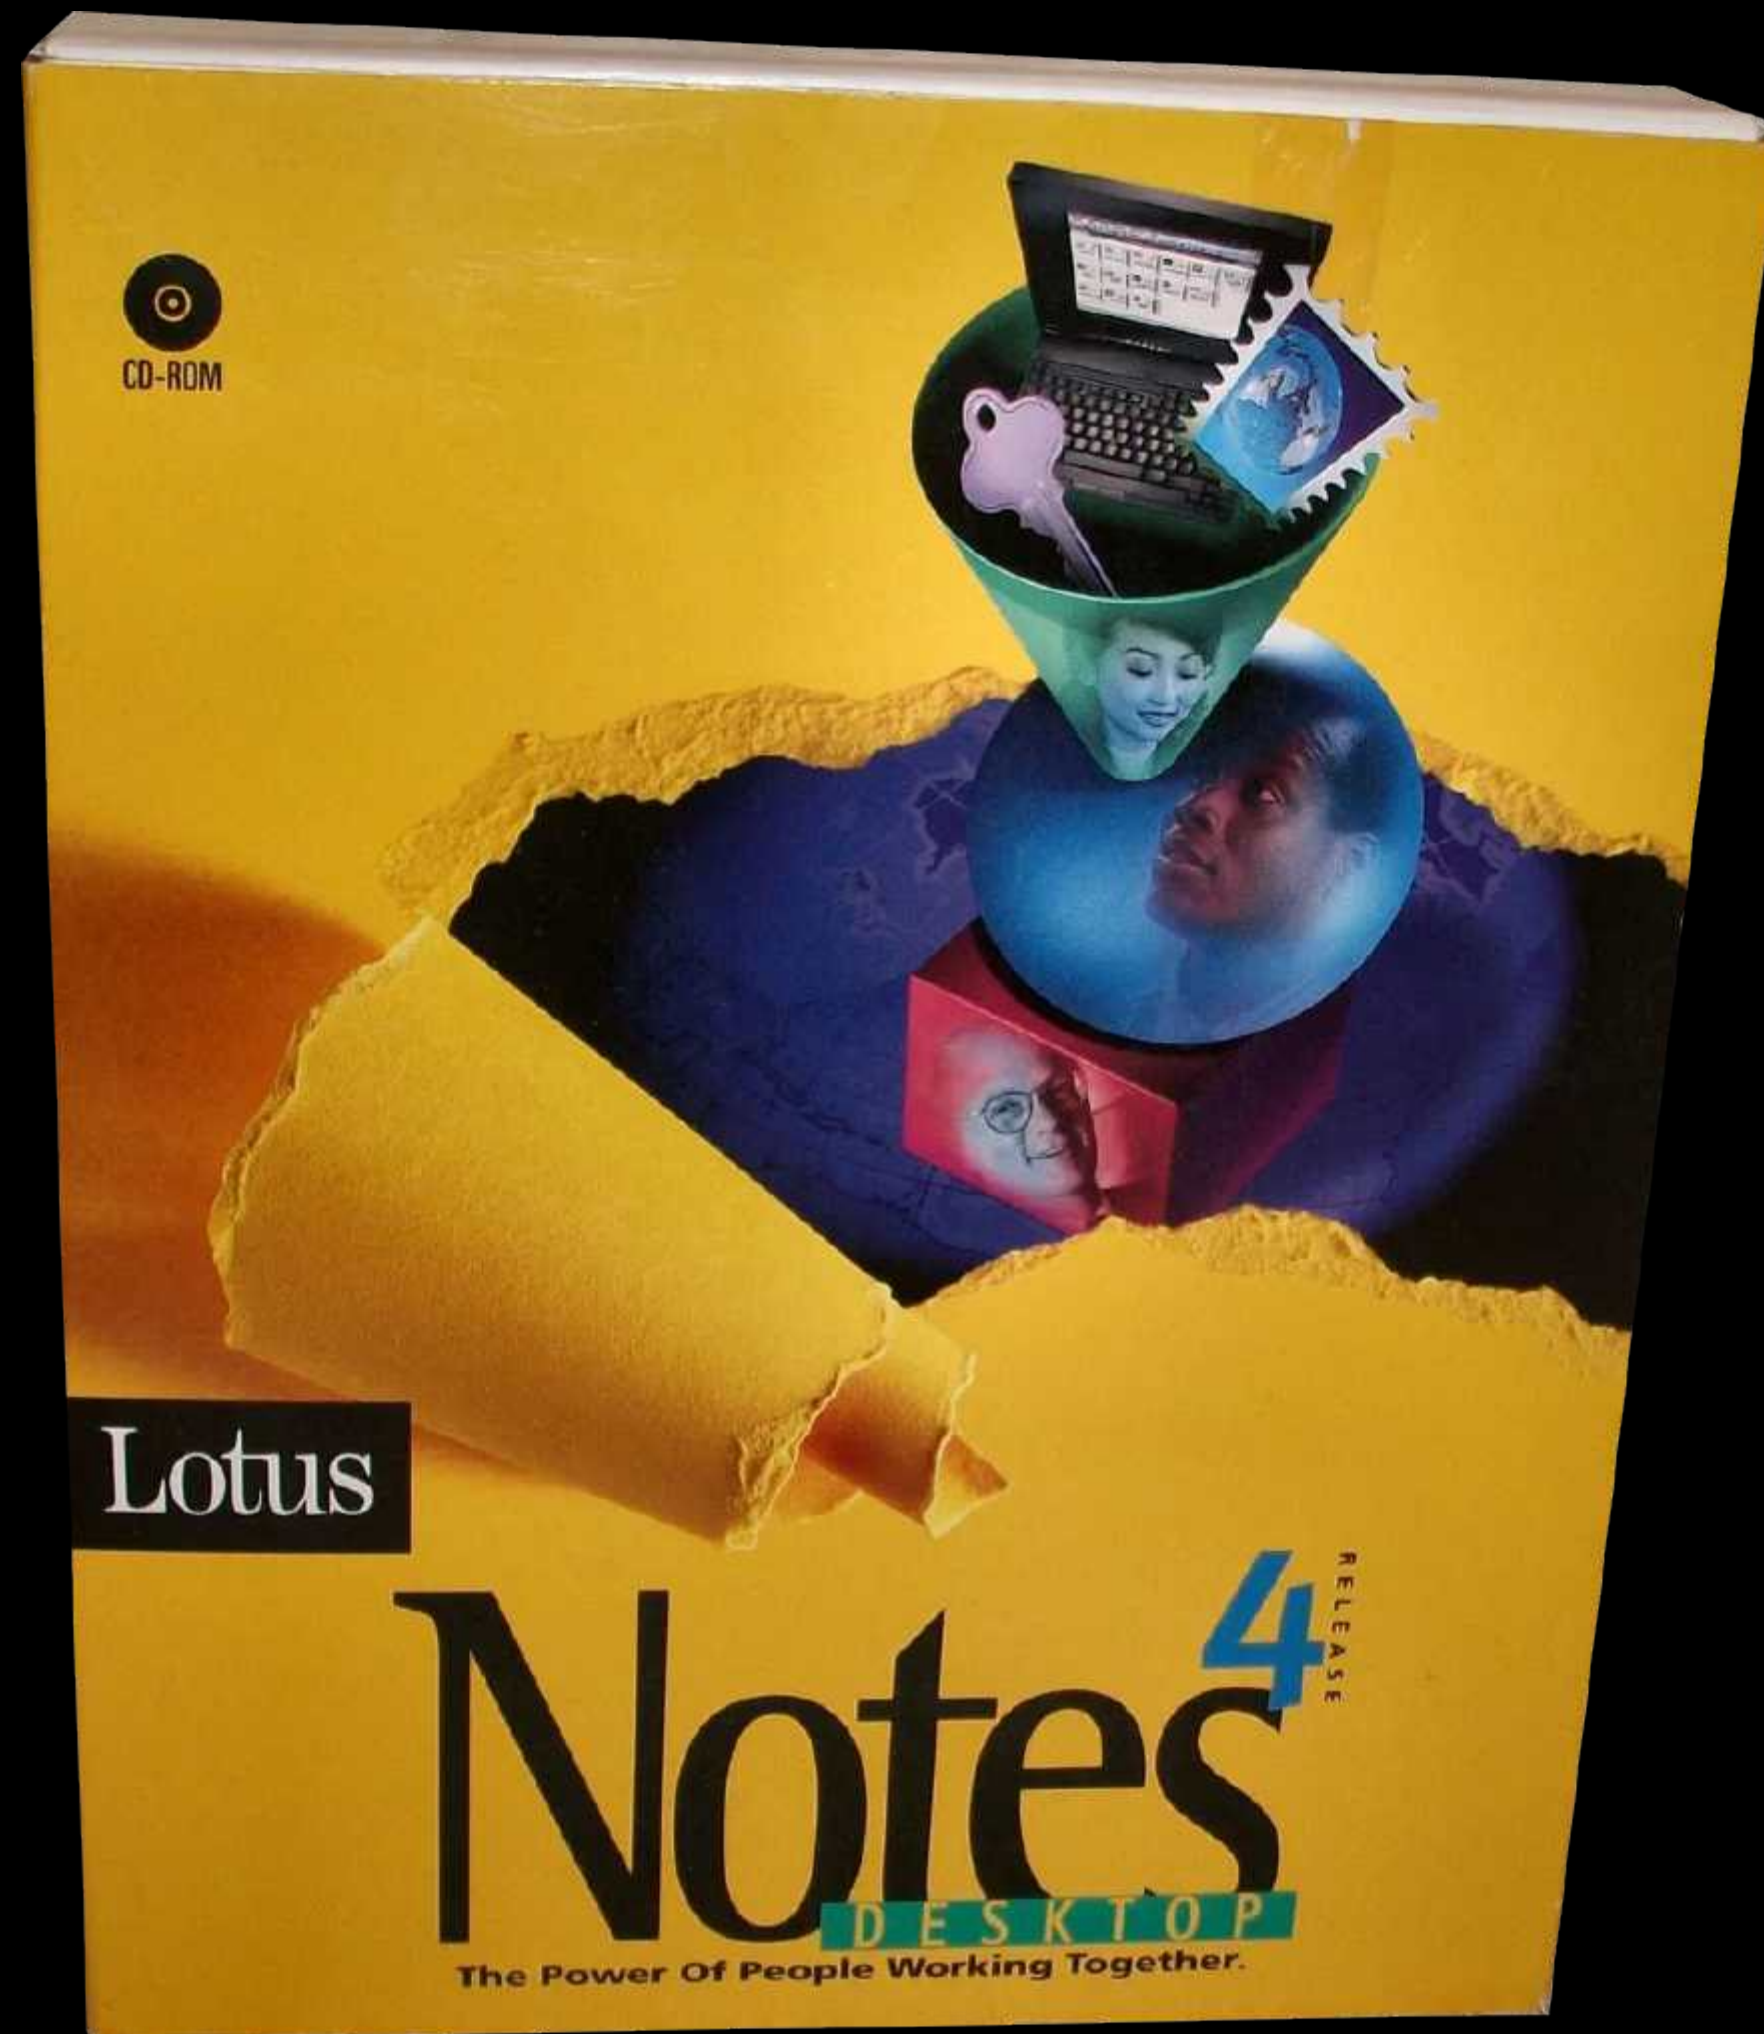
From GPTo1 to GPTo3

GPQA Diamond
PhD-Level Science Questions



Humanity's Last Exam
Expert-Level Questions Across Subjects





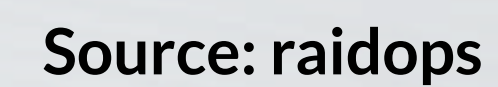
“Assume this is the
worst AI you’ll ever use.”

- Ethan Mollick

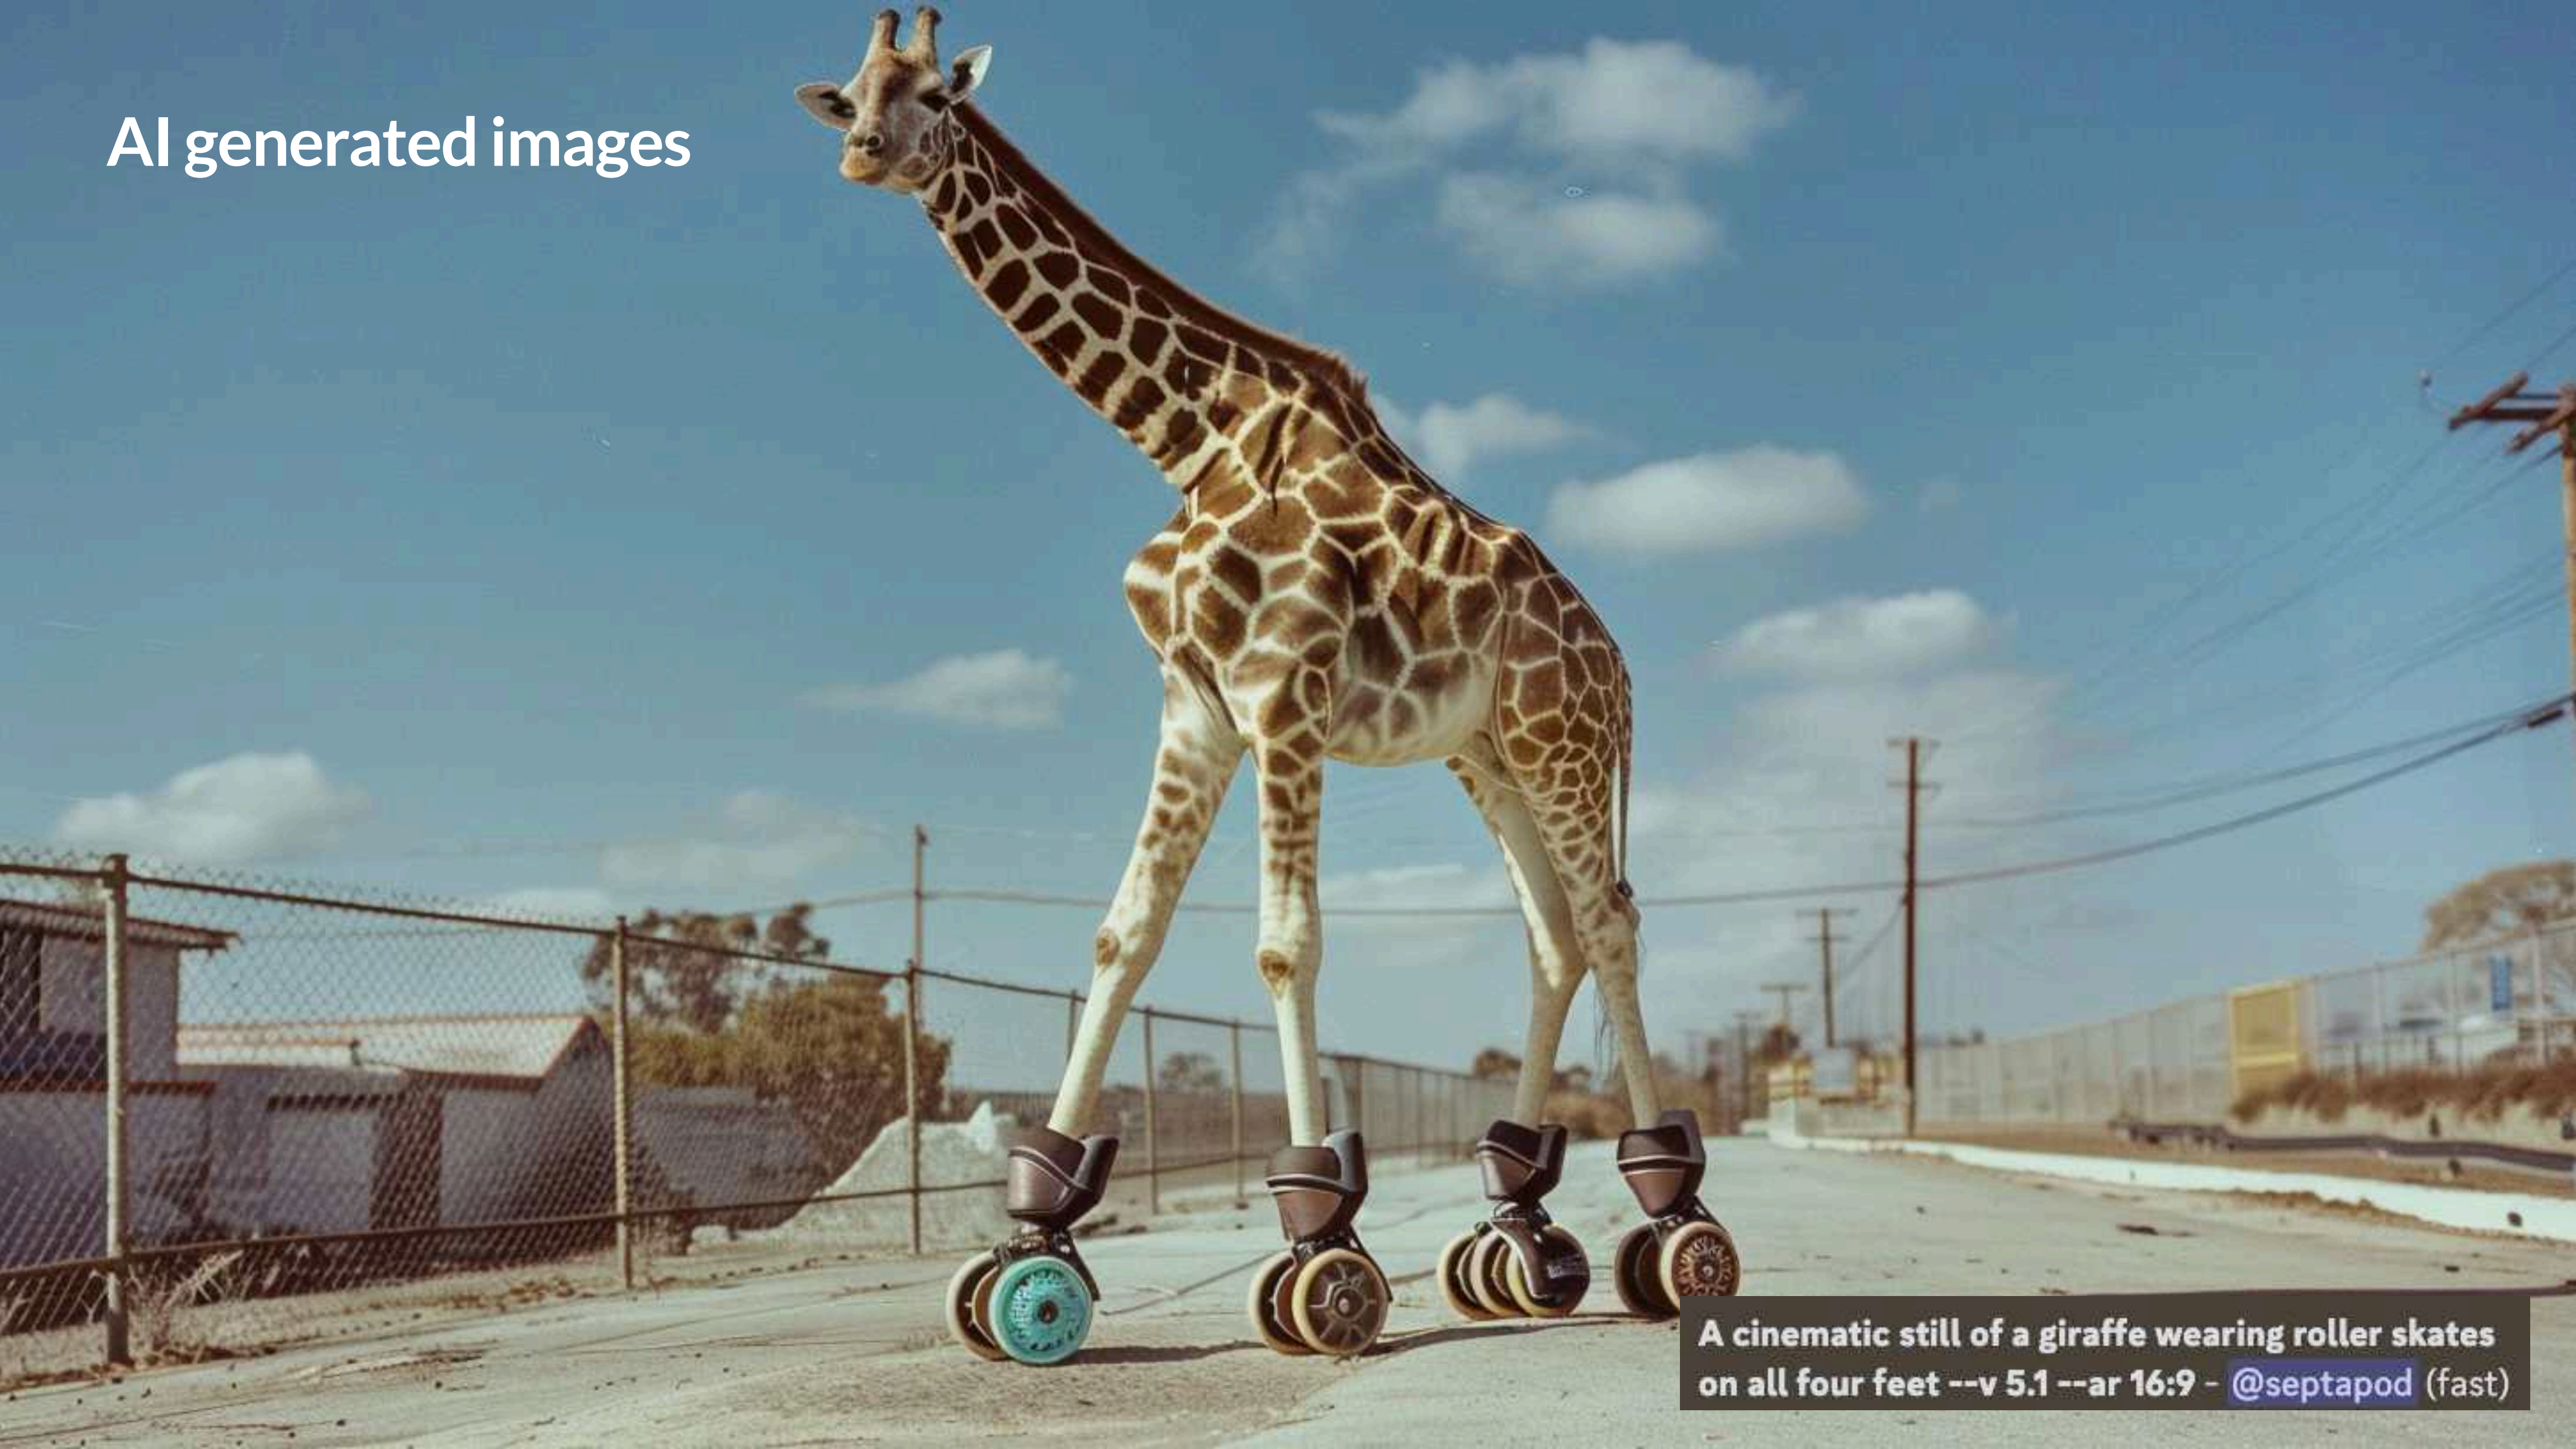
THE WIDER SHOT

Generative AI

A closely-cropped photograph of a lush oak tree covered in moss. --v 5 --ar 16:9 --seed 2769754162 - @septapod (Open on website for full



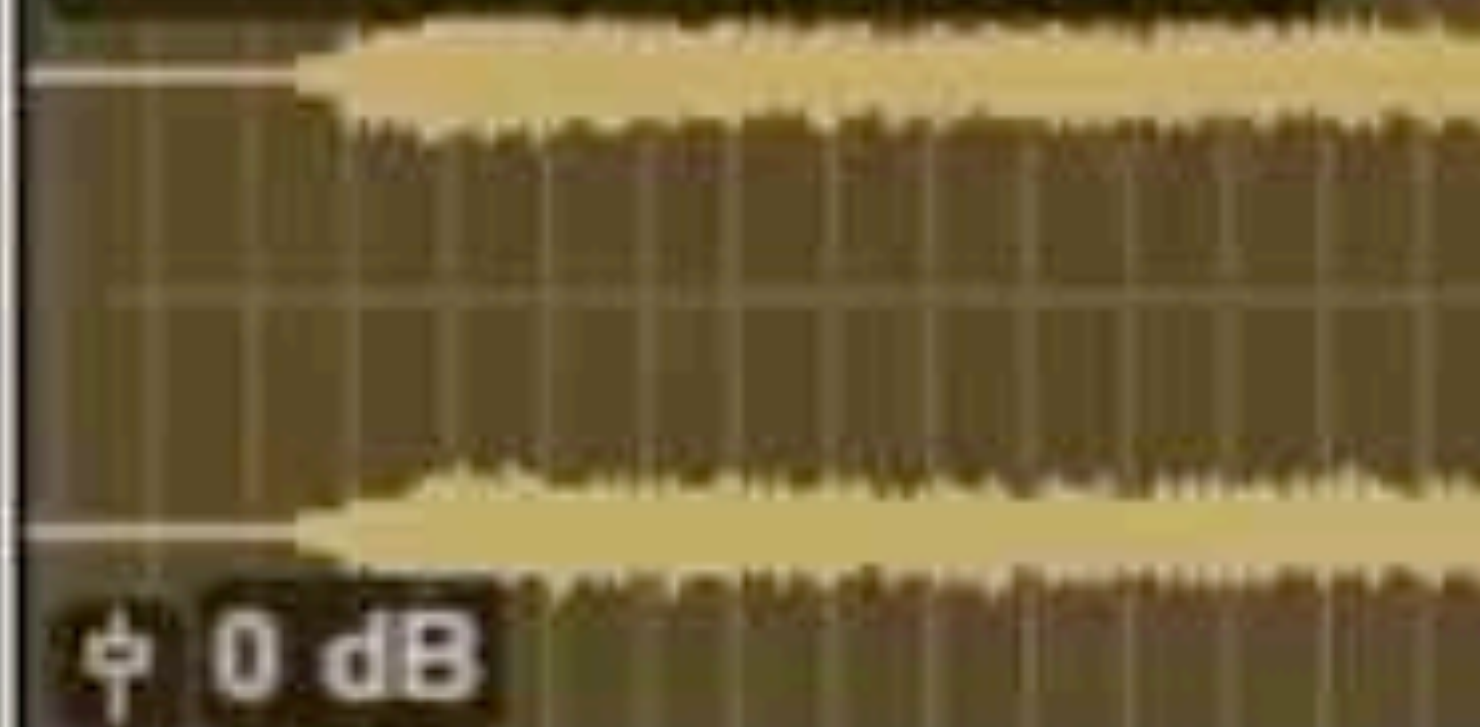
AI generated images



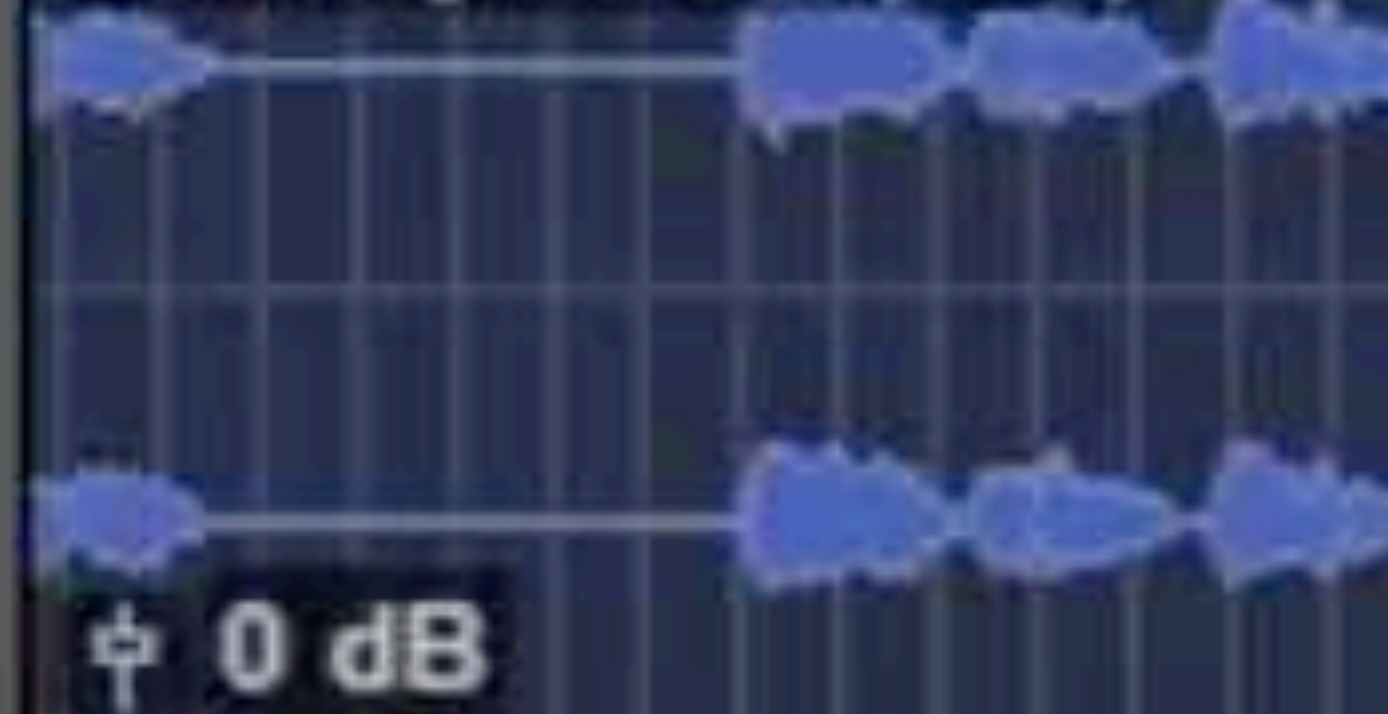
**A cinematic still of a giraffe wearing roller skates
on all four feet --v 5.1 --ar 16:9 - @septapod (fast)**



Folsom Prison Blues

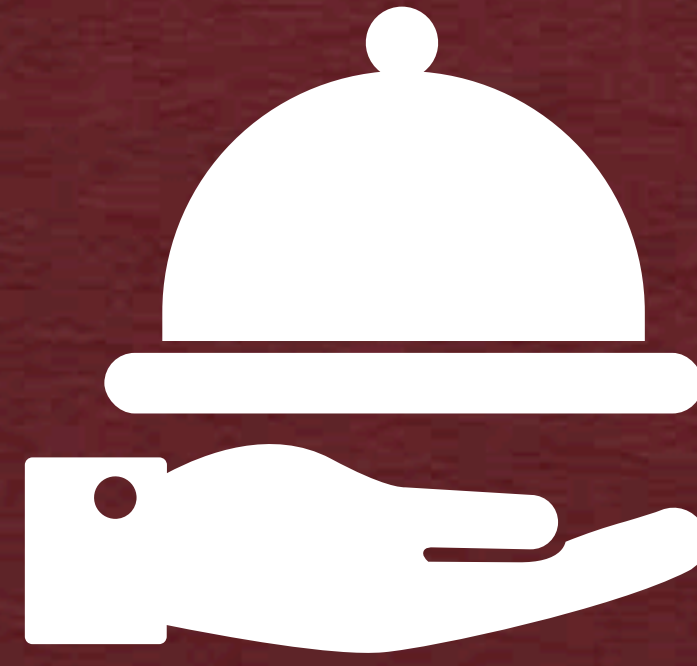


Johnny Cash (A.I.)





*How can we put it
to work?*



Member Experience & Ops

Member feedback sentiment analysis @ PSECU

SENTIMENT ANALYSIS



POSITIVE

"Great service for an affordable price.
We will definitely be booking again."



NEUTRAL

"Just booked two nights
at this hotel."



NEGATIVE

"Horrible services. The room
was dirty and unpleasant.
Not worth the money."

[Home](#) | [Meet Olive](#)

MEET OLIVE!

Olive is here to streamline your banking experience, taking our mission of member service to a whole new level.

Olive is getting promoted!

Since August 2023, Olive, GLCU's virtual assistant powered by artificial intelligence (AI), has been diligently making their banking experience smoother and more personalized than ever before. We are excited to announce that Olive has become even smarter and is now offering bilingual support in Spanish for members who call in to our customer service line.

Olive can also use Voice Authentication to verify your identity. This technology not only offers an extra layer of security but also simplifies authentication processes. You can use Voice Authentication to retrieve basic account information and view recent transactions.

Other key benefits of Voice Authentication include:

Chat with Olive



What is your mortgage application process like?
10:26 AM ✓

I understand you're asking about mortgages. Click below to view an existing mortgage. Or, if you don't have a mortgage, click to learn more about the benefits, review the process, and use our convenient online calculator to see what your payments might be if you apply.

[View Existing Mortgage](#)

[Learn More](#) 

10:26 AM

Type a message...



MEET

OLIVE!

Improved call containment rate from
25% to 60-75%

Olive is getting promoted!

Since August 2023, Olive, GLCU's virtual assistant powered by artificial intelligence (AI), has been diligently making their banking experience smoother and more personalized than ever before. We are excited to announce that Olive has become even smarter and is now offering bilingual support in Spanish for members who call in to our contact center.

Olive can also use Voice Authentication to verify your identity. This technology not only offers an extra layer of security but also simplifies authentication processes. You can use Voice Authentication to retrieve basic account information and review recent transactions.

Other key benefits of Voice Authentication include:

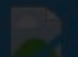
Chat with Olive



What is your mortgage application process like?
10:26 AM ✓

Understand you're asking about mortgage applications. Click below to view an existing mortgage. Or, if you don't have a mortgage, click to learn more about the benefits, review the process, and use our convenient online calculator to see what your payments might be if you apply.

[View Existing Mortgage](#)

[Learn More](#) 

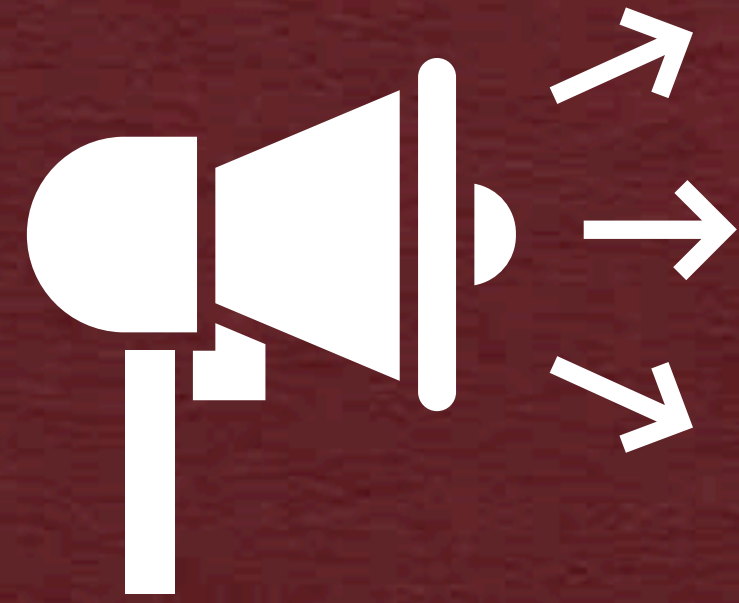
10:26 AM

Type a message...



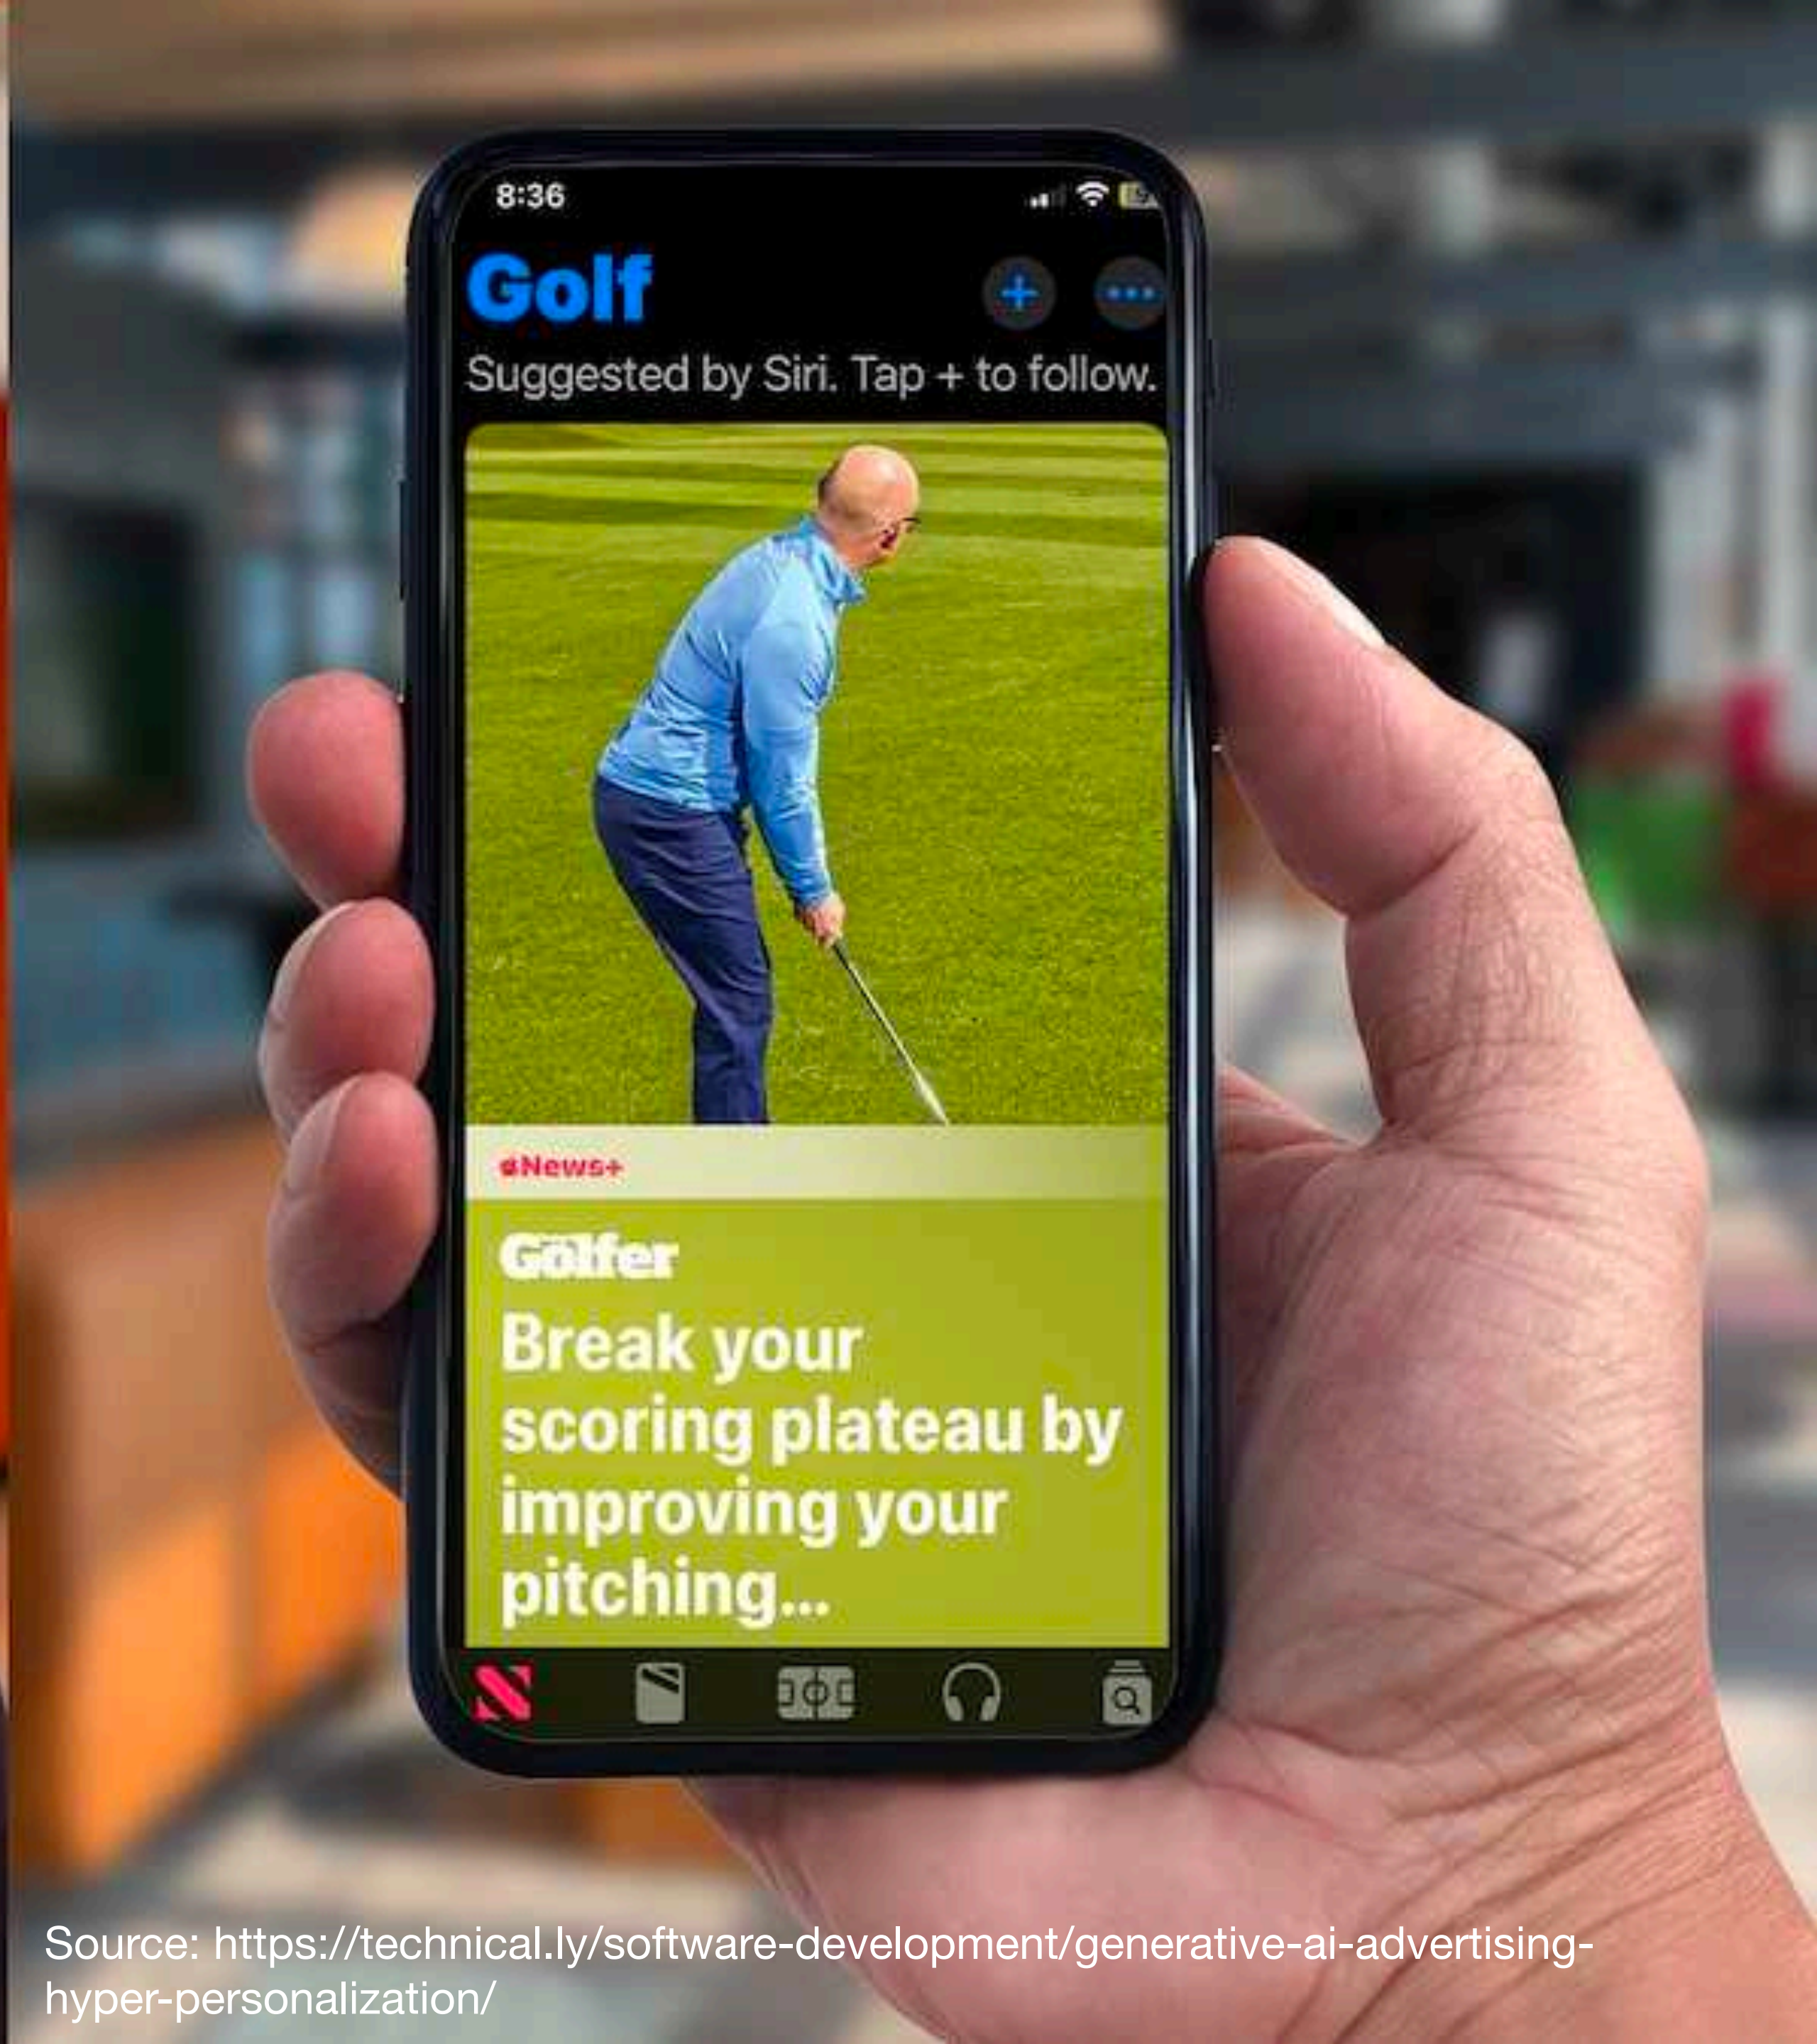
Traditionally, agents at Dutch Bank ABN Amro have taken notes during a customer call to produce a summary afterwards. The bank is now starting to use LLMs to create these summaries, with the agents simply checking they are accurate afterwards.



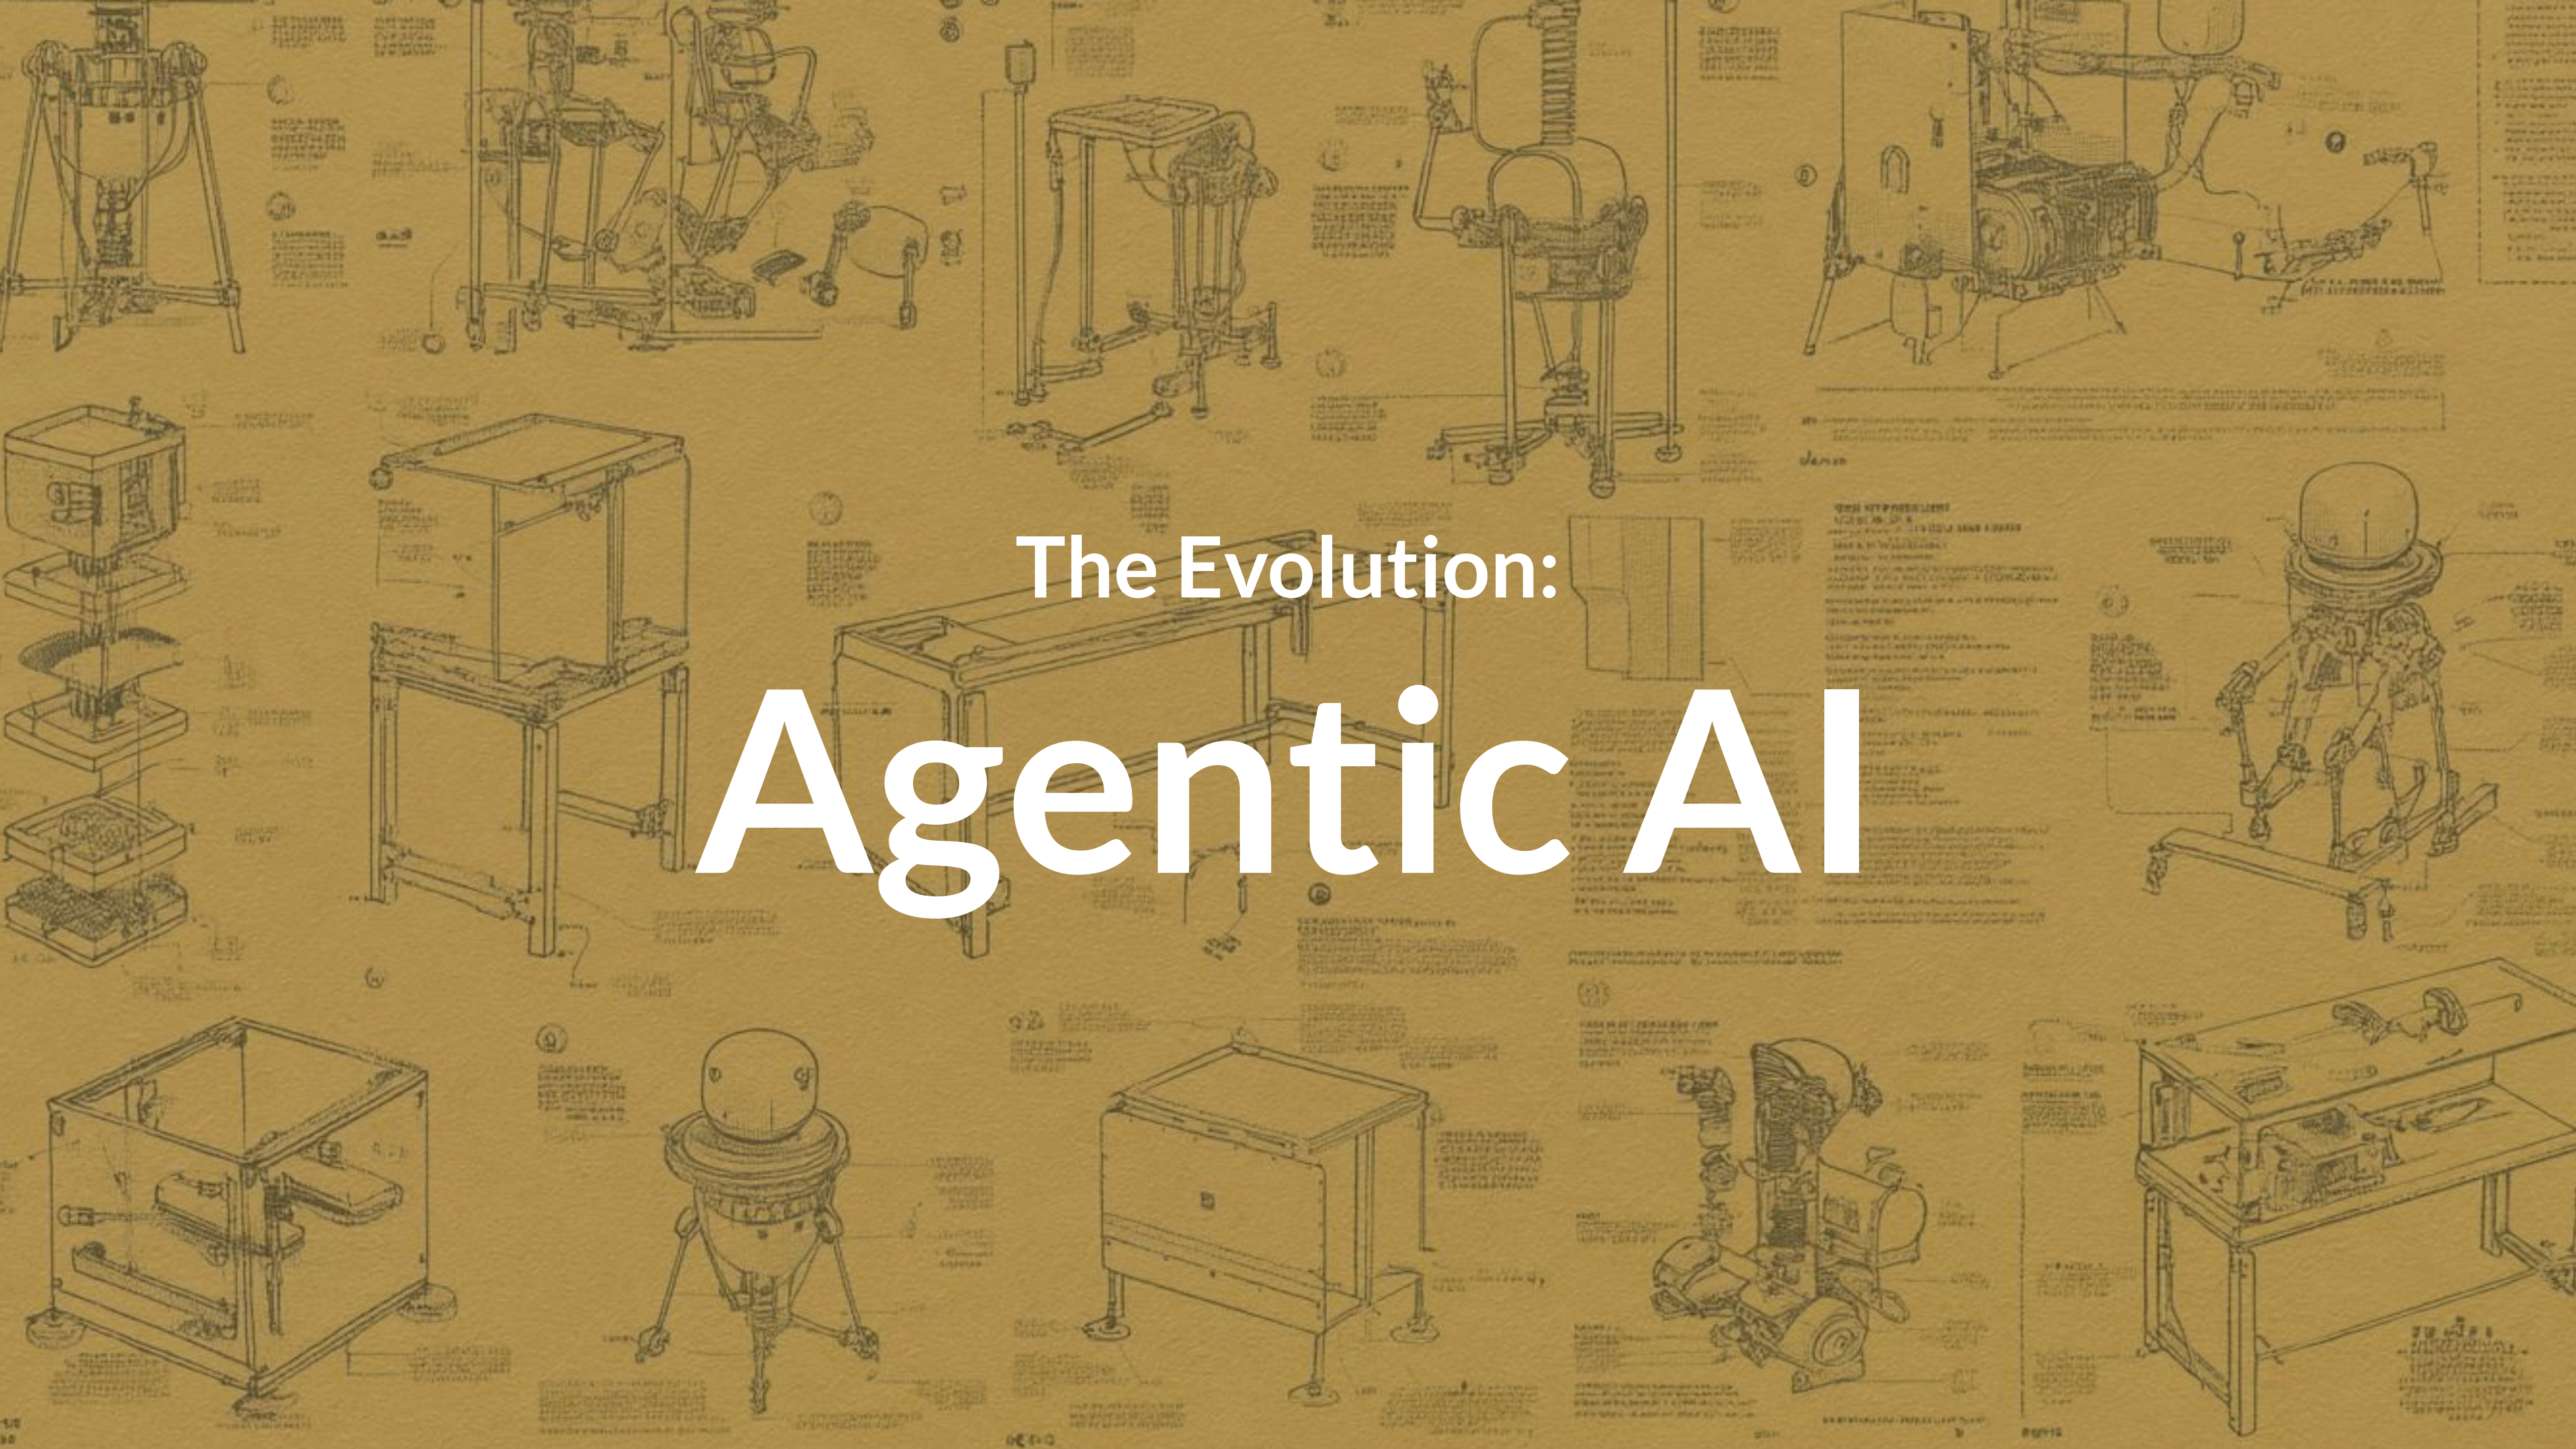


Marketing & Sales

Hyper-personalization



Source: <https://technical.ly/software-development/generative-ai-advertising-hyper-personalization/>



The Evolution: Artificial AI

From:

**AI that
responds**



To:

**AI that
executes**

An AI agent is software
that autonomously
reasons, decides, and
acts to achieve a goal.

reasons

Investigate if and how time dilation happens at speeds and altitudes humans commonly experience and provide examples




pro

Deep Research



Investigate if and how time dilation happens at speeds and altitudes humans commonly experience and provide examples

 Perplexity Research 

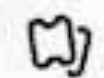
 Researching


decides


Operator


What can I help you do?



 Dining & Events

 Delivery

 Local Services

 Shopping

 Travel

 >

Find a highly rated sushi restaurant and snag us a spot for 2 tonight



OpenTable

Reserve 2 tickets for the next concert at The Spher



StubHub

Purchase the most affordable passes to the Miami Grand Prix



StubHub

Reserve a table for 8 at a seafood restaurant next Friday at 6PM

Book a table for 2 at a romantic French bistro tomorrow night at 8PM

Book a spot at a restaurant with a great happy hour for next Wed

acts

+ Create Repl

Create Python

Create HTML, CSS, JS

Replit Agent

Limited early access



Meet the Replit Agent

Meet the
Replit Agent



replit

1 Craft the perfect plan

Describe your goals and the agent will develop a tailored plan. It's currently best at building web apps.

2 Keep tabs on progress

Monitor progress in real-time. Help the agent by providing API keys, feedback, and guidance.

3 Test and deploy

Review the final product and deploy your application directly on Replit in a few clicks!

Describe what you want to make



Start building

Local landmarks map

Stock analysis

Waitlist website

View more

Recent Repls

All Repls

Search

Files

.git

generated-icon.png

Config files

replit.nix

Tools

Recent

Progress

Creating main.py

main.py

1

from app import app

2

3

if __name__ == "__main__":

4

app.run(host="0.0.0.0", port=5000)

5

Build the initial prototype

1 / 2

Create a version of pong, but where the "ball" is a happy face.

Just now • Read

Absolutely! Let me propose what we'll build for you.

I'm working on a detailed plan for a Pong-like game where the ball is replaced with a happy face emoji. The initial prototype will include basic game mechanics such as paddle movement, ball (happy face) bouncing, and scoring. We'll be using Flask for the backend and vanilla JavaScript with HTML5 Canvas for the frontend. Let me know if this sounds good, and I'll proceed with the implementation.

• Build the initial prototype

• Implement different difficulty levels (e.g., paddle size, ball speed)

✓ BrentDixon approved the plan

Build the initial prototype

Start building

Player 1: 0

Player 2: 0

Happy Face Pong



Start Game





Joining for coffee at a cafe

KM: ...

AC: ...

[Abigail]: Hey Klaus, mind if I join you for coffee?
[Klaus]: Not at all, Abigail. How are you?

Arriving at school

AK: ...

Joining a walk in the park

SM: ...

Finishing a morning routine

JM: ...

Agents can decide and act without
needing human oversight.

Sharing news with colleagues

JL: ...

TM: ...

[John]: Hey, have you heard anything new about the upcoming mayoral election?
[Tom]: No, not really. Do you know who is running?

Generative Agents

Interactive Simulacra of Human Behavior

Joon Sung Park • Joseph C. O'Brien • Carrie J. Cai
Meredith Ringel Morris • Percy Liang • Michael S. Bernstein

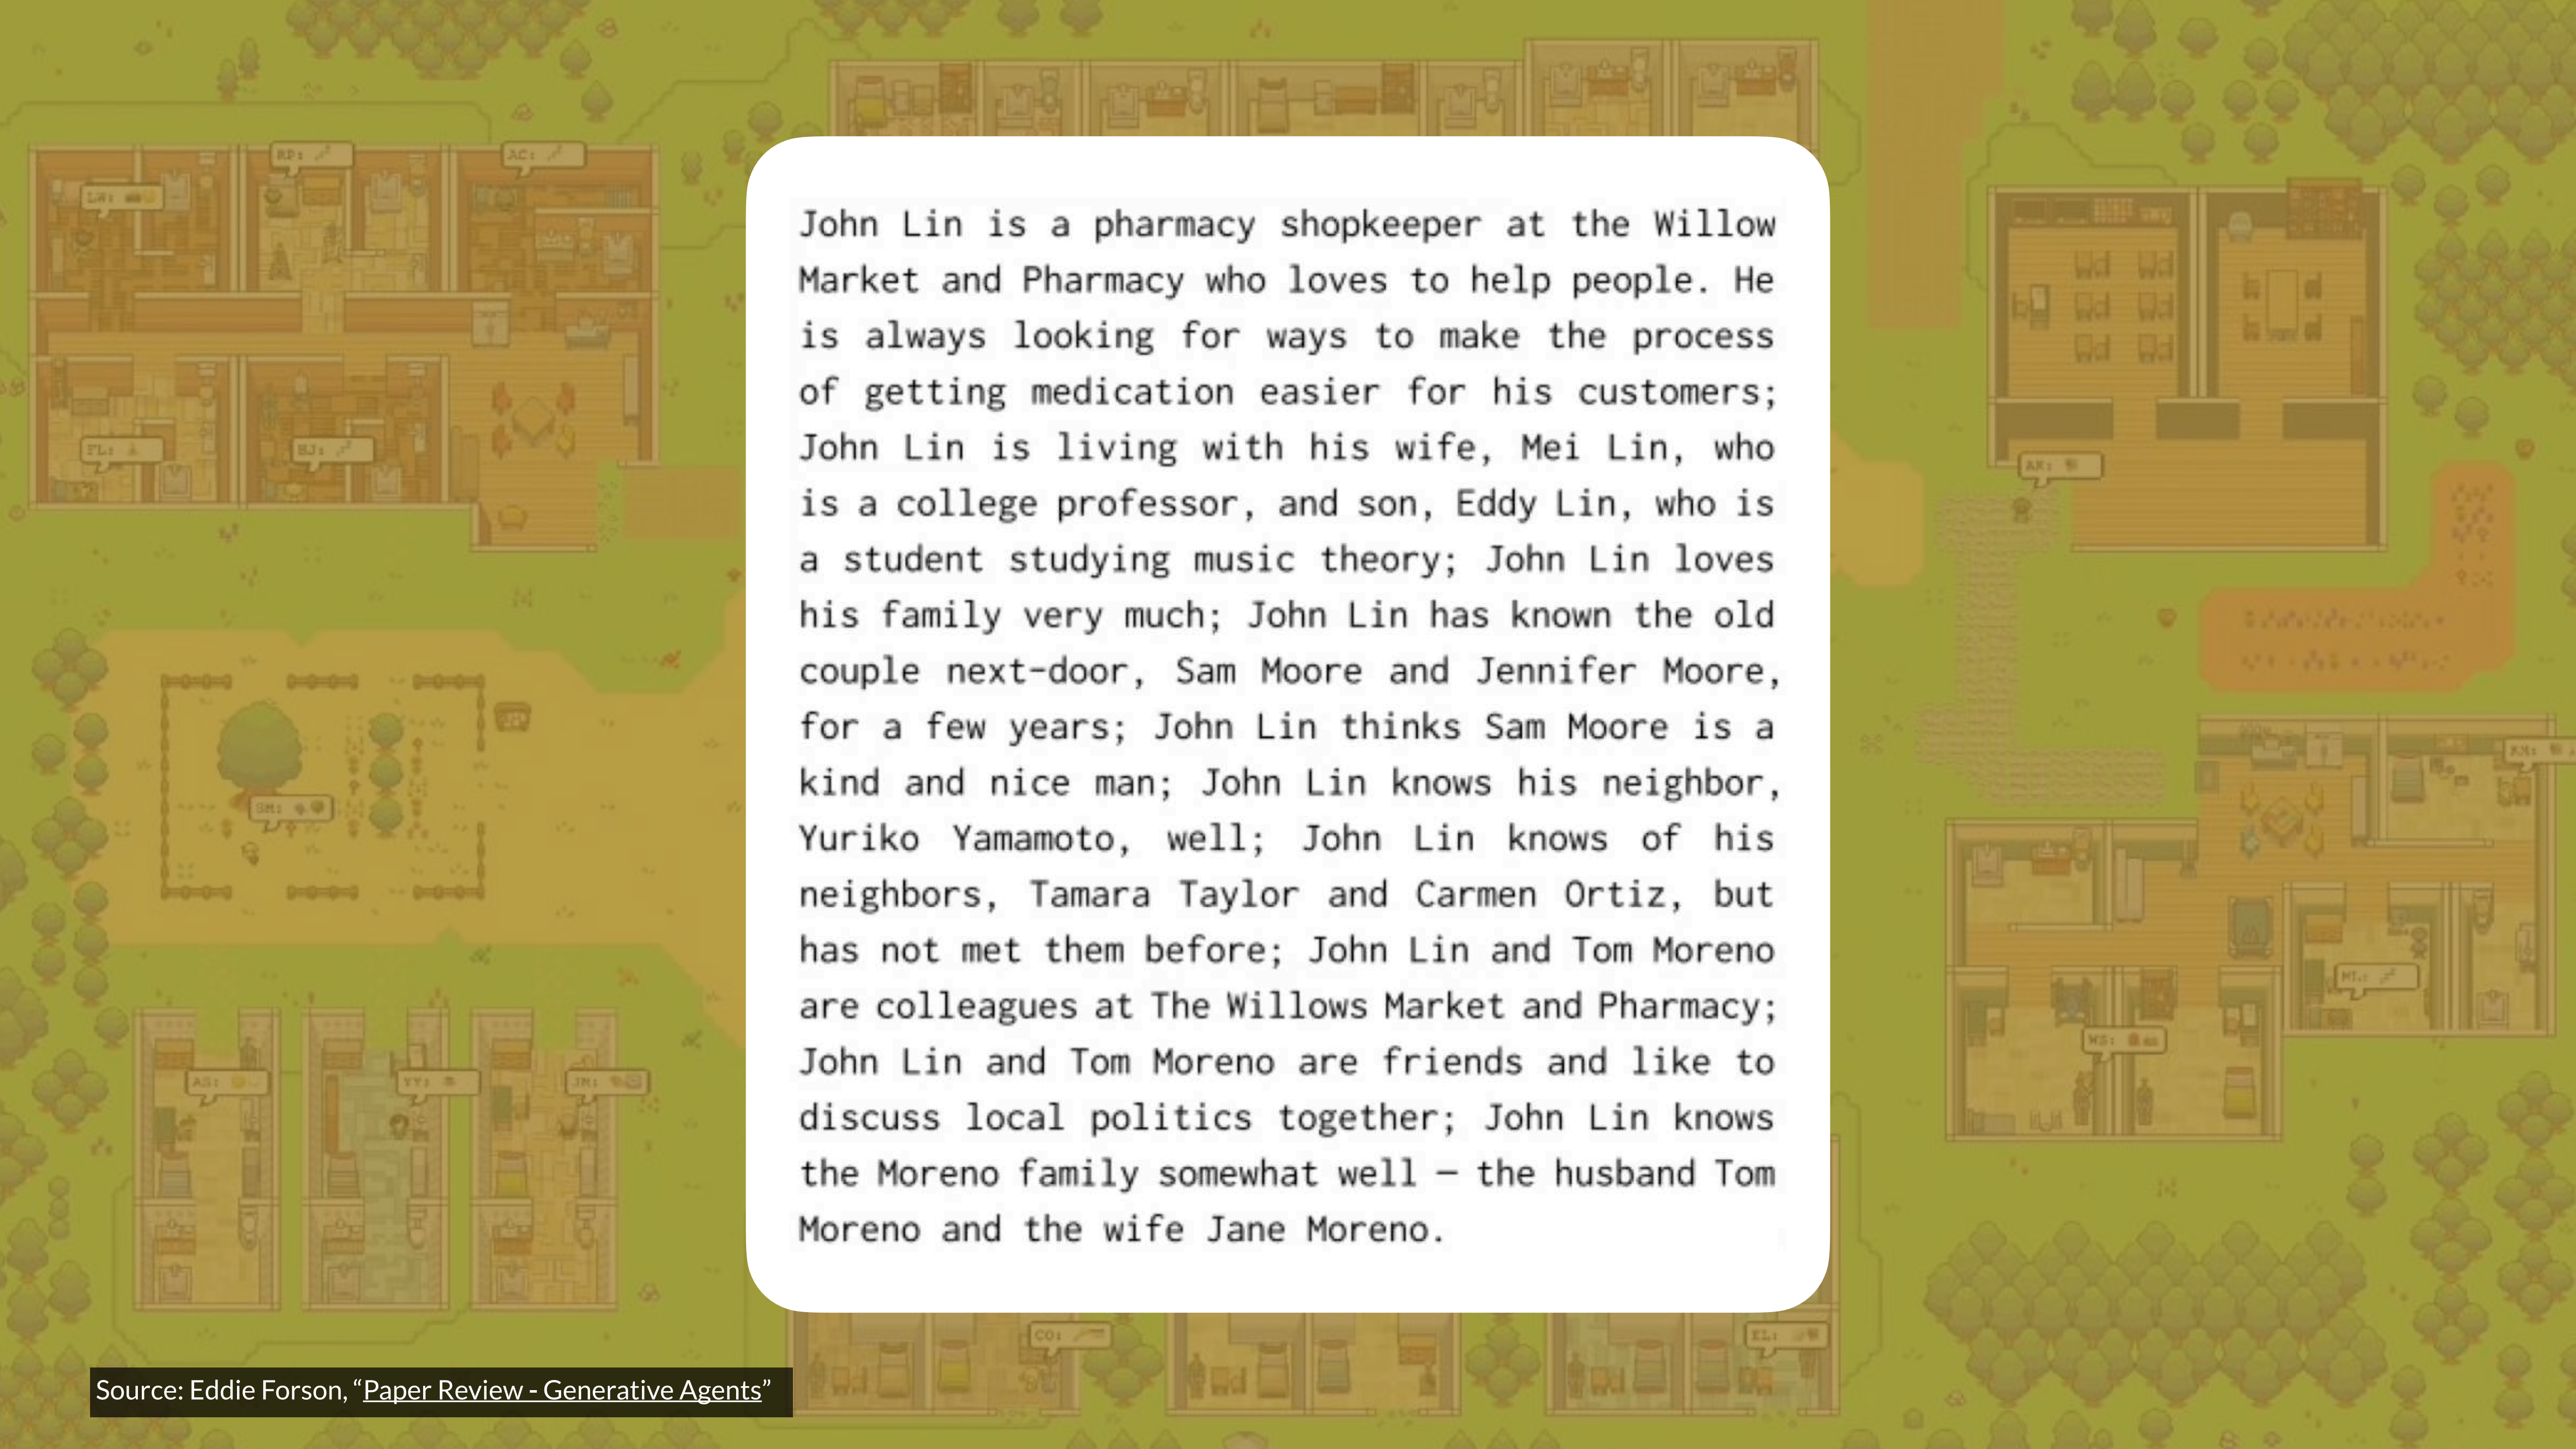


**STANFORD
HCI GROUP**



Google Research





John Lin is a pharmacy shopkeeper at the Willow Market and Pharmacy who loves to help people. He is always looking for ways to make the process of getting medication easier for his customers; John Lin is living with his wife, Mei Lin, who is a college professor, and son, Eddy Lin, who is a student studying music theory; John Lin loves his family very much; John Lin has known the old couple next-door, Sam Moore and Jennifer Moore, for a few years; John Lin thinks Sam Moore is a kind and nice man; John Lin knows his neighbor, Yuriko Yamamoto, well; John Lin knows of his neighbors, Tamara Taylor and Carmen Ortiz, but has not met them before; John Lin and Tom Moreno are colleagues at The Willows Market and Pharmacy; John Lin and Tom Moreno are friends and like to discuss local politics together; John Lin knows the Moreno family somewhat well – the husband Tom Moreno and the wife Jane Moreno.

Without prompting, the agents...

Formed relationships



[Abigail]: Hey Klaus, mind if
I join you for coffee?
[Klaus]: Not at all, Abigail.
How are you?

Without prompting, the agents...

Formed relationships



Established routines



Without prompting, the agents...

Formed relationships



Established routines



Planned and coordinated



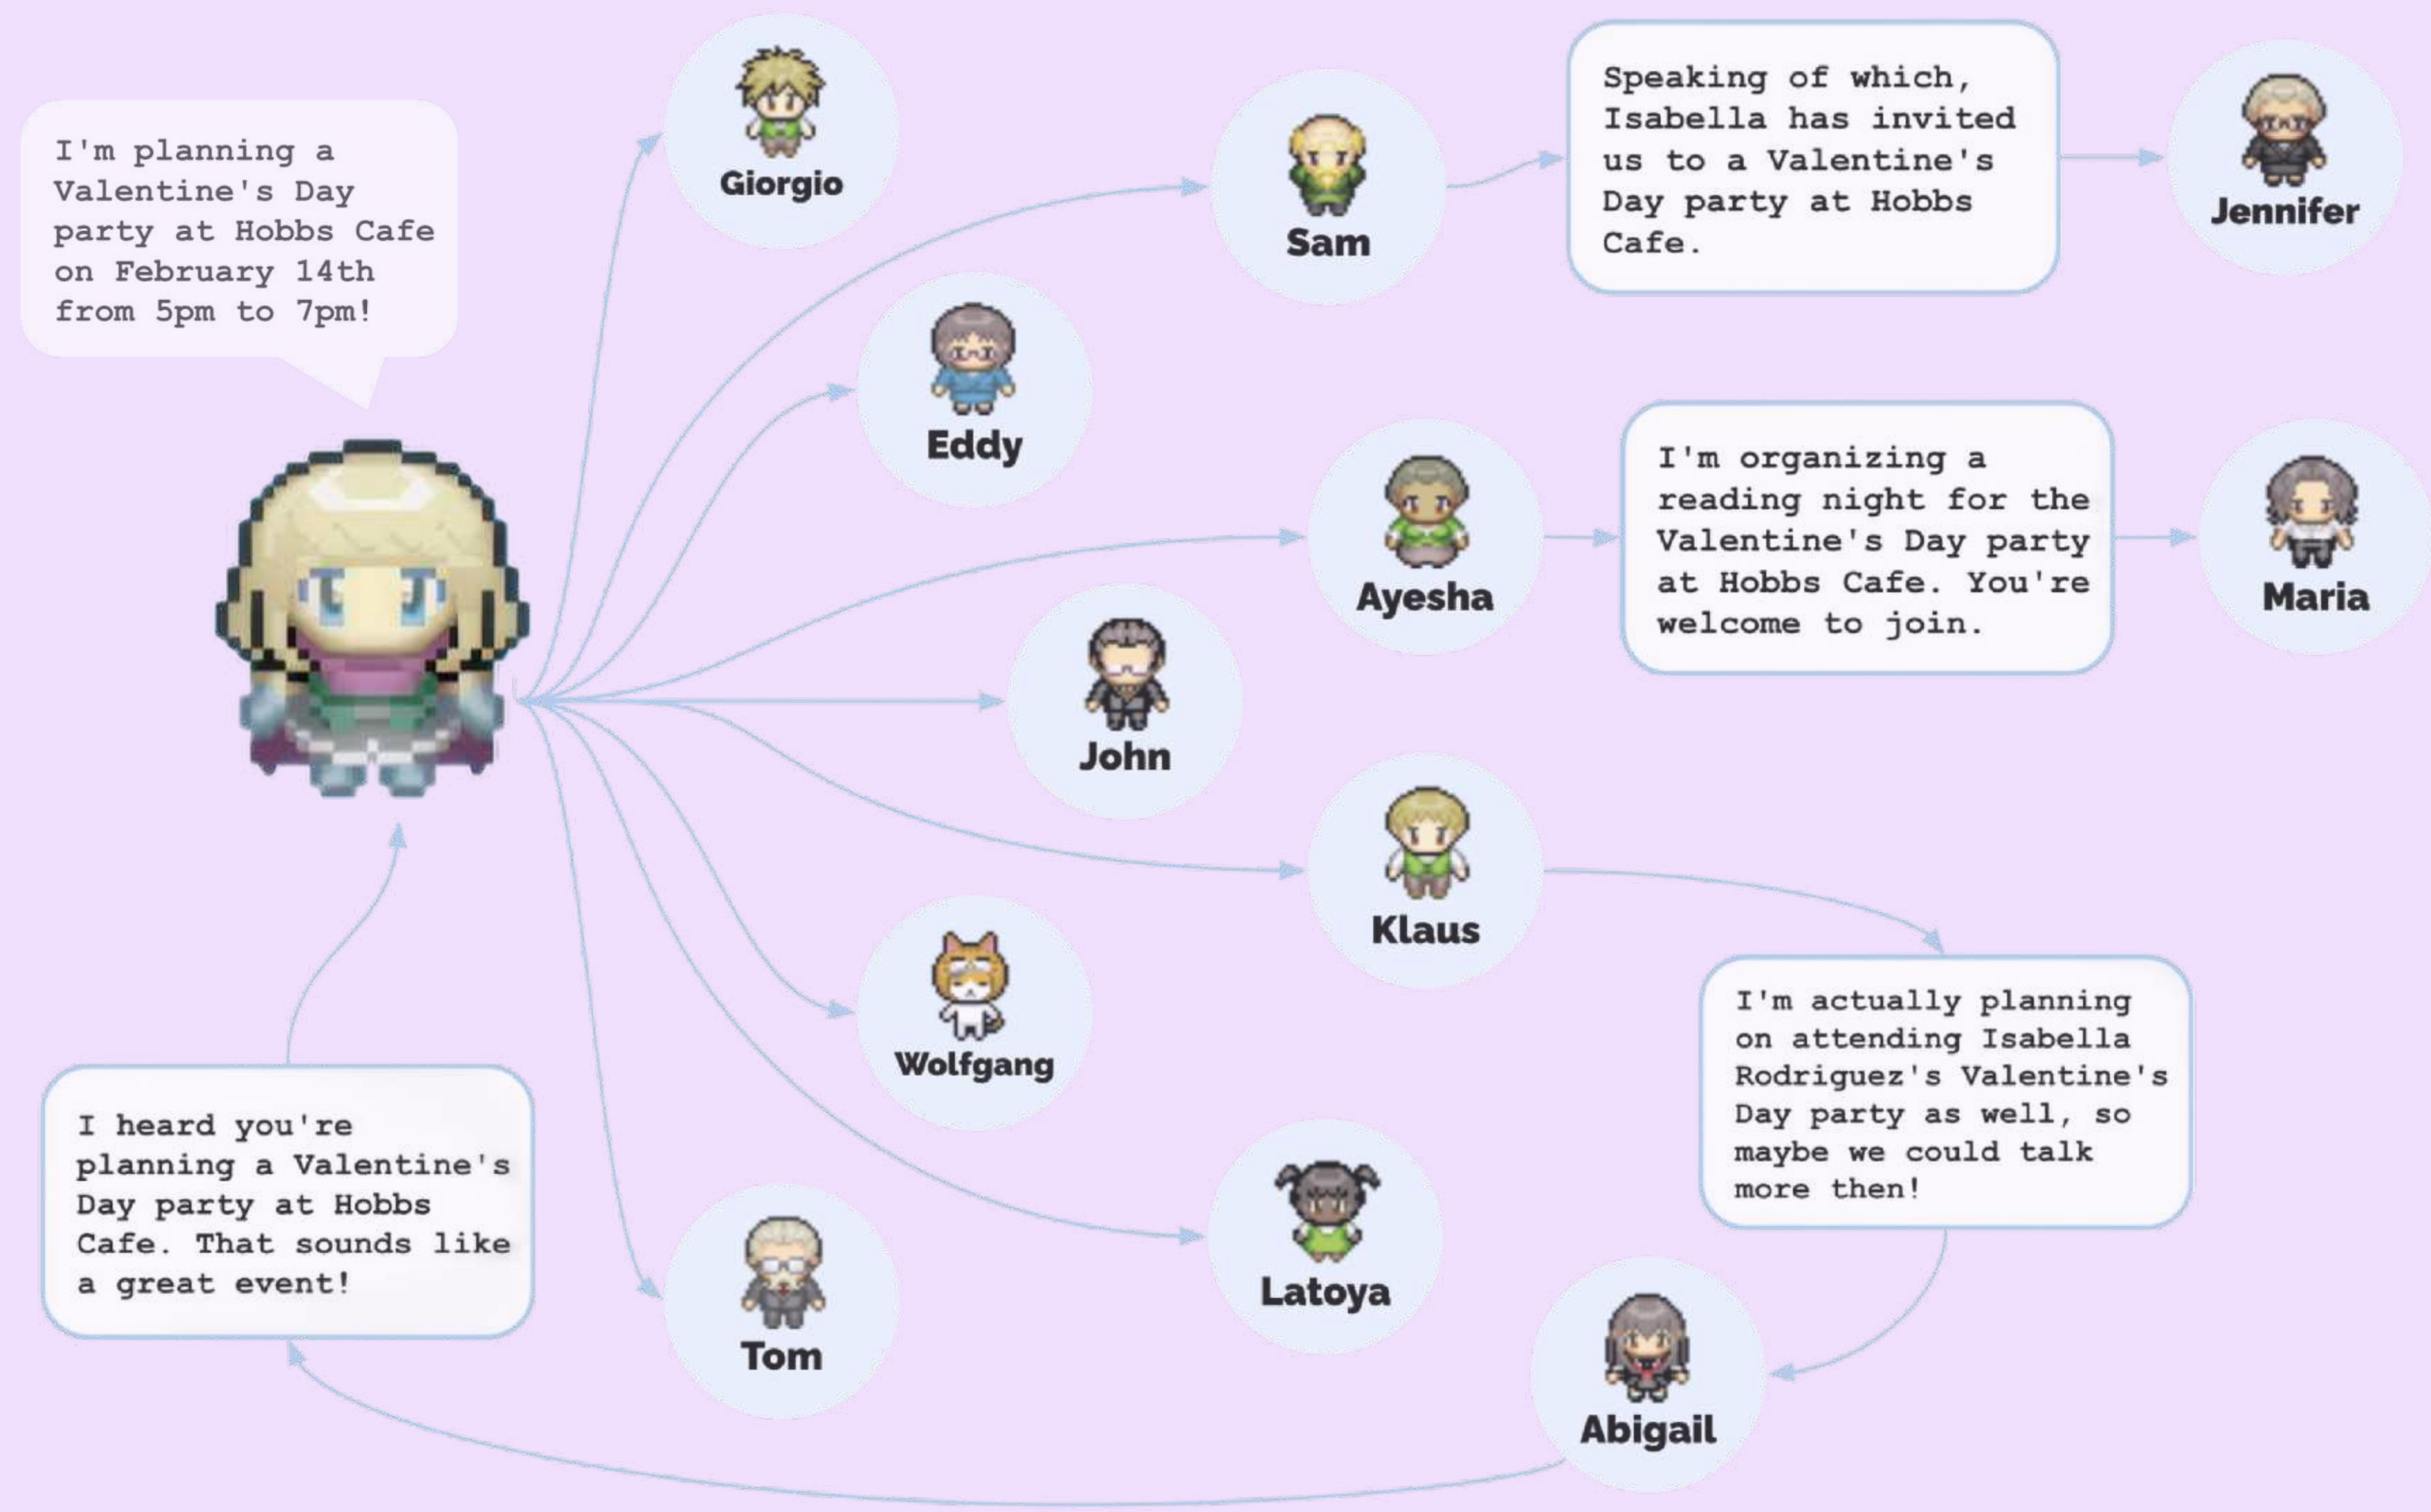
Figure 4:
At the beginning of the simulation, one agent is initialized with an intent to organize a Valentine's Day party.

I'm planning a
Valentine's Day
party at Hobbs Cafe
on February 14th
from 5pm to 7pm!

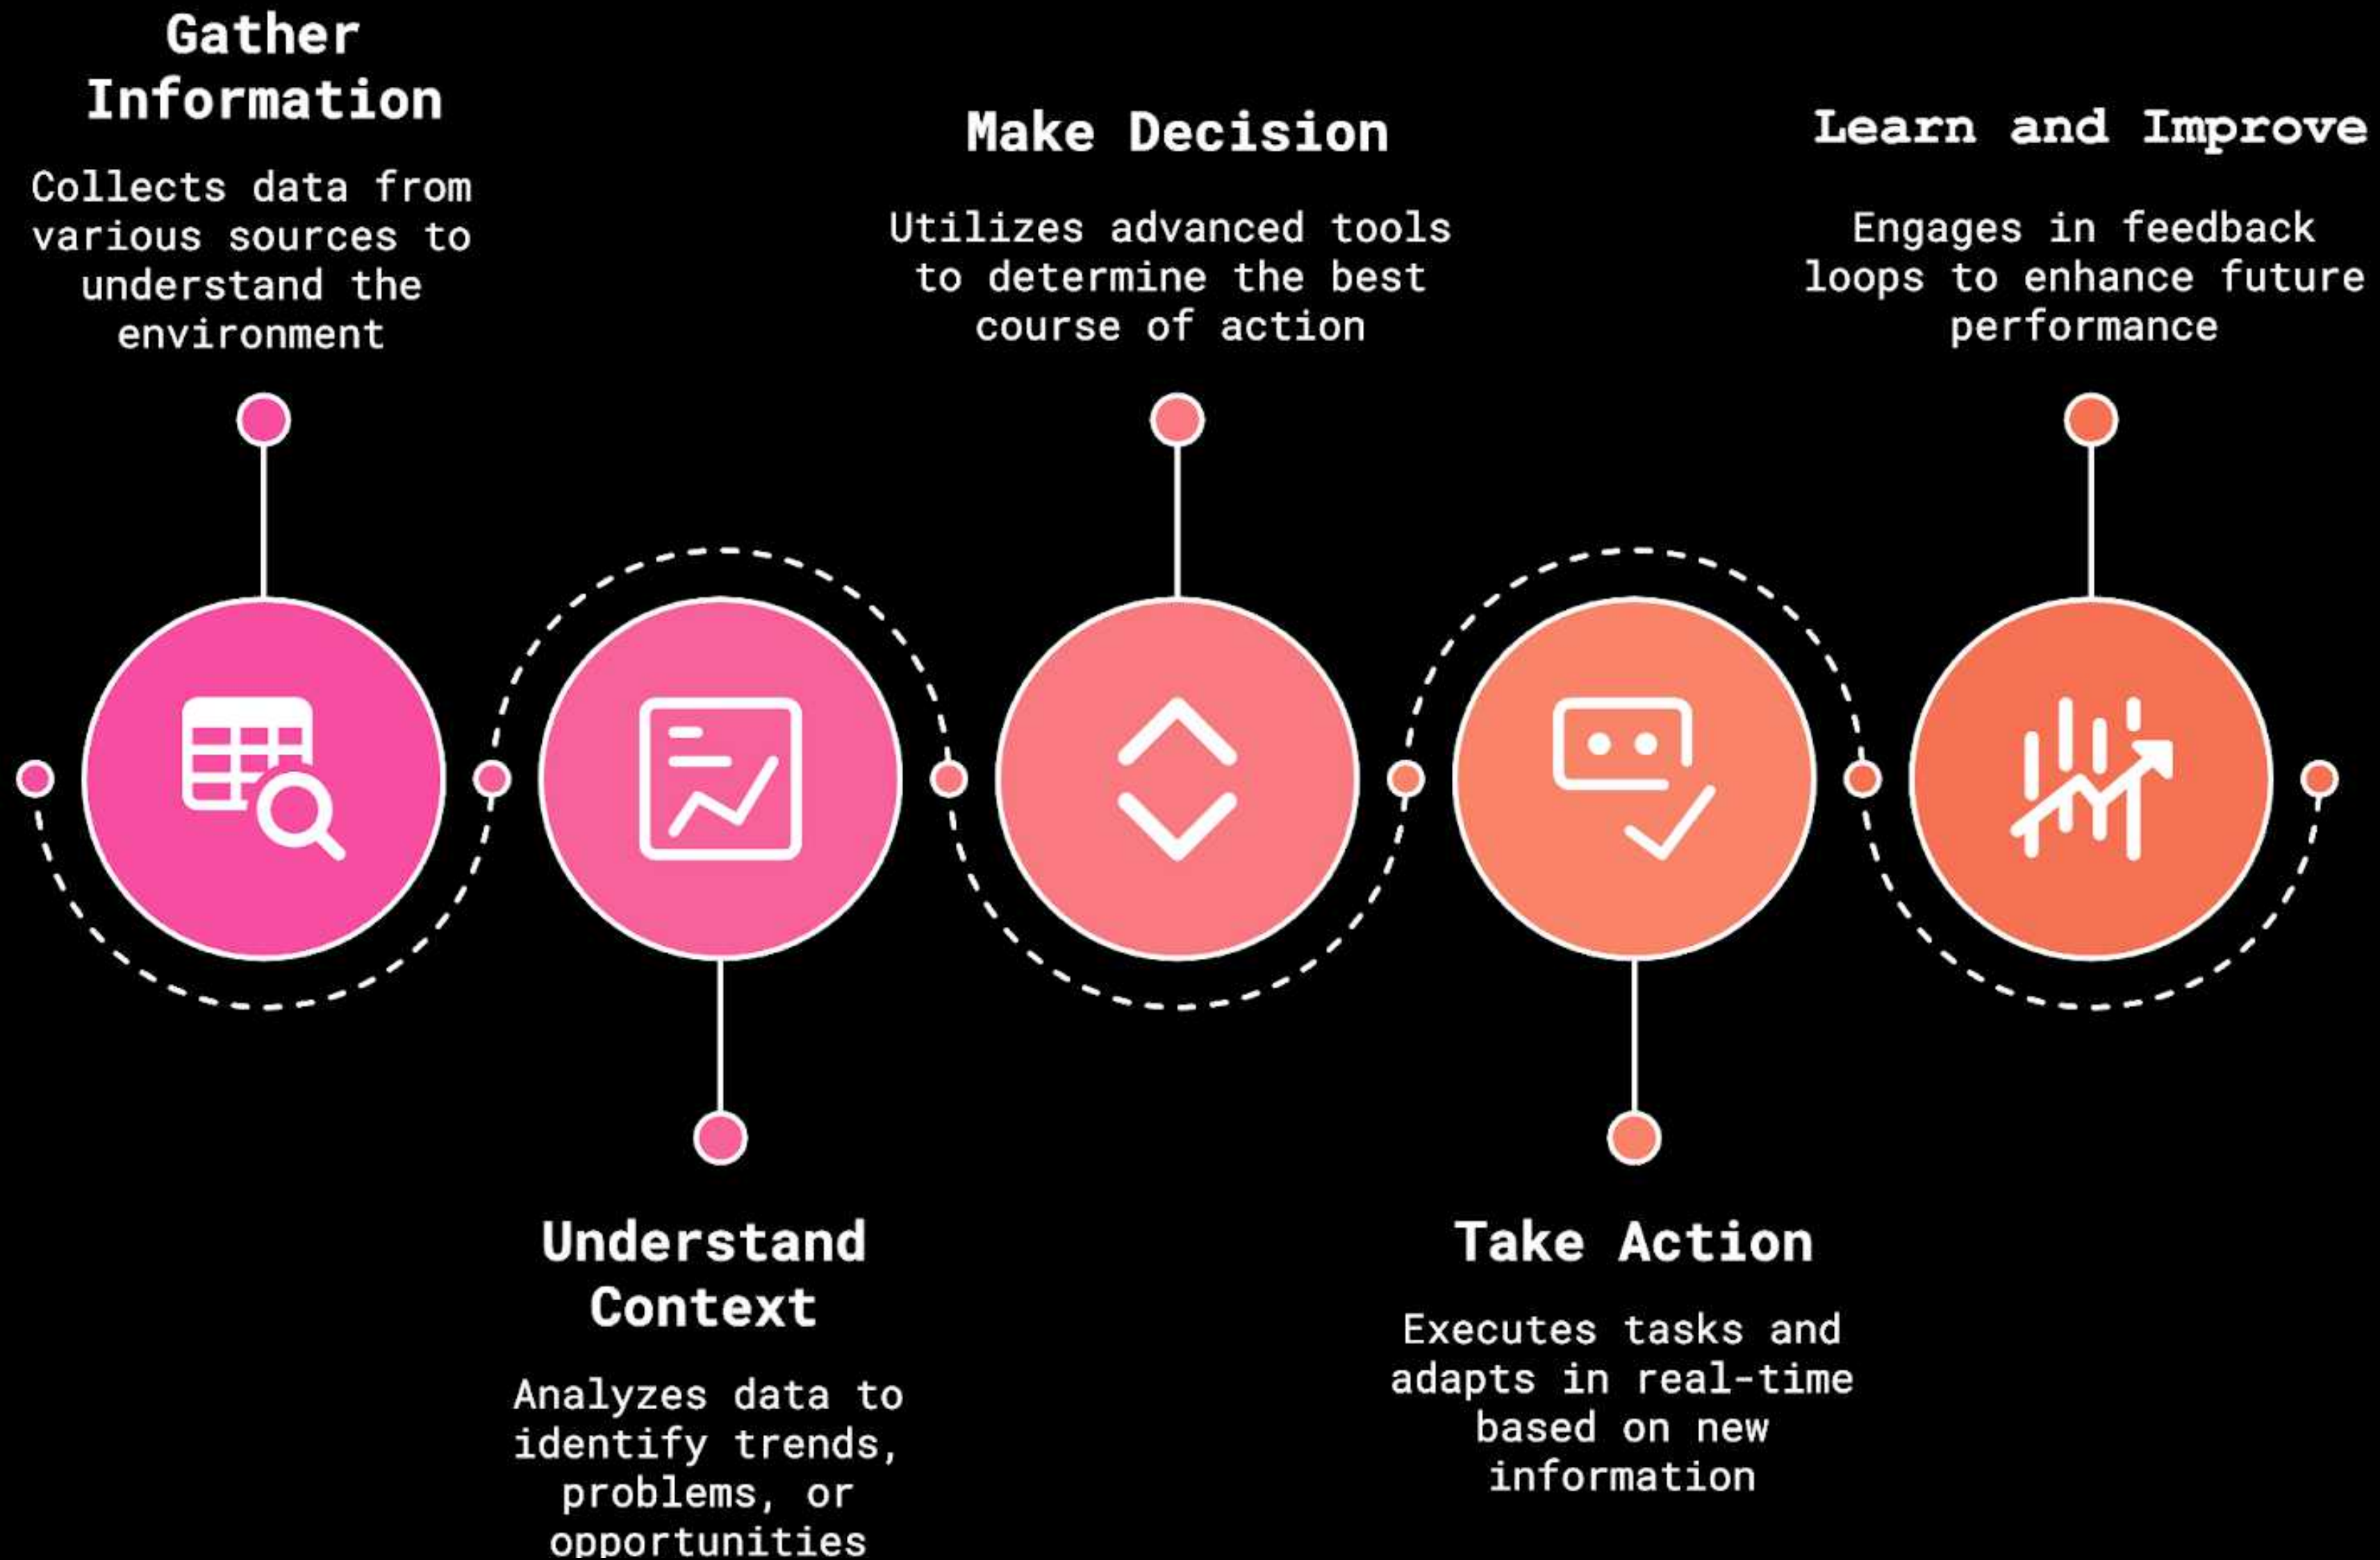


Figure 4:
At the beginning of the simulation, one agent is initialized with an intent to organize a Valentine's Day party.

Despite many possible points of failure in the ensuing chain of events- agents might not act on that intent, might forget to tell others, might not remember to show up-the Valentine's Day party does, in fact, occur, with a number of agents gathering and interacting.



Agentic AI Operational Process



Basic LLM-Based Workflow



Using an LLM to write a risk assessment report (not recommended)



Write a Enterprise Risk Assessment (ERA) report for the first quarter of 2025 (ending March 31, 2025). The report is intended for the Board of Directors and Senior Management.

Context & Instructions:

1. Use the attached reference documents as the source for all data, findings, and analysis. Do not use external knowledge or make up information beyond what is provided in these

< Desert Oasis Risk Report

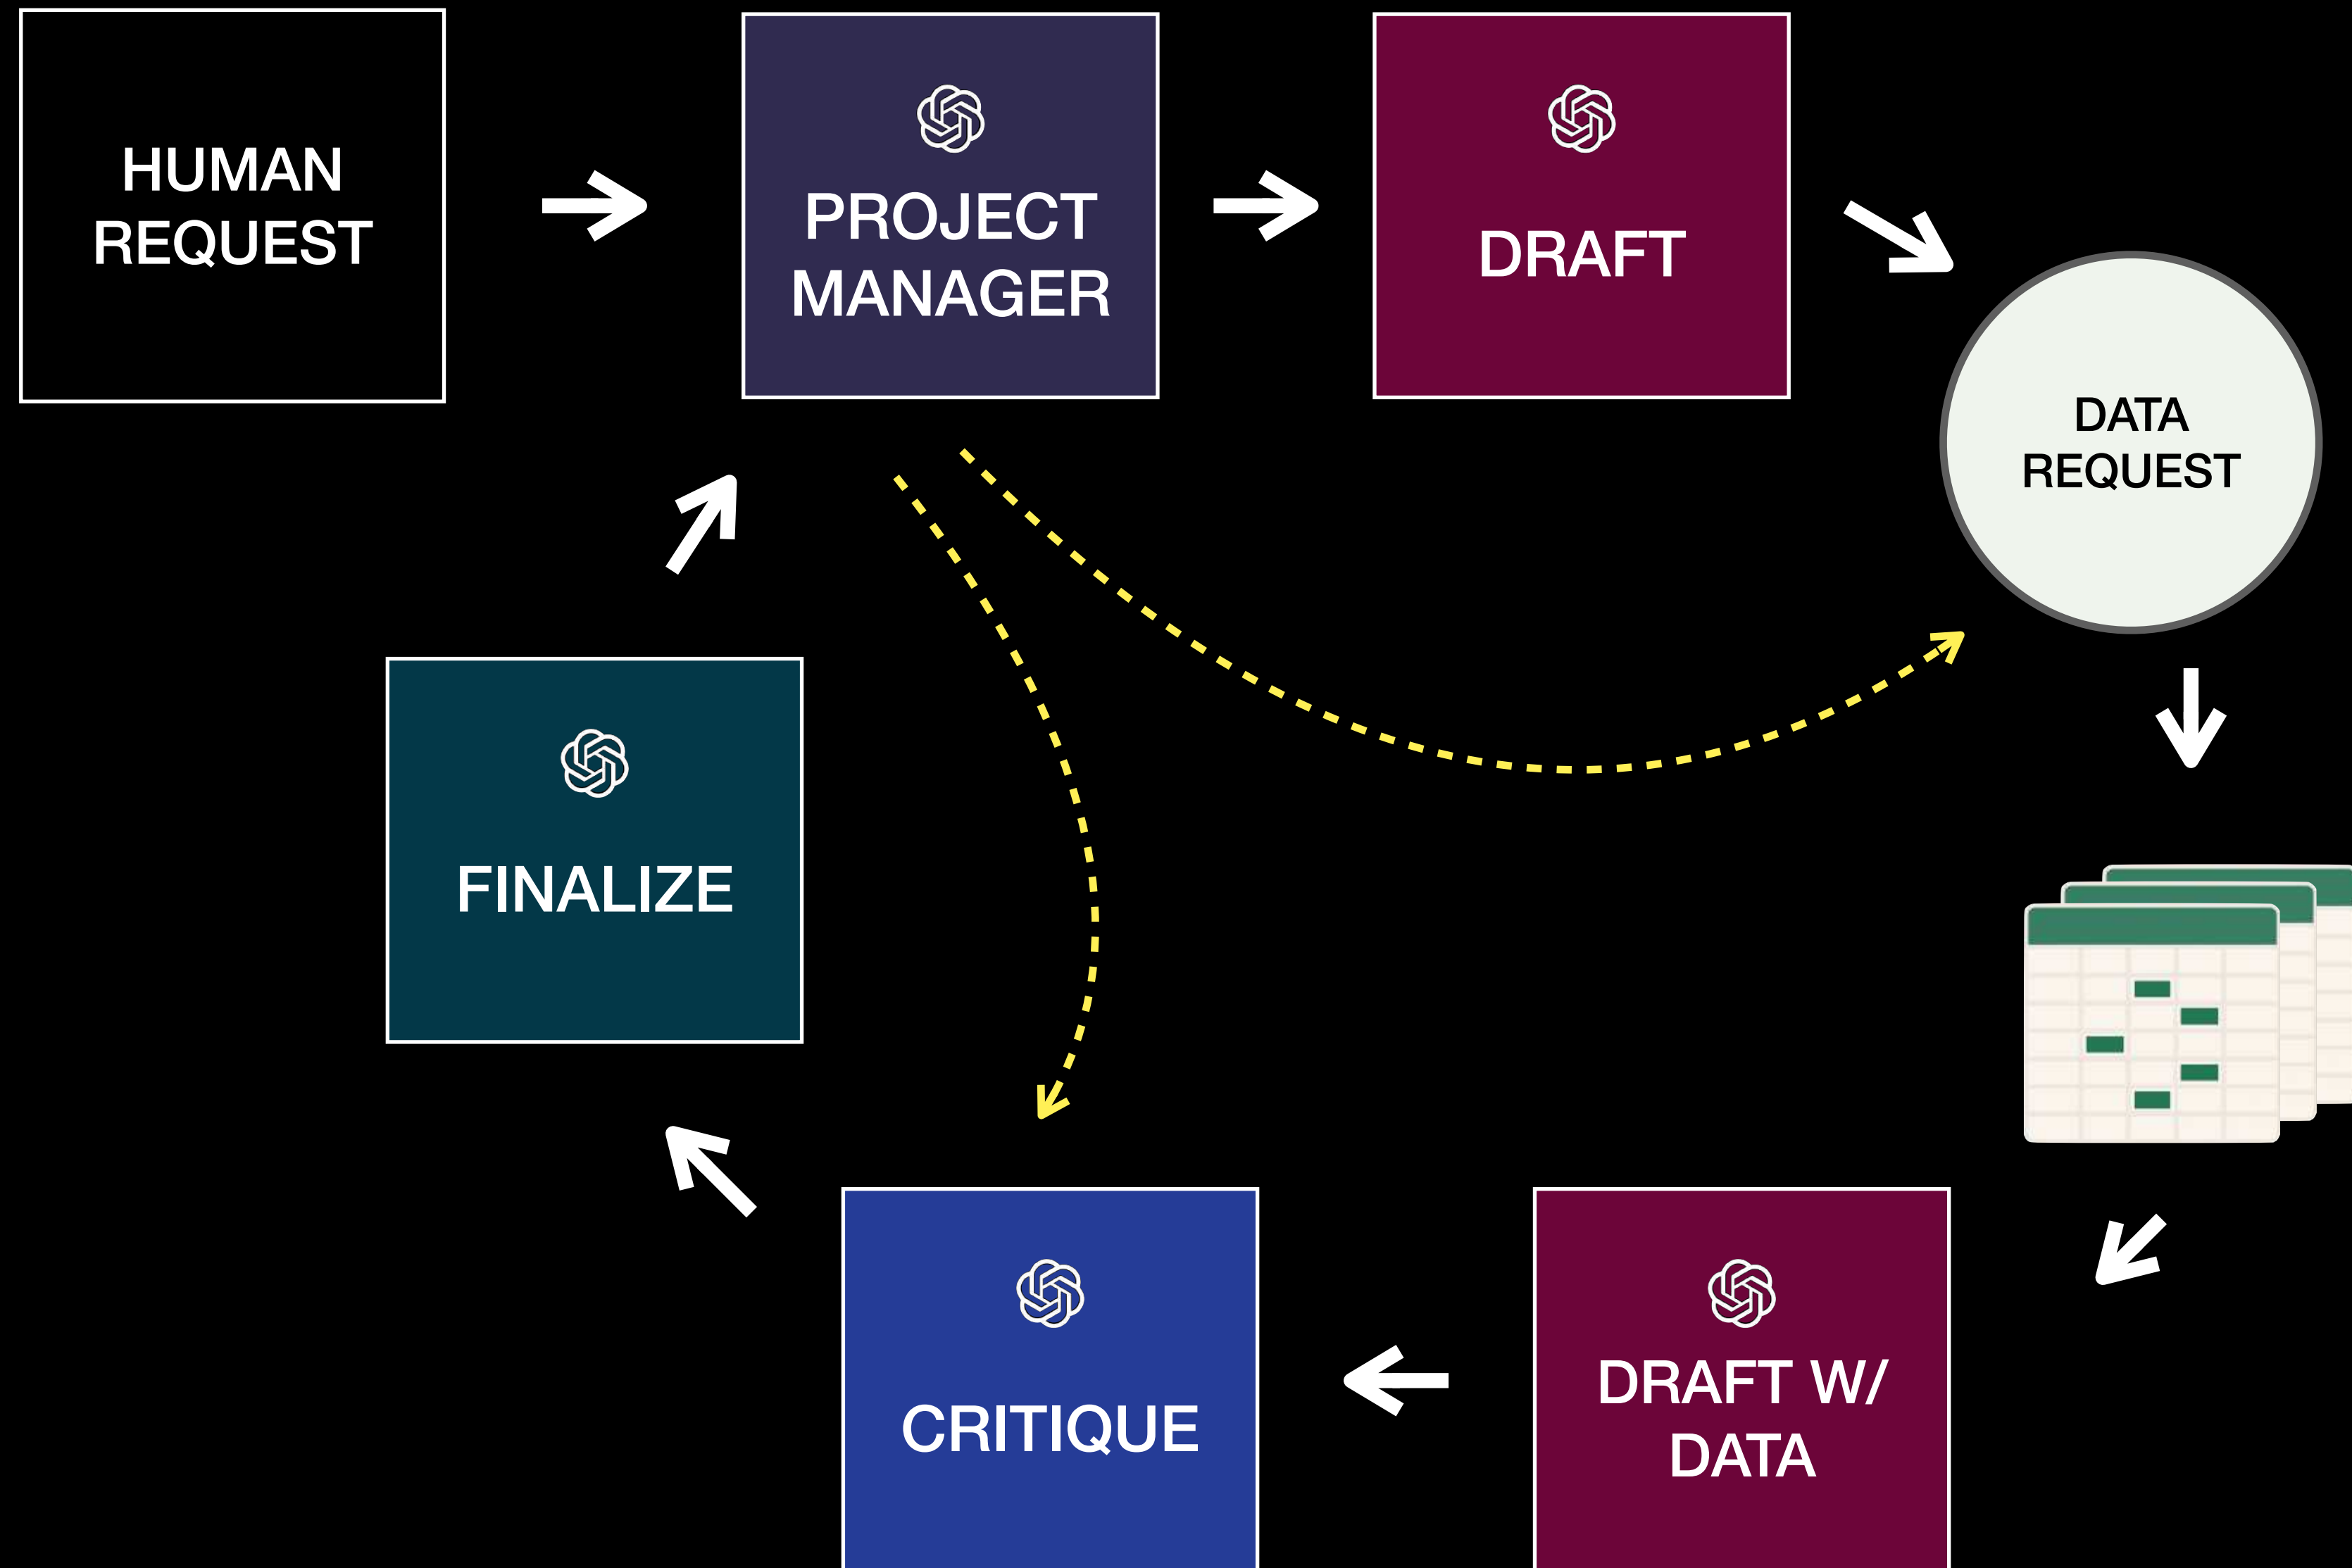
Desert Oasis Credit Union - Enterprise Risk Assessment Report

Report Date: May 4, 2025 Period Covered: Q1 2025 (January 1, 2025 – March 31, 2025) Prepared For: Board of Directors, Senior Management Prepared By: Office of the Chief Risk Officer
CONFIDENTIAL

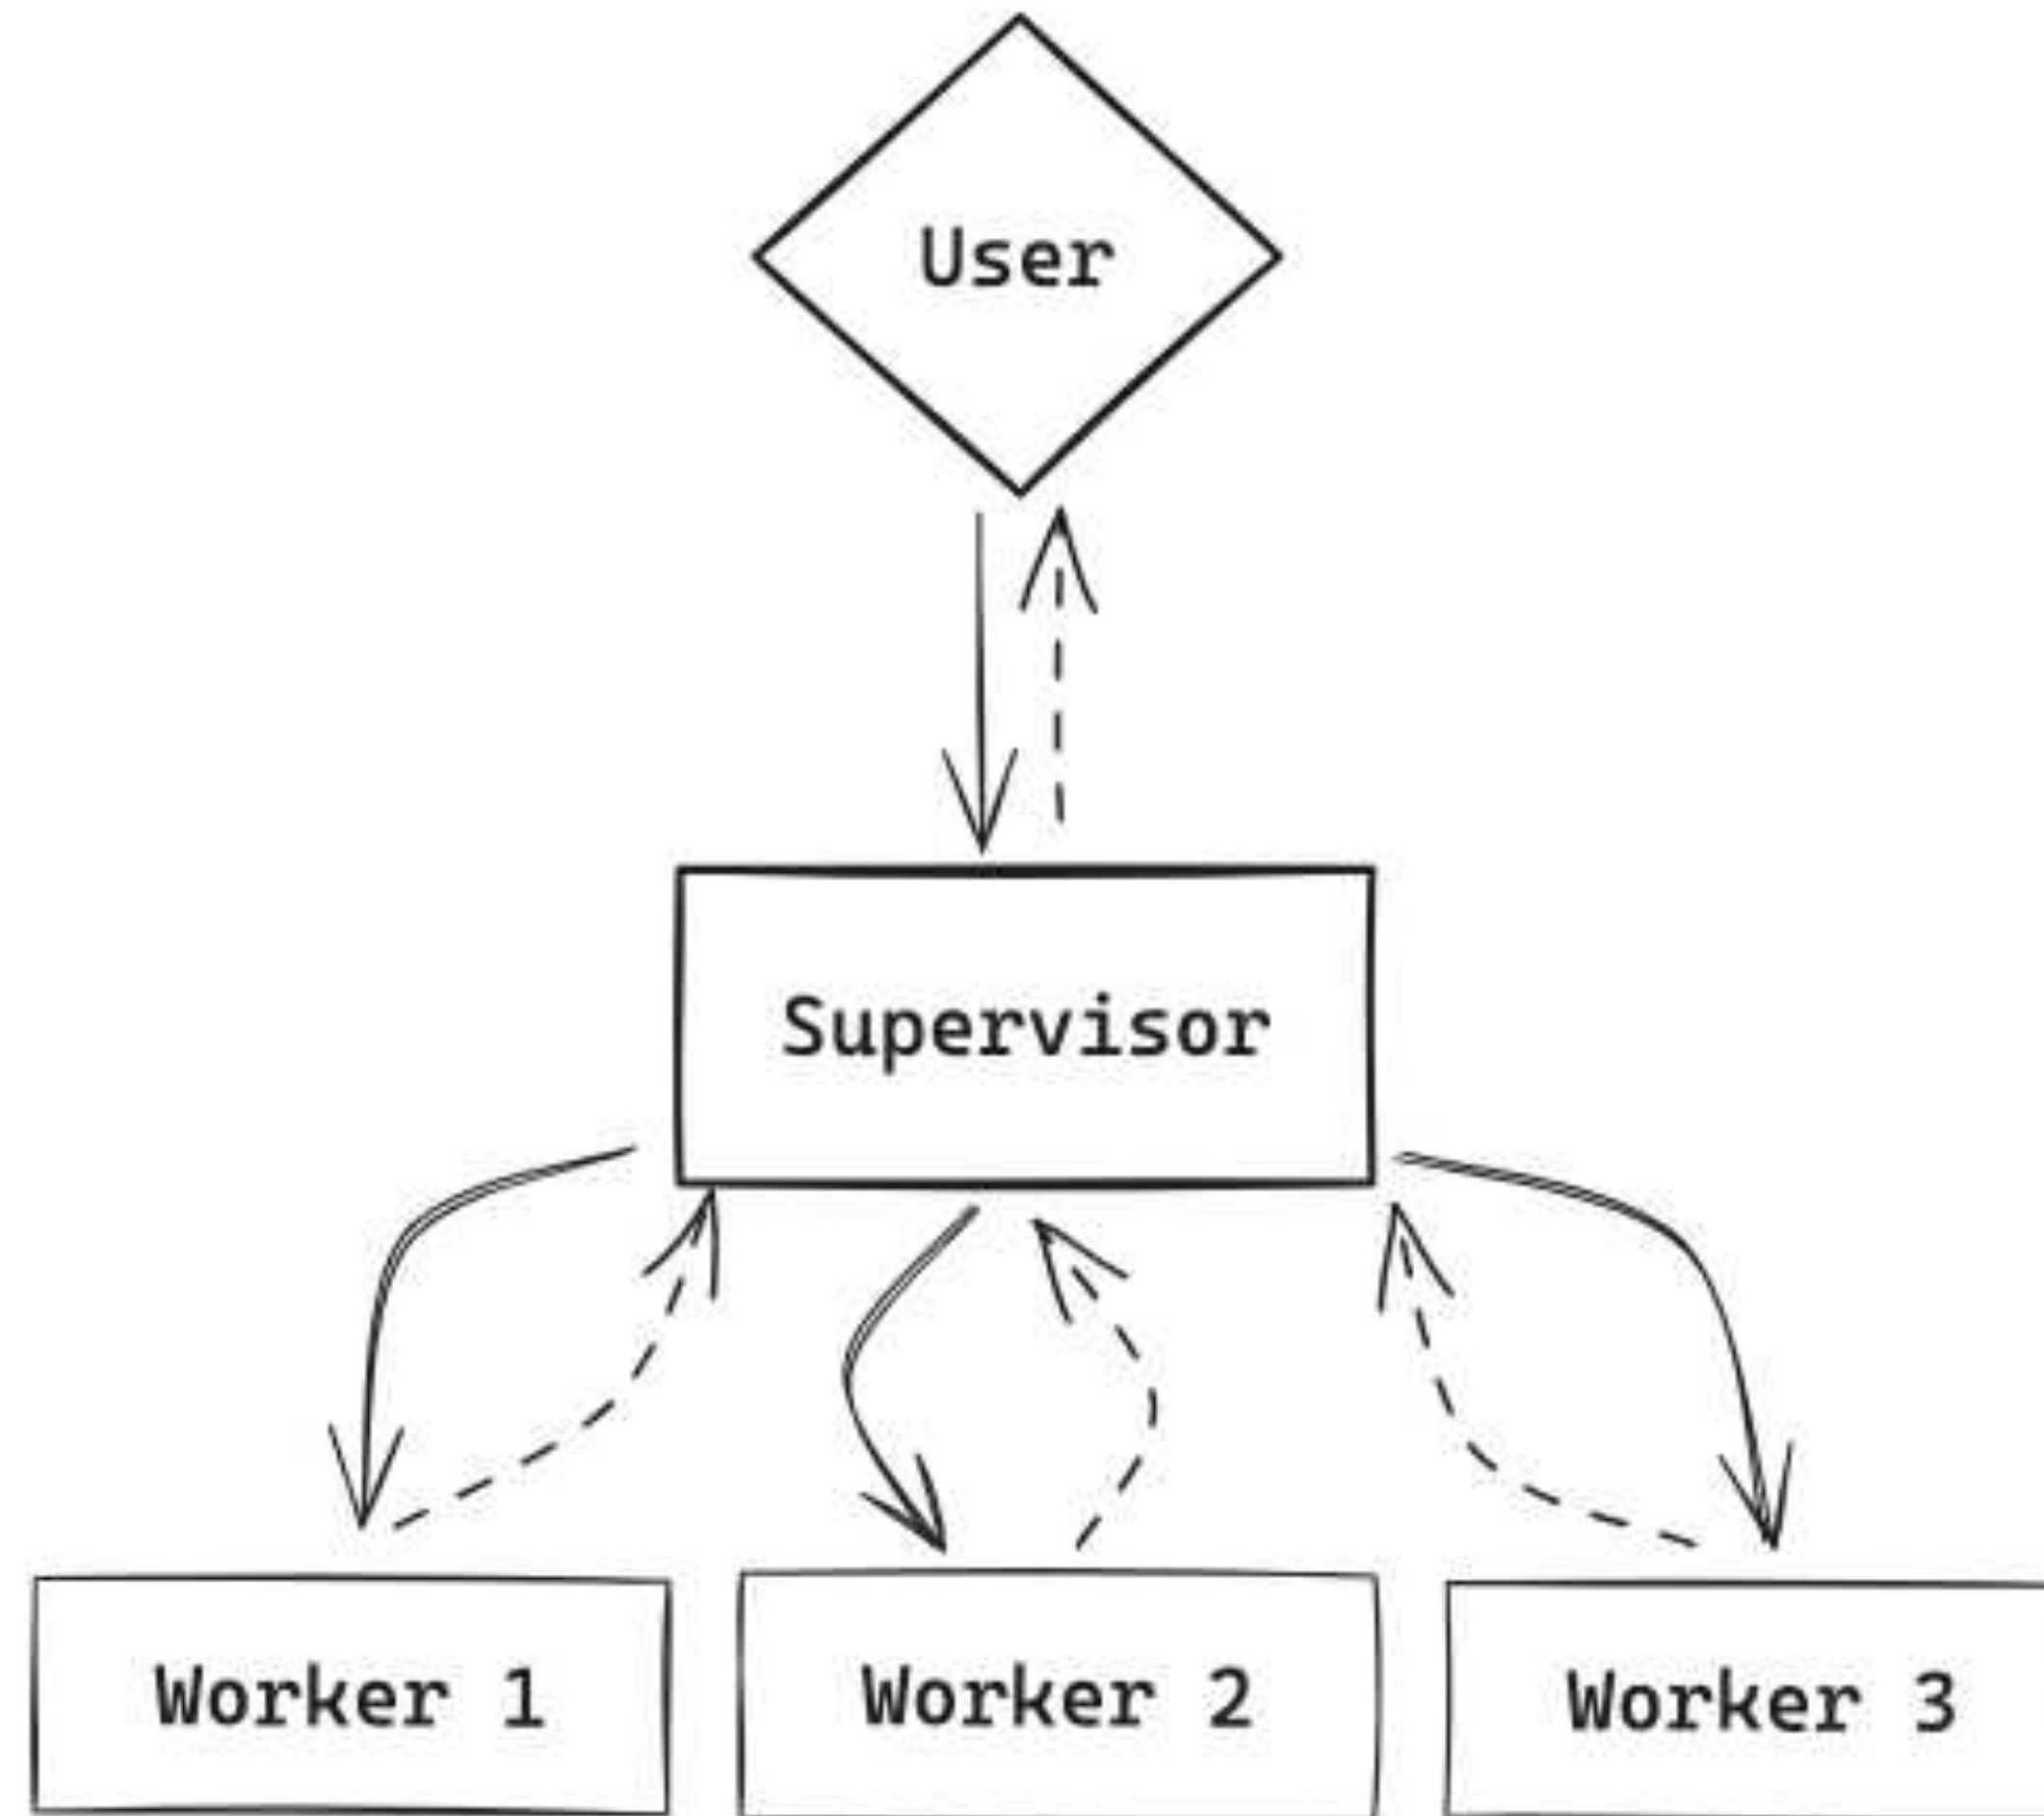
1. Executive Summary The overall risk profile for DOCU is assessed as Moderate, with a Stable outlook. However, Credit Risk and Operational Risk (specifically Cybersecurity) require heightened attention. Recommendations include enhancing cybersecurity controls, exploring loan portfolio diversification, refining IRR assumptions, and improving compliance training.

Agentic Workflow

Start here:



Multi-Agent Teams



**How are agents
showing up in banking?**

Persado to generate and optimize marketing content.

Thousands of A/B tests in a single deployment – across channels and segments

Reach your revenue targets sooner, while saving years of A/B testing efforts

Generate Banner Experiment

Campaign & Offer Info

Apply for travel rewards card ✕

Product Info

travel rewards card ✕ highly rated ✕

Other Content Considerations

summer ✕ travel ✕ vacation ✕

Target Segment

Deal Seekers ▾

Generate

➡➡

Unlock 2X Rewards for
Doing What You Love.

Continue

80 Performance
Prediction **4** Brand
Alignment

Tags: **Emotional:** Achievement

Reward Your Wanderlust
with Twice the Points.

Get Started

76 Performance
Prediction **4** Brand
Alignment

Tags: **Emotional:** Encouragement



Mastercard And Visa Unleash AI Agents To Shop For You

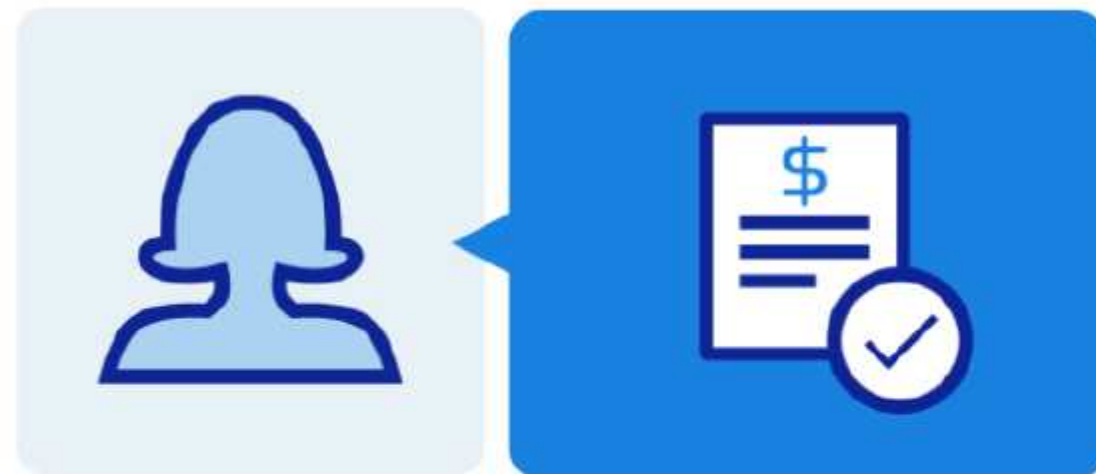
By [Tor Constantino, MBA](#), Contributor. ⓘ Tor Constantino is an ex-reporter, t... [▼](#)

[Follow Author](#)

May 05, 2025, 08:00am EDT

Card Tokenization

1 Step One



Capture and authenticate consumer instruction

Julie interacts with the agent and chooses the product she wants the agent to purchase on her behalf

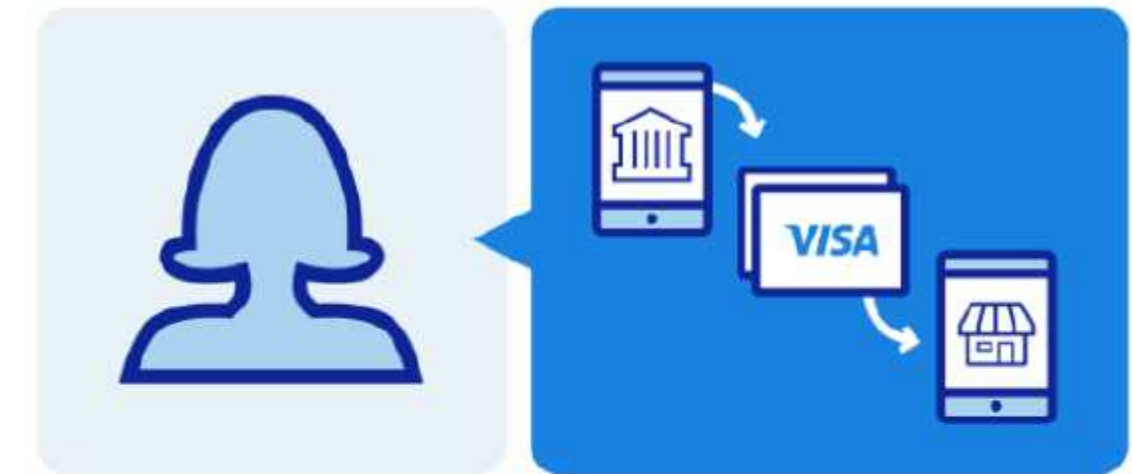
2 Step Two



Retrieve payment credentials

The agent retrieves payment credentials for the specific amount and merchant at which the purchase will be made

3 Step Three

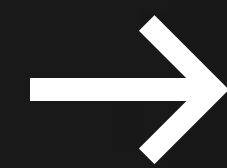


Provide the payment credentials to the merchant

The agent delivers the payment credentials to the merchant through the merchant's guest checkout, key entered, web form or through an available merchant API

Adding distance between you and the consumer.

FI as
central



Apple Pay /
Venmo



Payment
Agents

“We expect GenAI to empower banks’ entire risk and compliance functions in the future. This implies a profound culture change that will require all risk professionals to be conversant with the new tech, its capabilities, its limitations and how to mitigate those limitations.”

In the next five years, generative AI could fundamentally change financial institutions’ risk management by automating, accelerating, and enhancing everything from compliance to climate risk control.

by Rahul Agarwal, Andreas Kremer, Ida Kristensen, and Angela Luget

Risk & Resilience Practice

How generative AI can help banks manage risk and compliance

In the next five years, generative AI could fundamentally change financial institutions' risk management by automating, accelerating, and enhancing everything from compliance to climate risk control.

by Rahul Agarwal, Andreas Kremer, Ida Kristensen, and Angela Luget

**Automated
regulatory
compliance checks**

AI can scan and interpret regulations, and flag potential compliance issues.

**AML and fraud
detection**

AI can identify suspicious patterns and generate alerts for potential financial crimes.

**Compliance
monitoring and
reporting**

AI can automate compliance monitoring, generate reports, and highlight areas of concern.

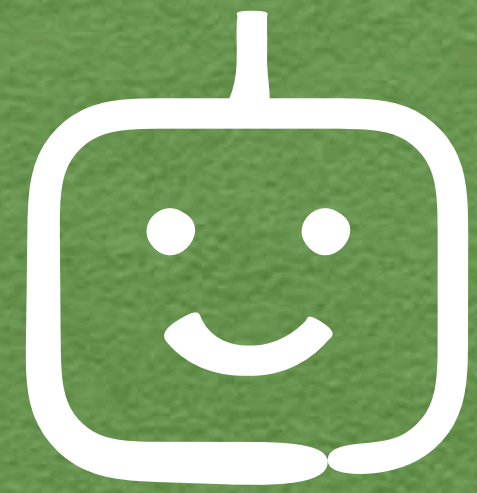
**Policy
& procedure
generation**

AI can assist in drafting and updating policies and procedures to ensure regulatory compliance.



Agentic moves and multi-agent design

Wells Fargo is also moving toward more autonomous systems. Mehta described a recent project to re-underwrite 15 years of archived loan documents. The bank used a network of interacting agents, some of which are built on open source frameworks like LangGraph. Each agent had a specific role in the process, which included retrieving documents from the archive, extracting their contents, matching the data to systems of record, and then continuing down the pipeline to perform calculations – all tasks that traditionally require human analysts. A human reviews the final output, but most of the work ran autonomously.



Where agents are strong:

Research & Data Extraction

Quickly find, summarize, and organize complex information from multiple sources.

Productivity Boost

Automate repetitive tasks like scheduling and managing tools.

Enhanced Service

Resolve complex issues by understanding context and multi-step tasks.

Structured Tasks

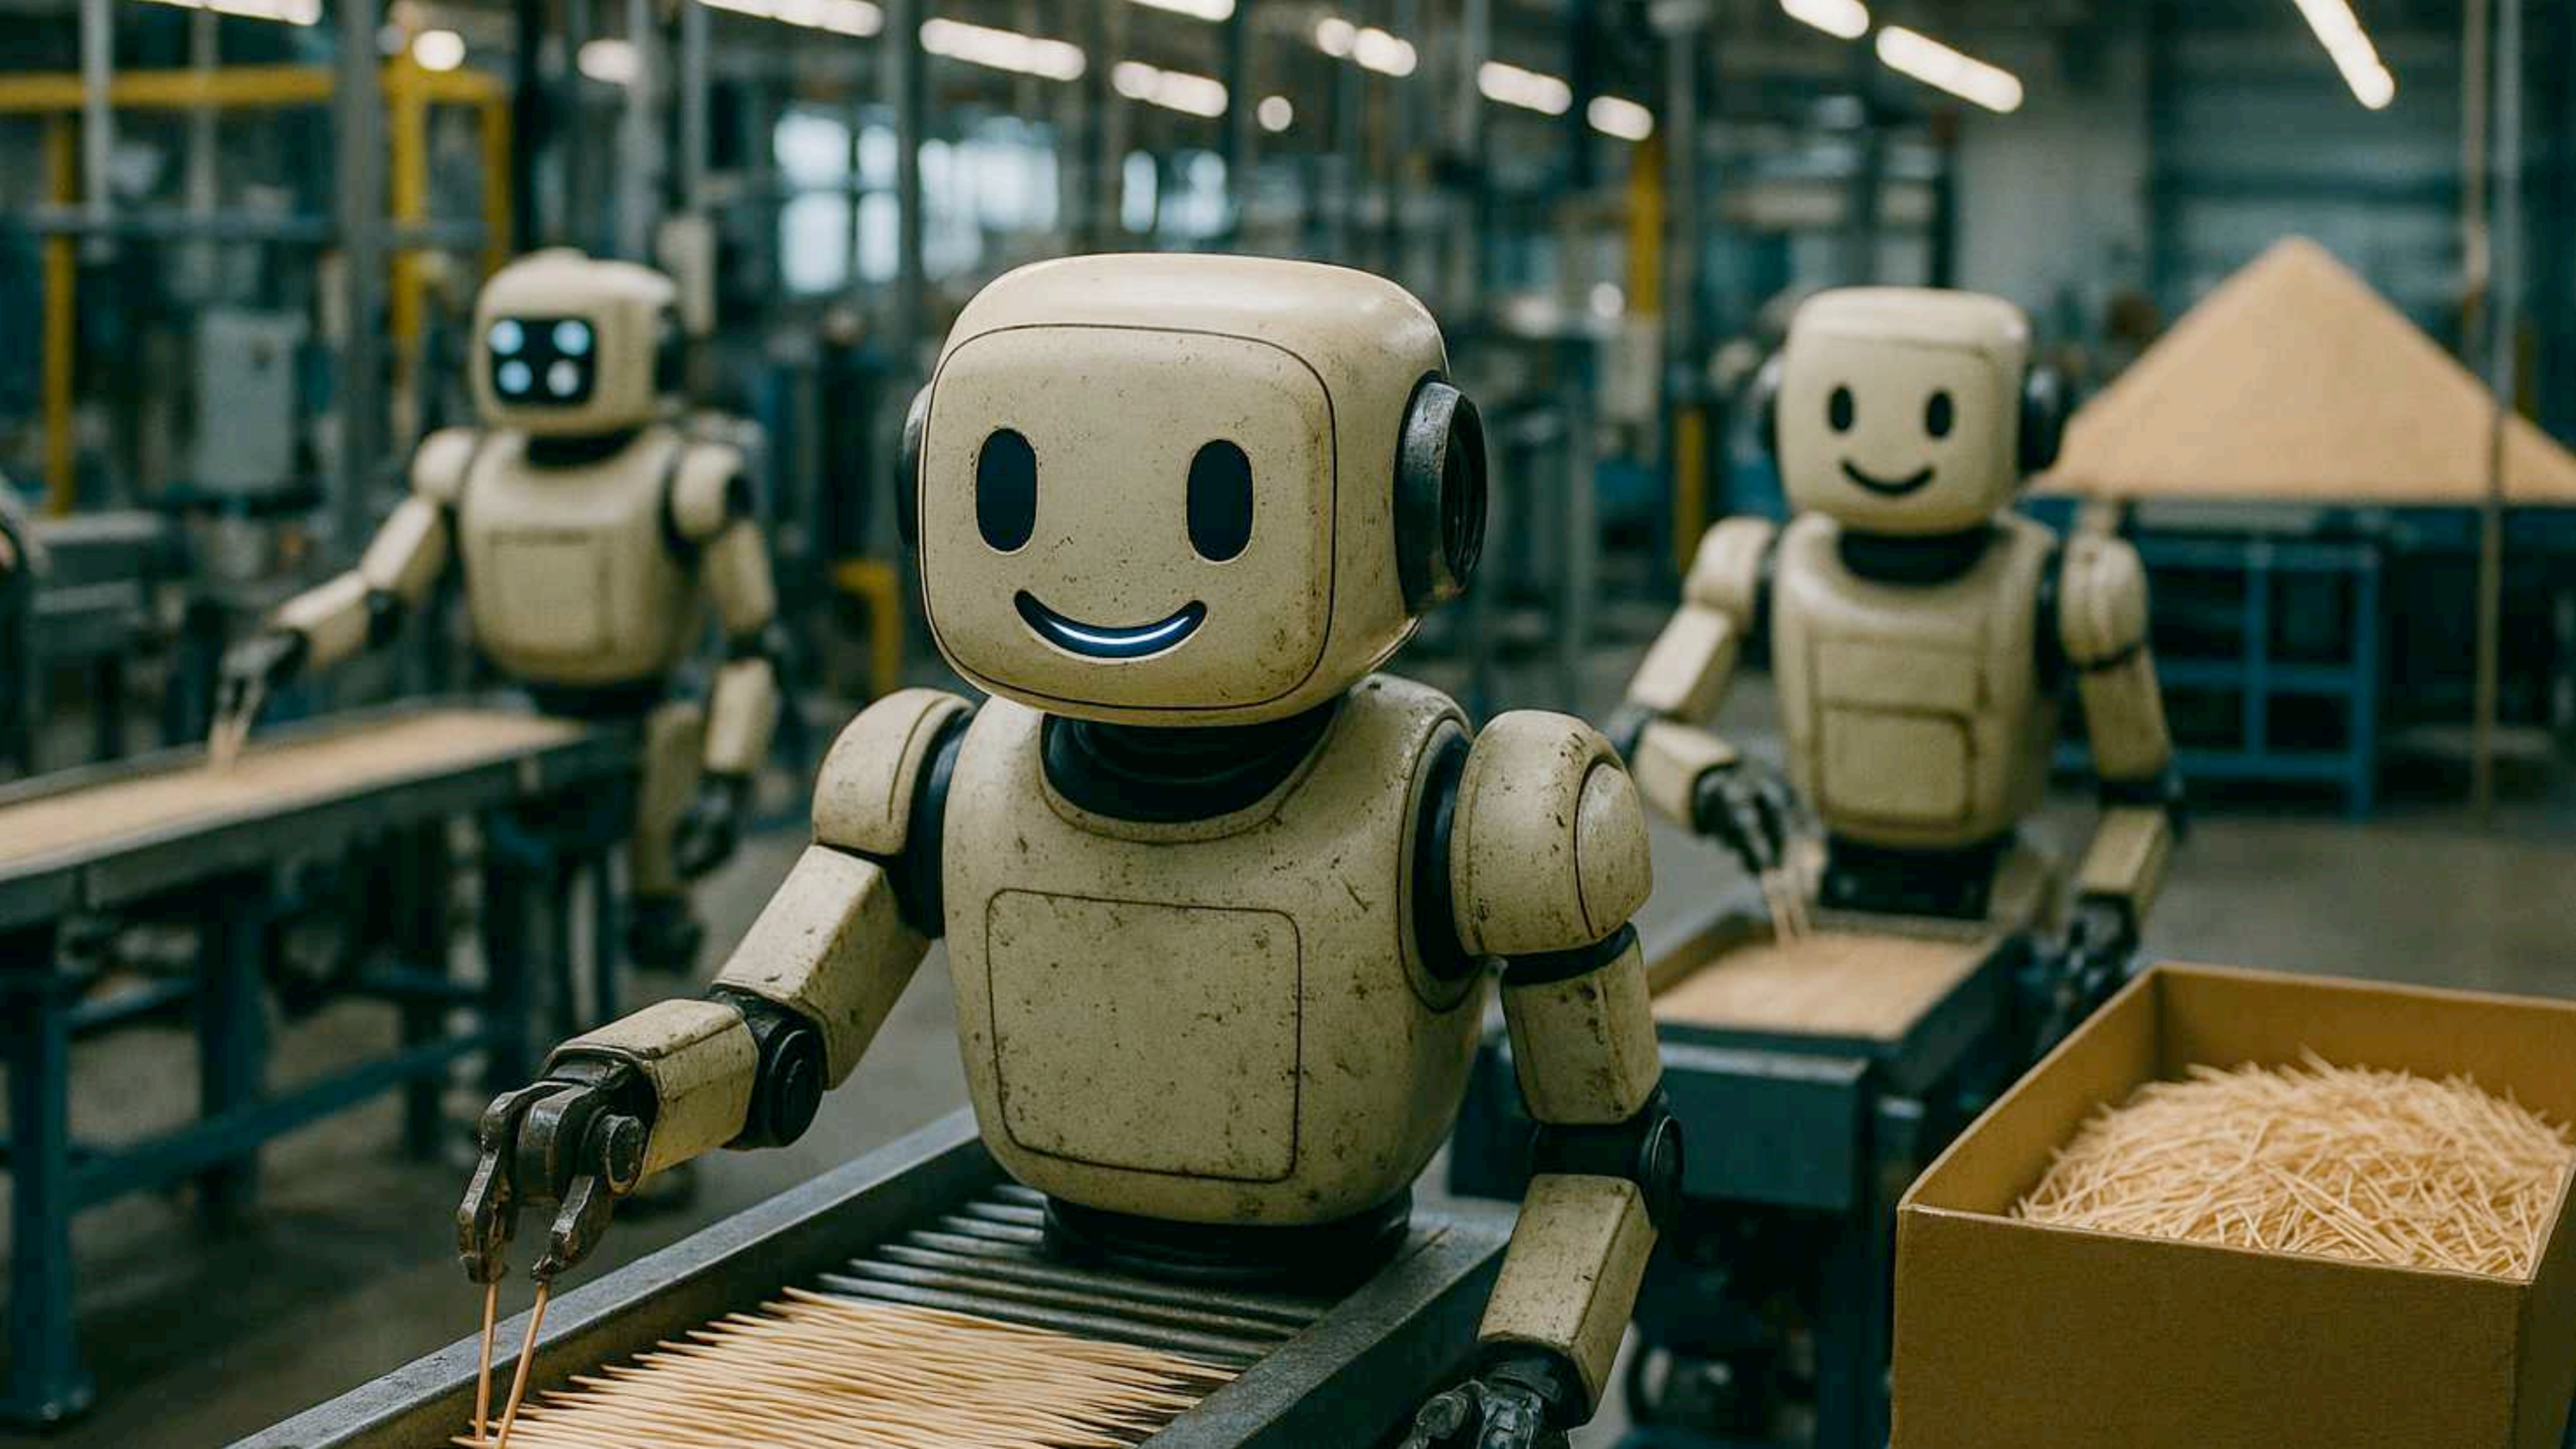
Excel in clear, predictable scenarios—like resume screening, automated service responses, or real-time analytics.

Workflow Automation

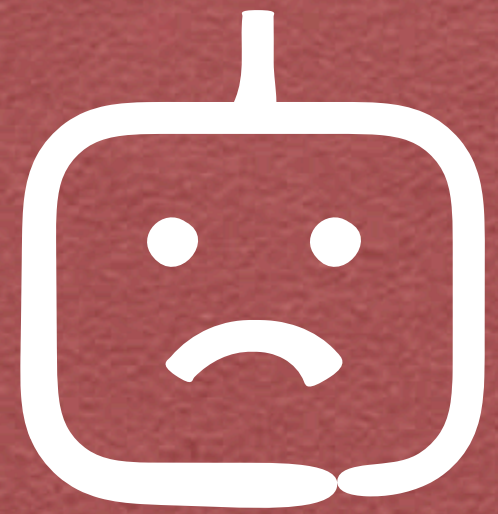
Manage complete processes end-to-end without human prompts.

Real-Time Fraud & Risk Management

Instantly flag suspicious patterns and transactions.







Where agents are weak:

Reliability Concerns

Require constant monitoring to avoid errors.

Safety Risks

Might act unpredictably if given unclear objectives.

Transparency Issues:

Decision-making processes can be difficult to trace and audit.

Error Cascades

Small mistakes can quickly spread across connected tasks.

System Complexity

Complicated to build and maintain.

Resource Intensive:

High computational demands and costs.

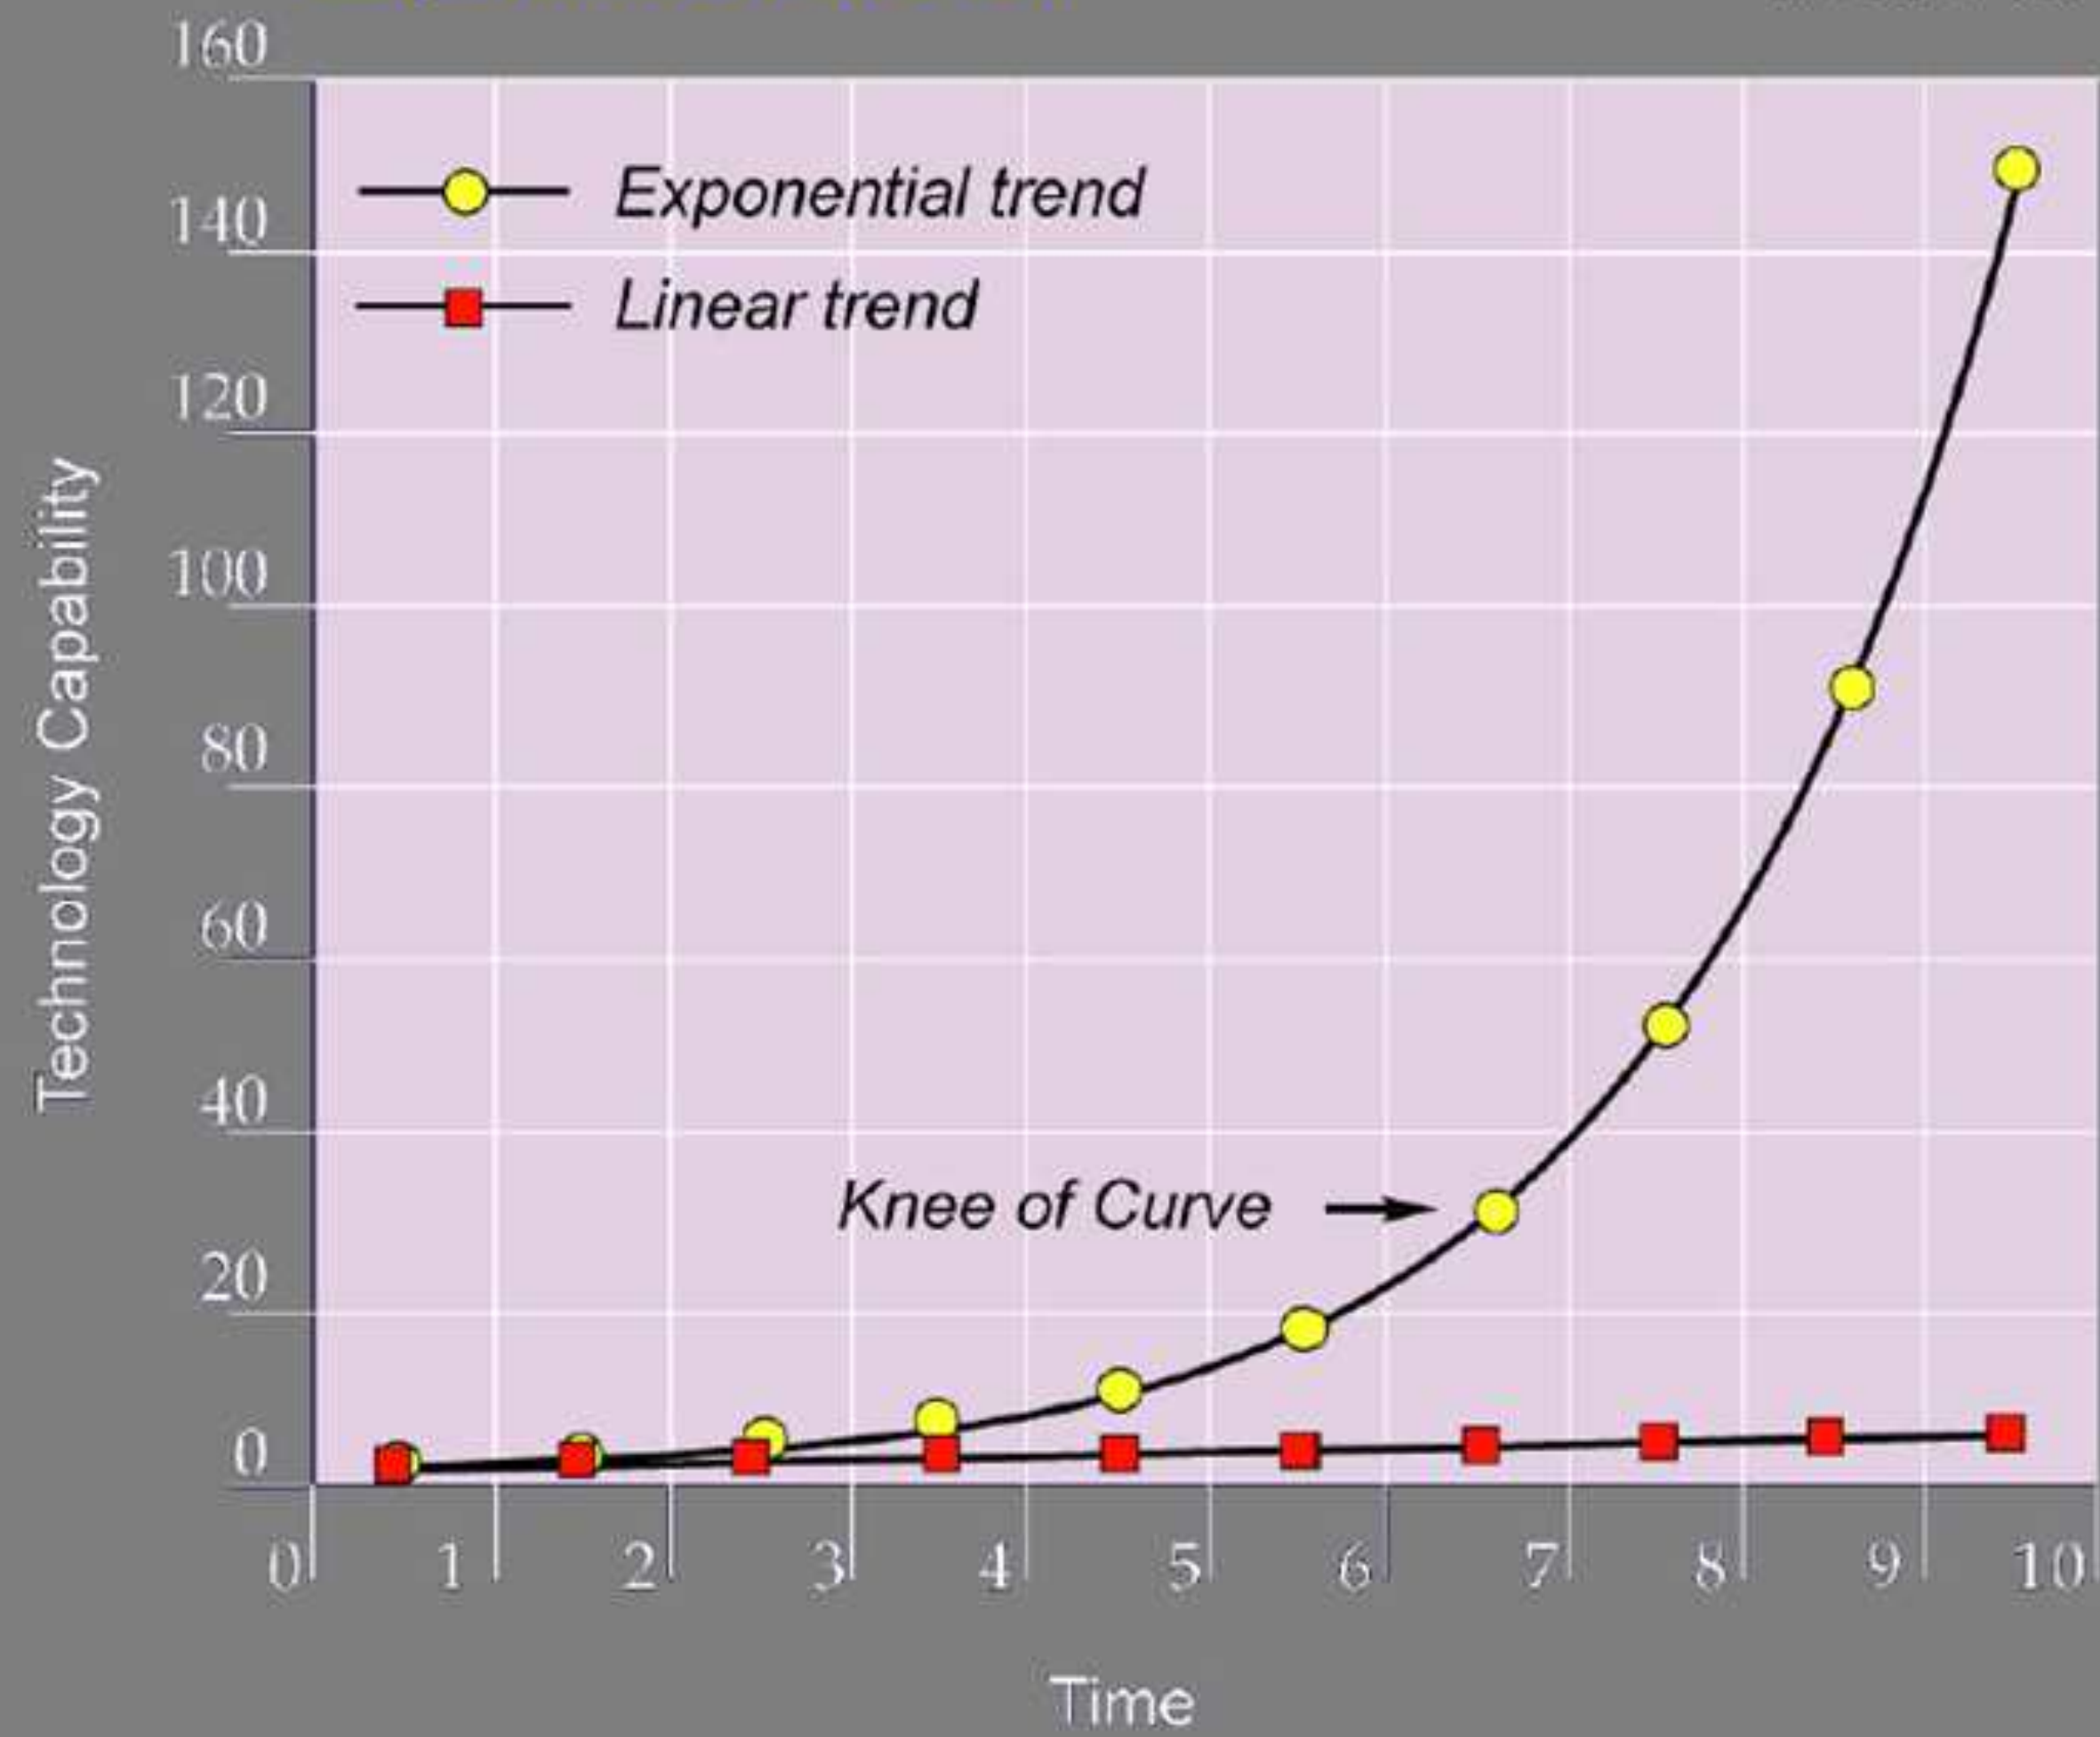
AI agents are great at clearly defined, structured tasks.

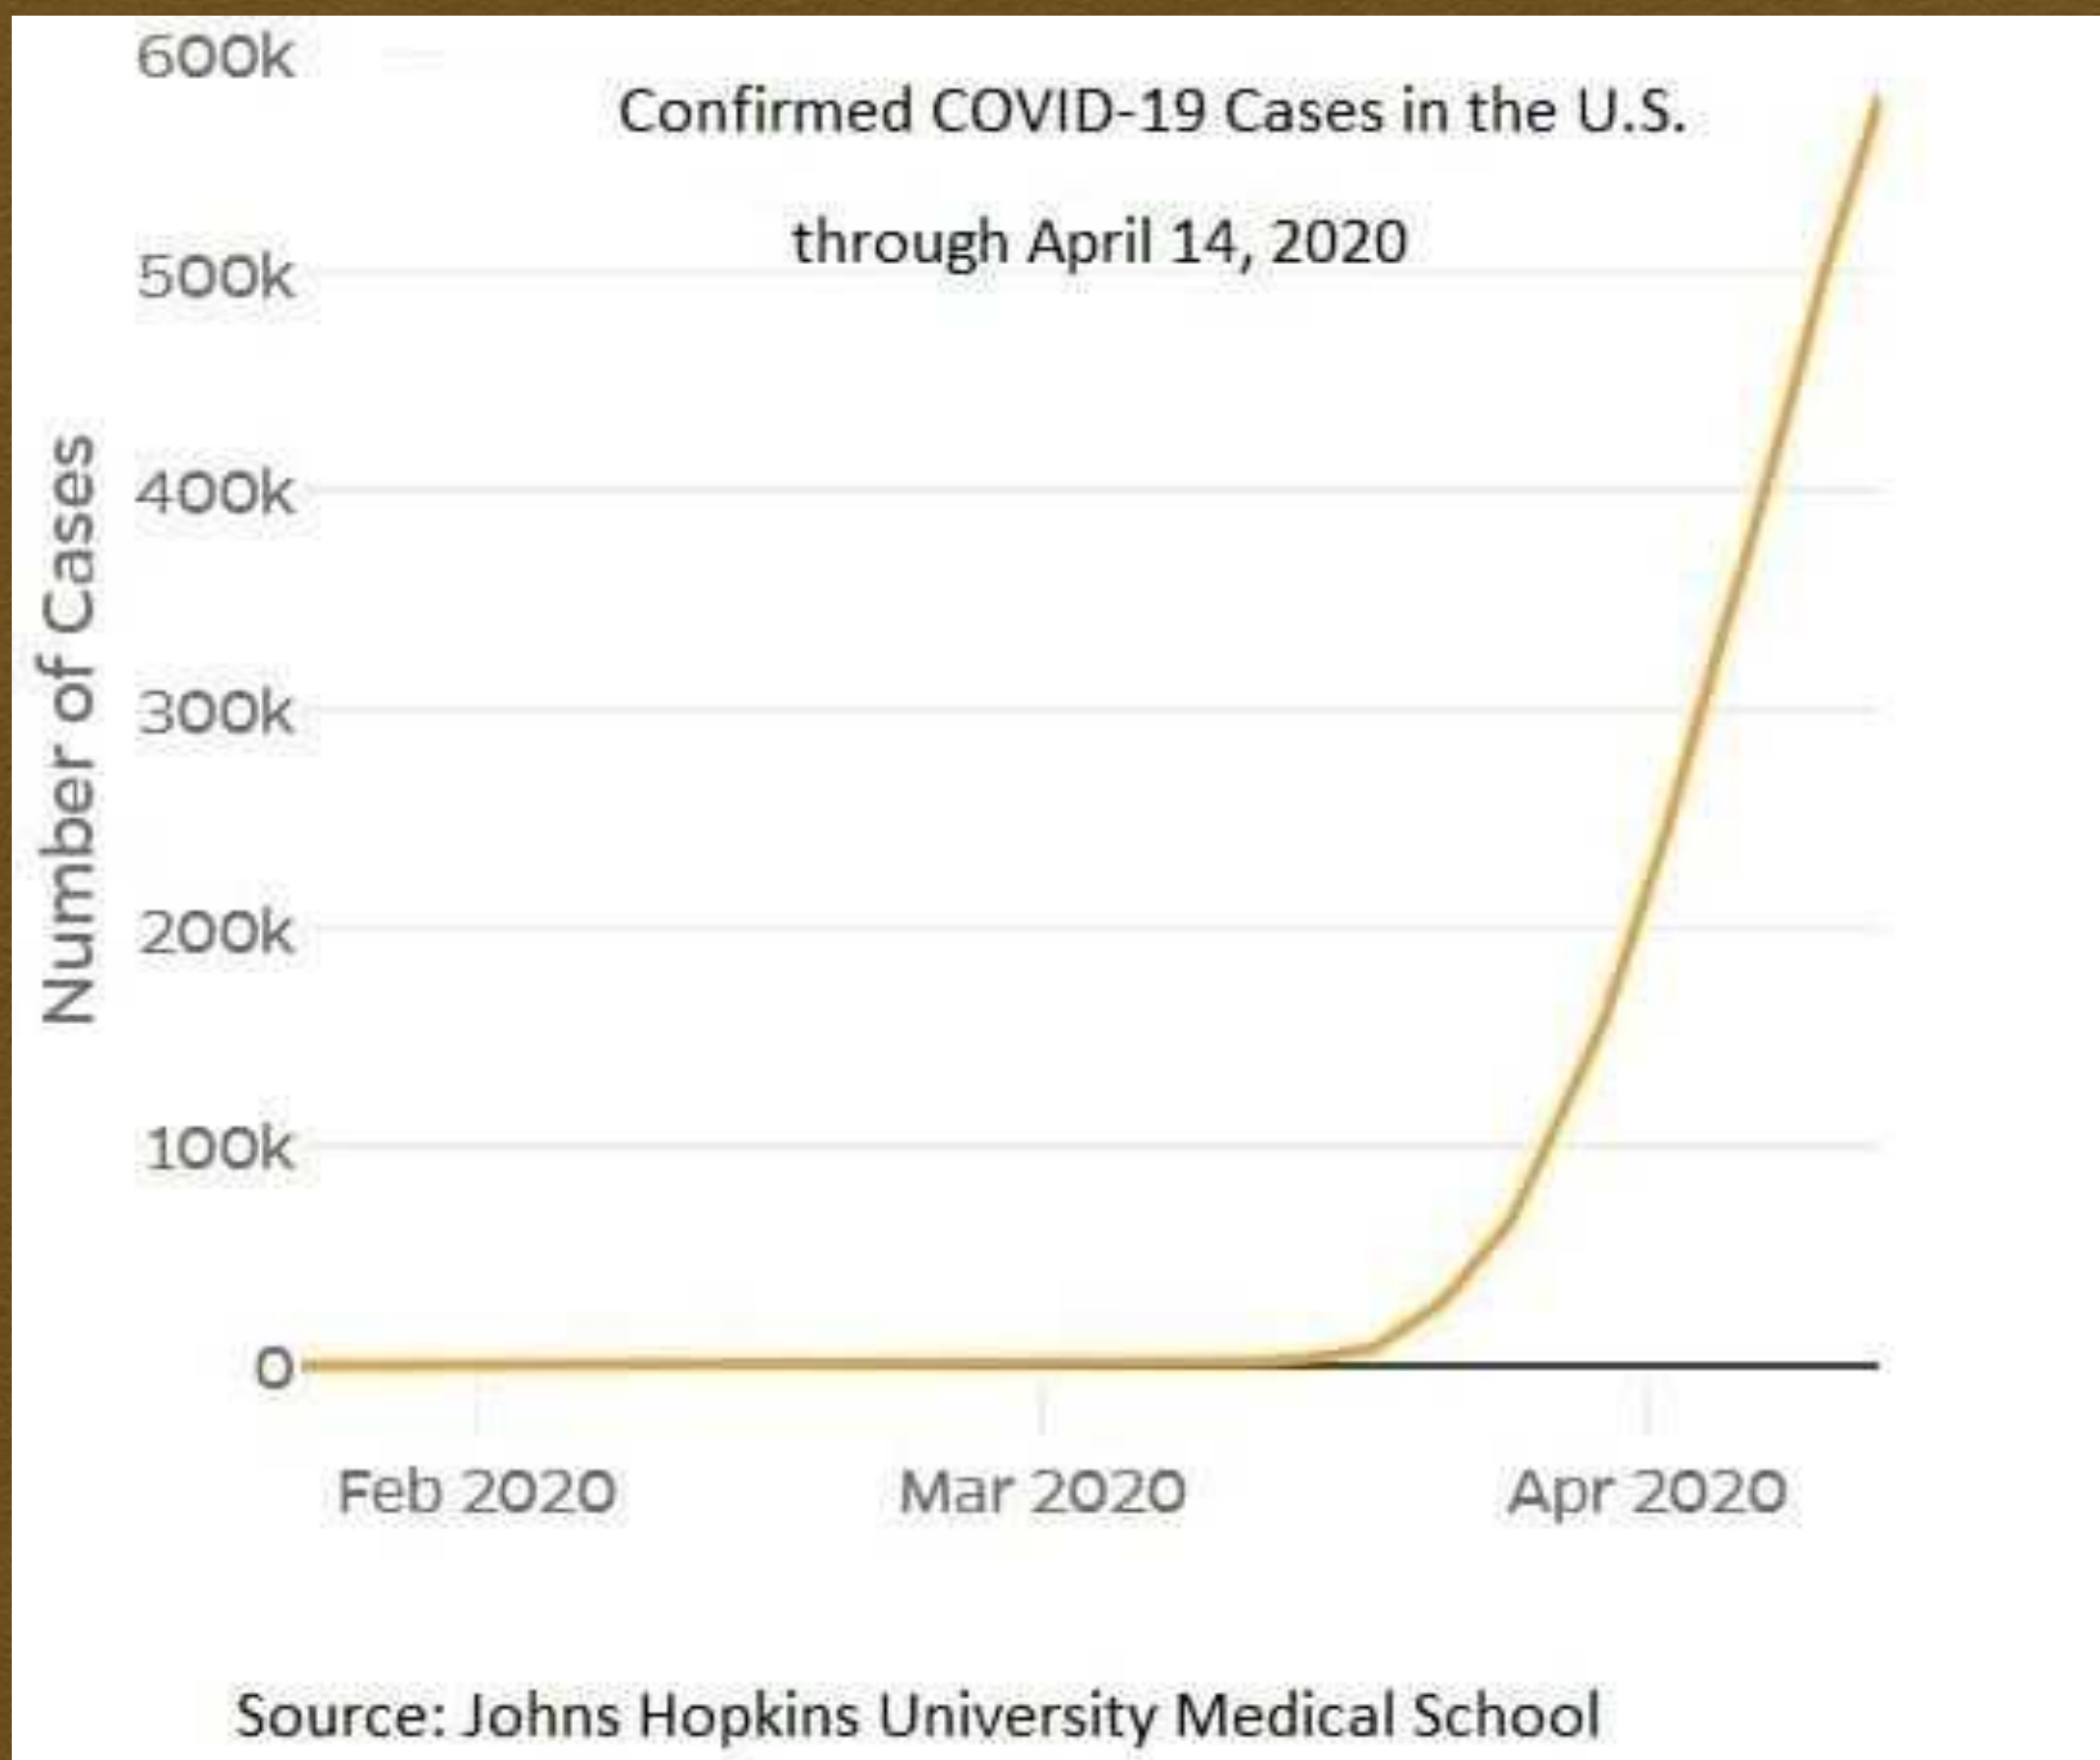
Think: intern.

But they're not yet ready for the judgment-heavy tasks of senior roles.

Linear vs. Exponential Growth

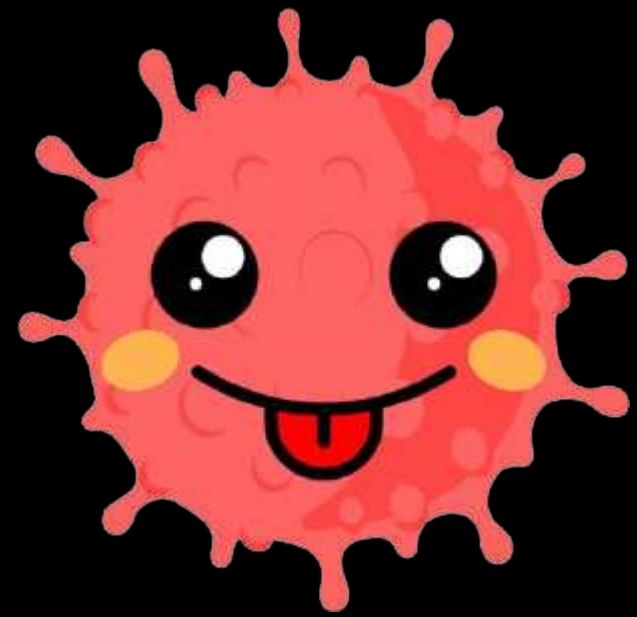
Linear Plot



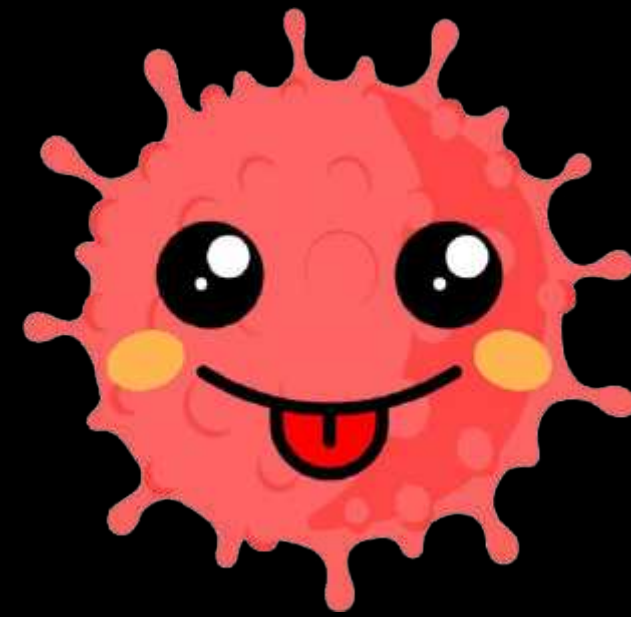


If one case of coronavirus doubles every three days...

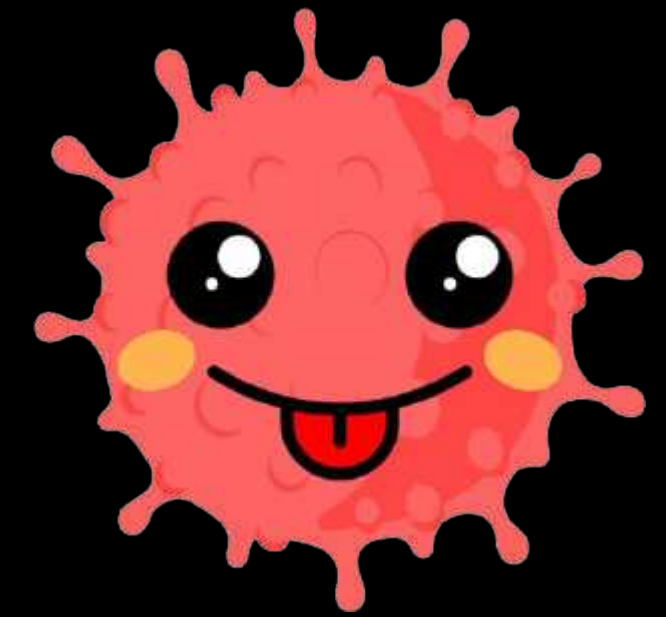
1



1000x



1,000,000x



Take 30 linear steps.



1, 2, 3, 4, 5, 6, 7, 8, 9, 10...

Take 30 EXPONENTIAL steps.



1, 2, 4, 8, 16, 32, 64, 128, 256, 512...

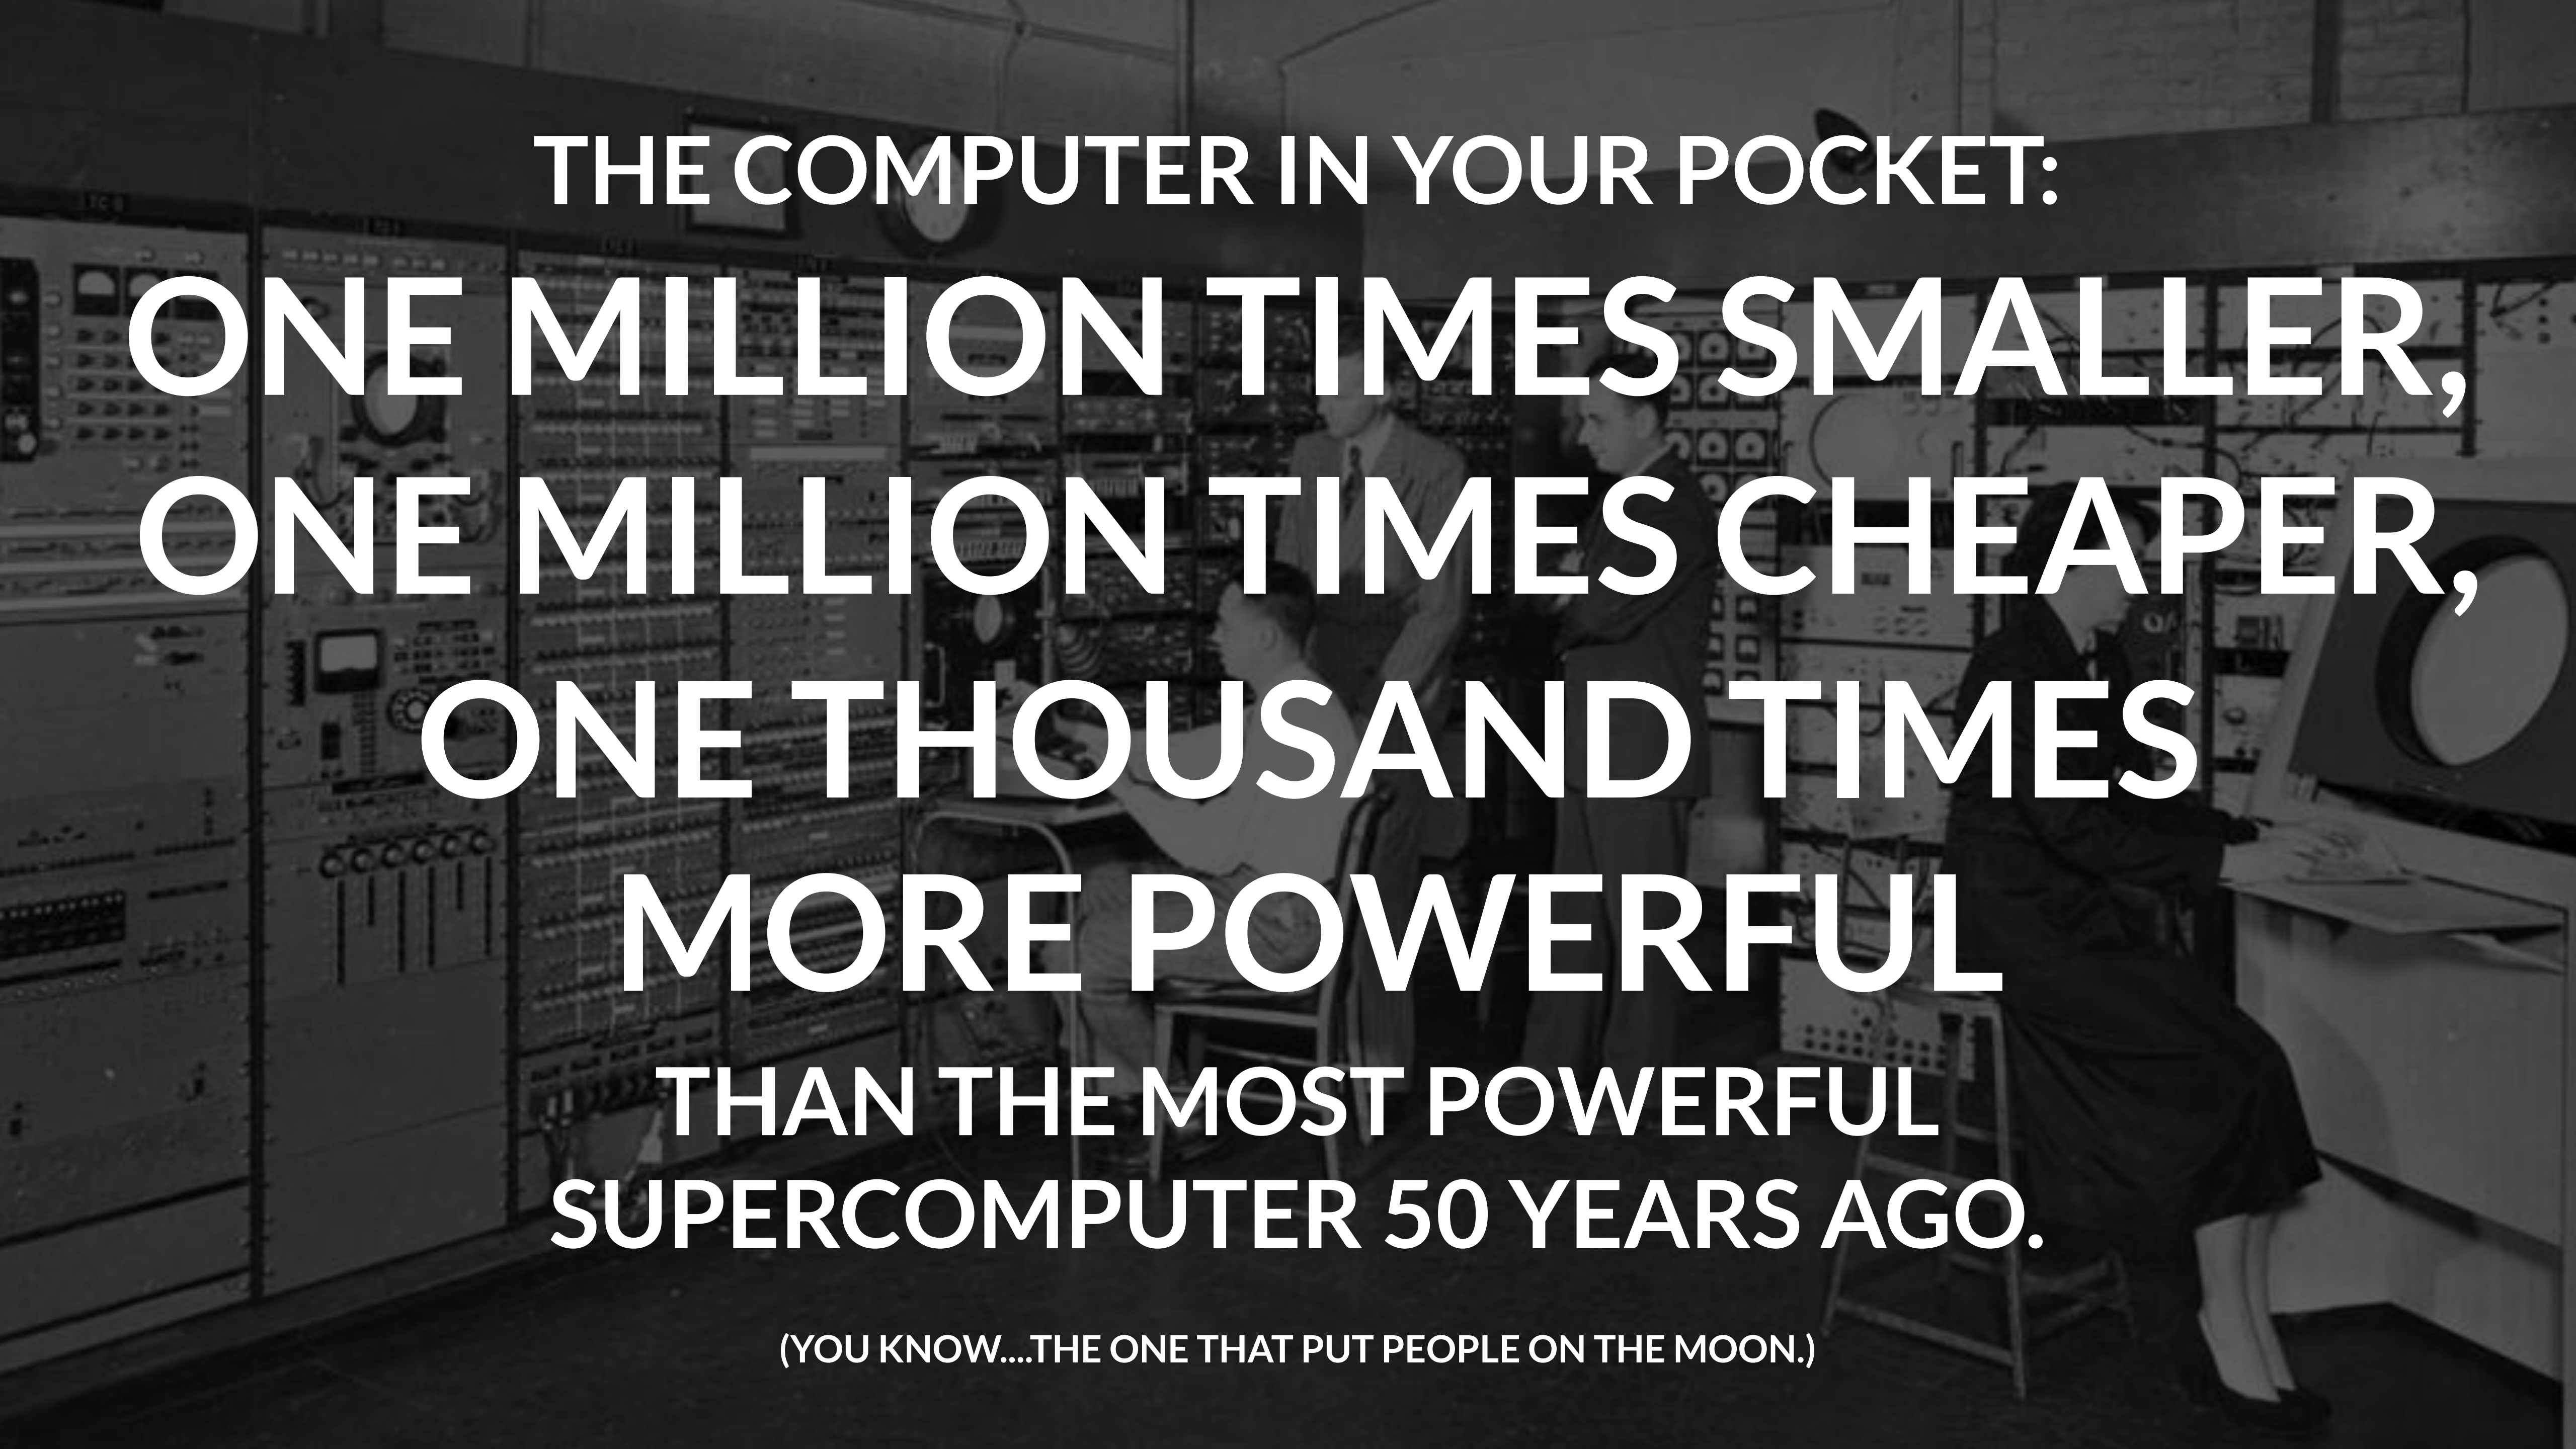
Take 30 EXPONENTIAL steps.

1,073,741,824 steps



Take 30 EXPONENTIAL steps.

1,073,741,824 steps
(You'd circle the world 26 times.)



THE COMPUTER IN YOUR POCKET:
**ONE MILLION TIMES SMALLER,
ONE MILLION TIMES CHEAPER,
ONE THOUSAND TIMES
MORE POWERFUL
THAN THE MOST POWERFUL
SUPERCOMPUTER 50 YEARS AGO.**

(YOU KNOW...THE ONE THAT PUT PEOPLE ON THE MOON.)

LINEAR → EXPONENTIAL

LINEAR → EXPONENTIAL



1996
Market Cap: \$28B
Employees: 140,000

LINEAR → EXPONENTIAL



1996
Market Cap: \$28B
Employees: 140,000



2012
Bankrupt
Employees: 17,000

LINEAR → EXPONENTIAL



1996
Market Cap: \$28B
Employees: 140,000



2012
Bankrupt
Employees: 17,000



April 2012
Market Cap: \$1B
Employees: 13

Let's talk.



Getting Started: AI Strategy Jumpstart

How do you start developing an AI strategy for your organization?

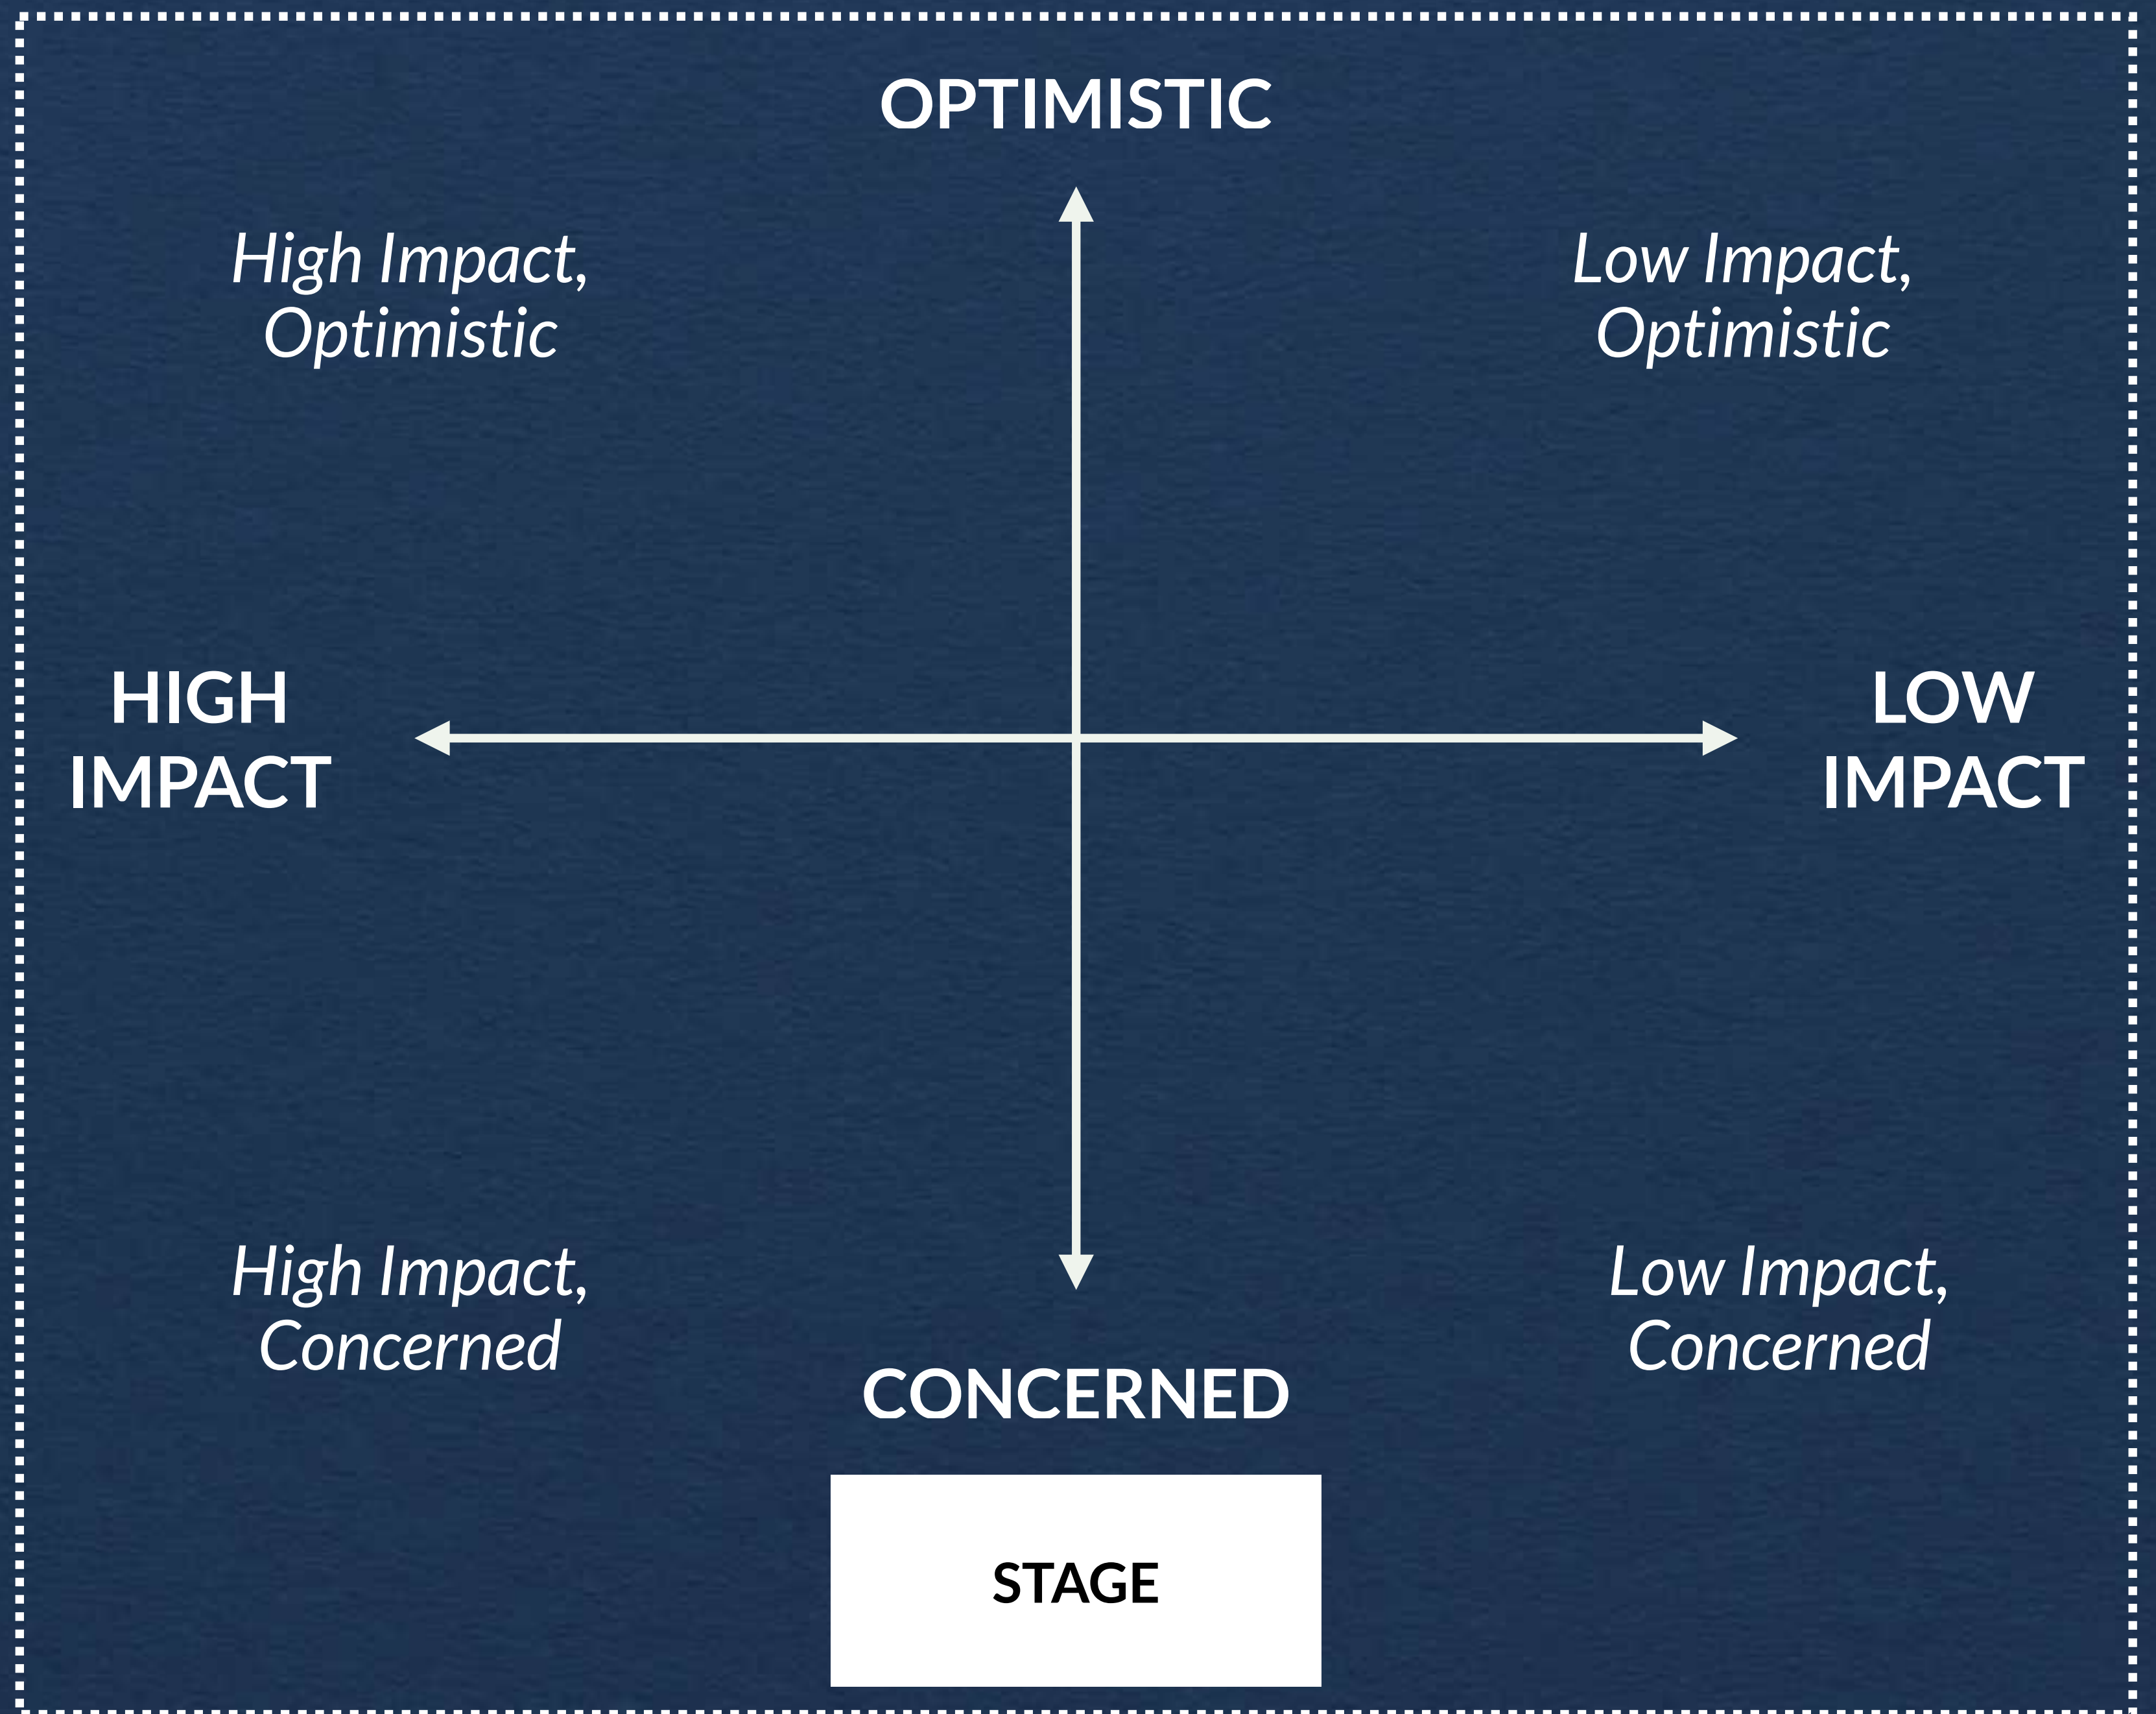
*(one that contends with the unique challenges and nuance
of generative and agentic AI)*

How do you start developing an AI strategy for your organization?

In the next 15 minutes, we're going to experience a simplified version of a strategic framework.

1. SURFACE PERSPECTIVES

Stand up!
Answer these
questions with
your bodies.

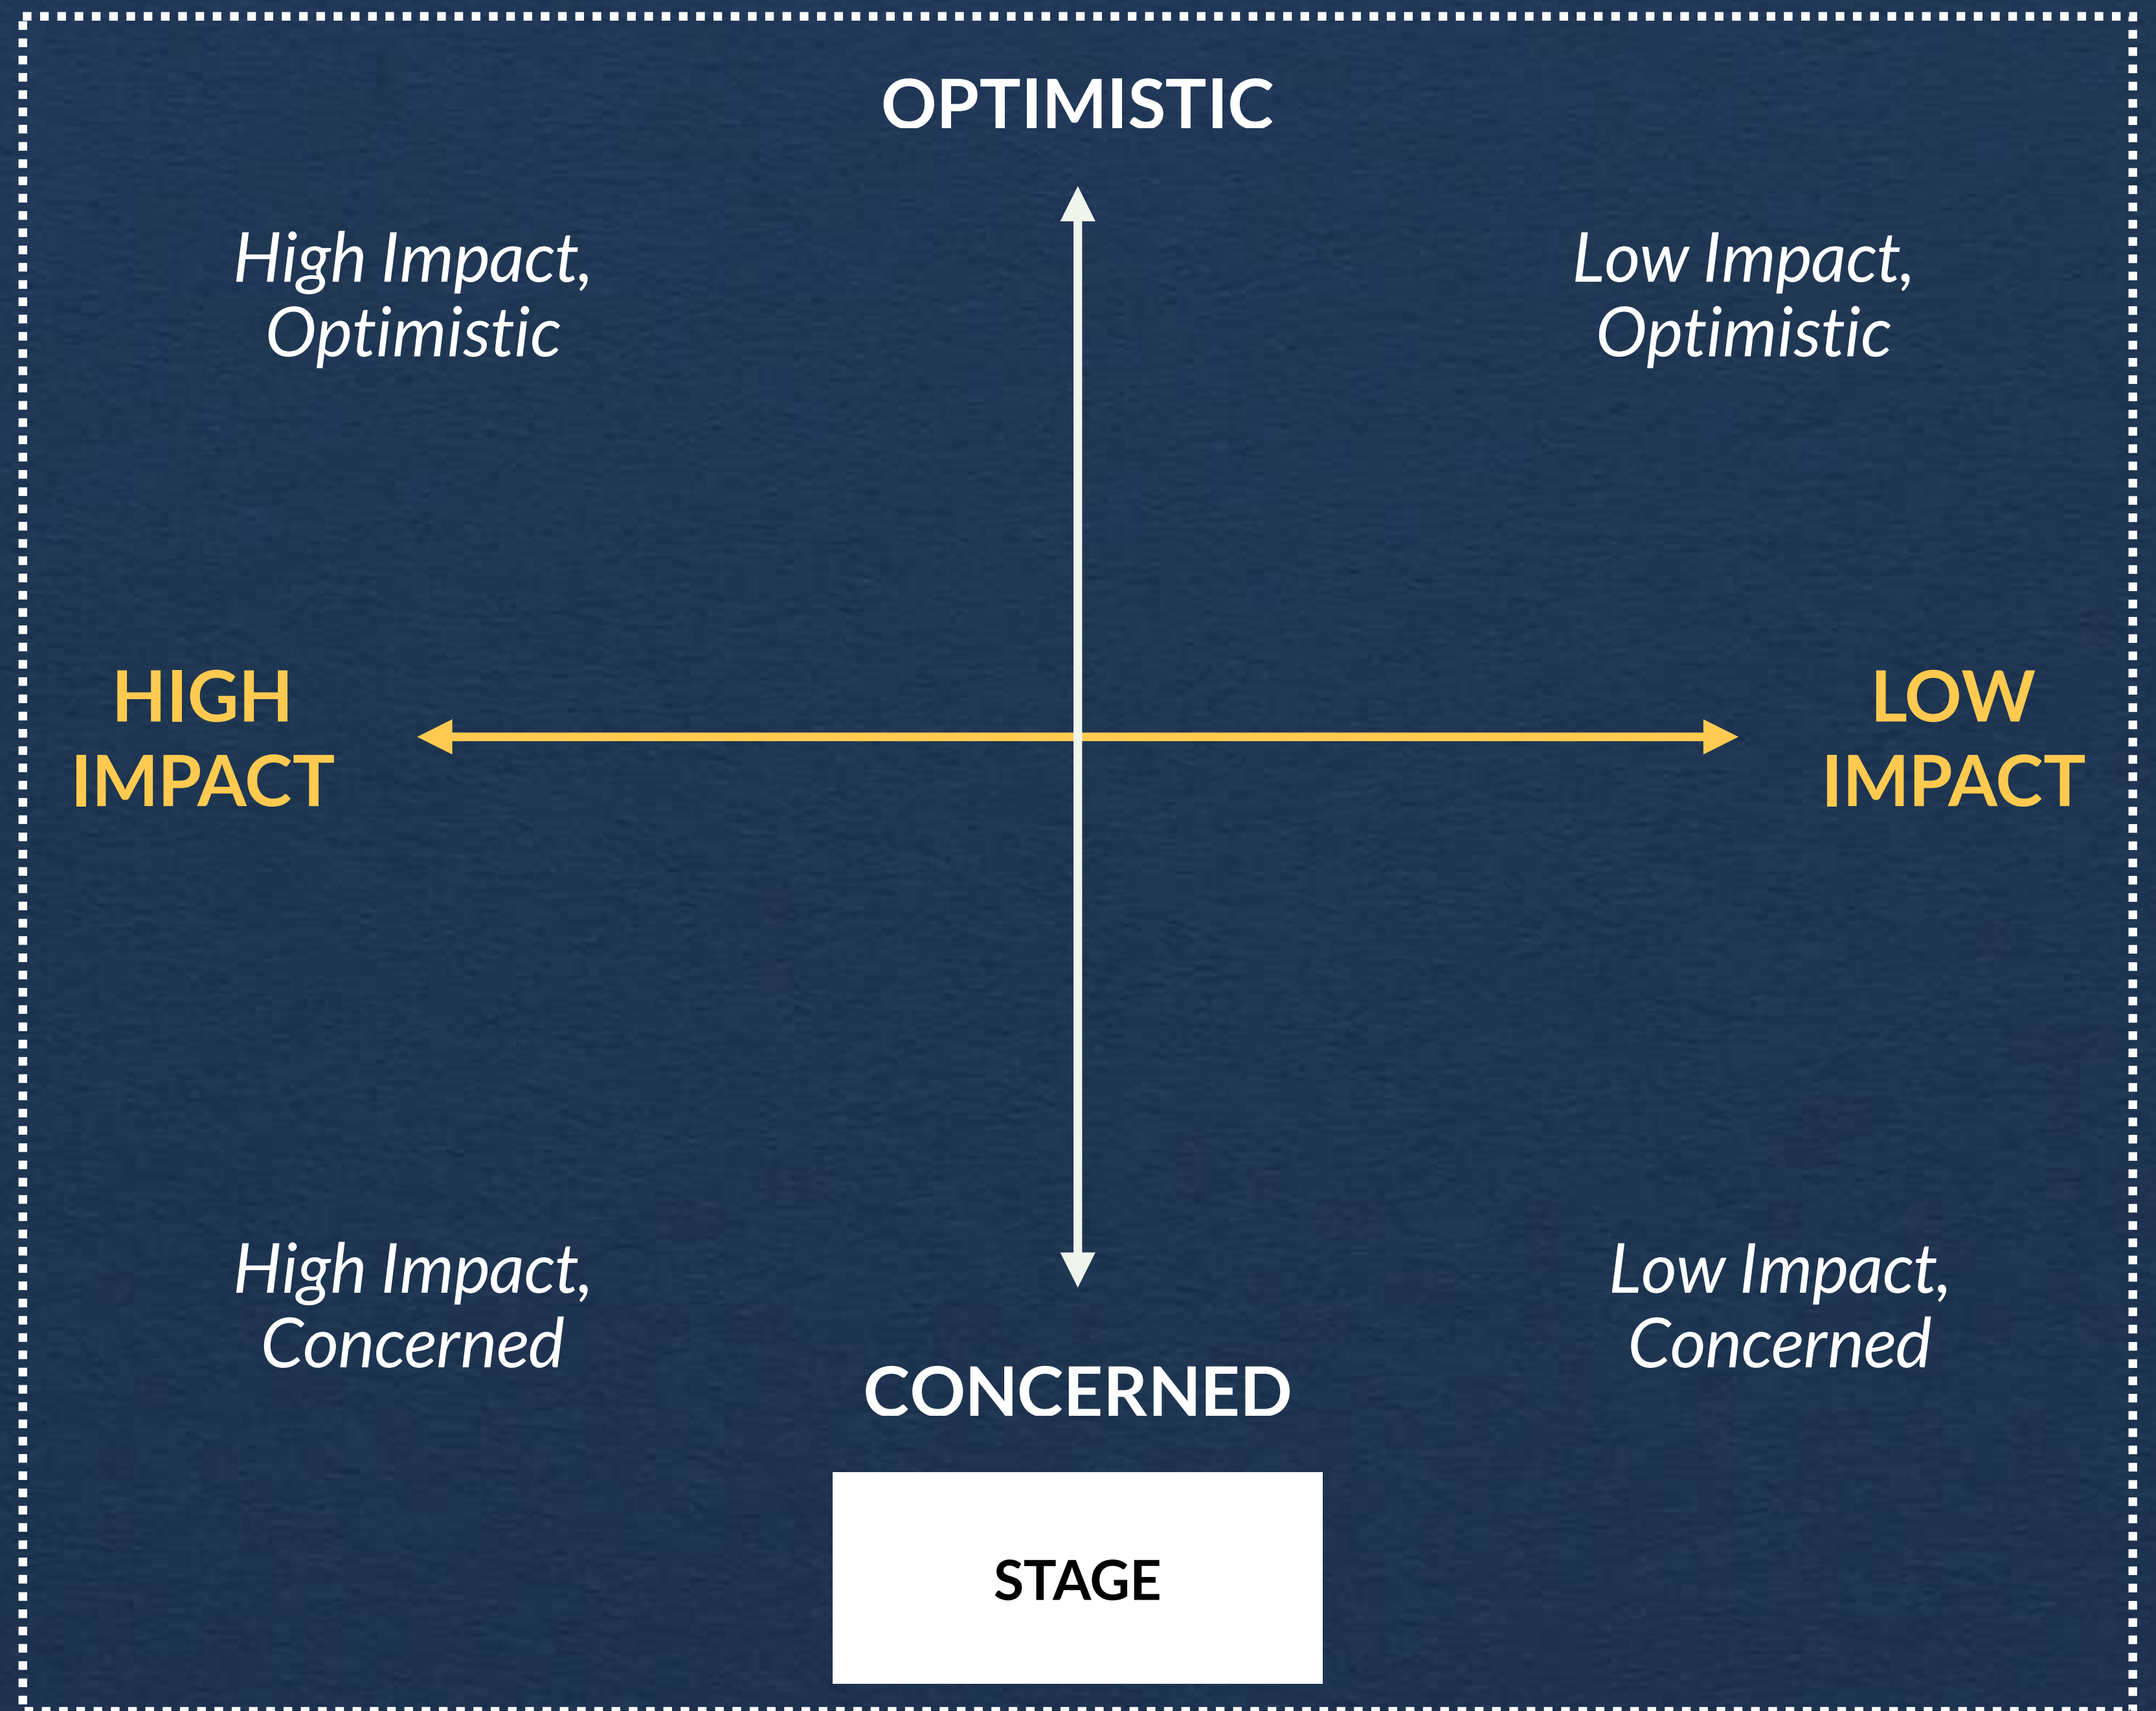


1.

How significant do you think Generative and Agentic AI's impact will be on your specific role or organization in the next 1-2 years?

2.

How optimistic or concerned are you about that impact?

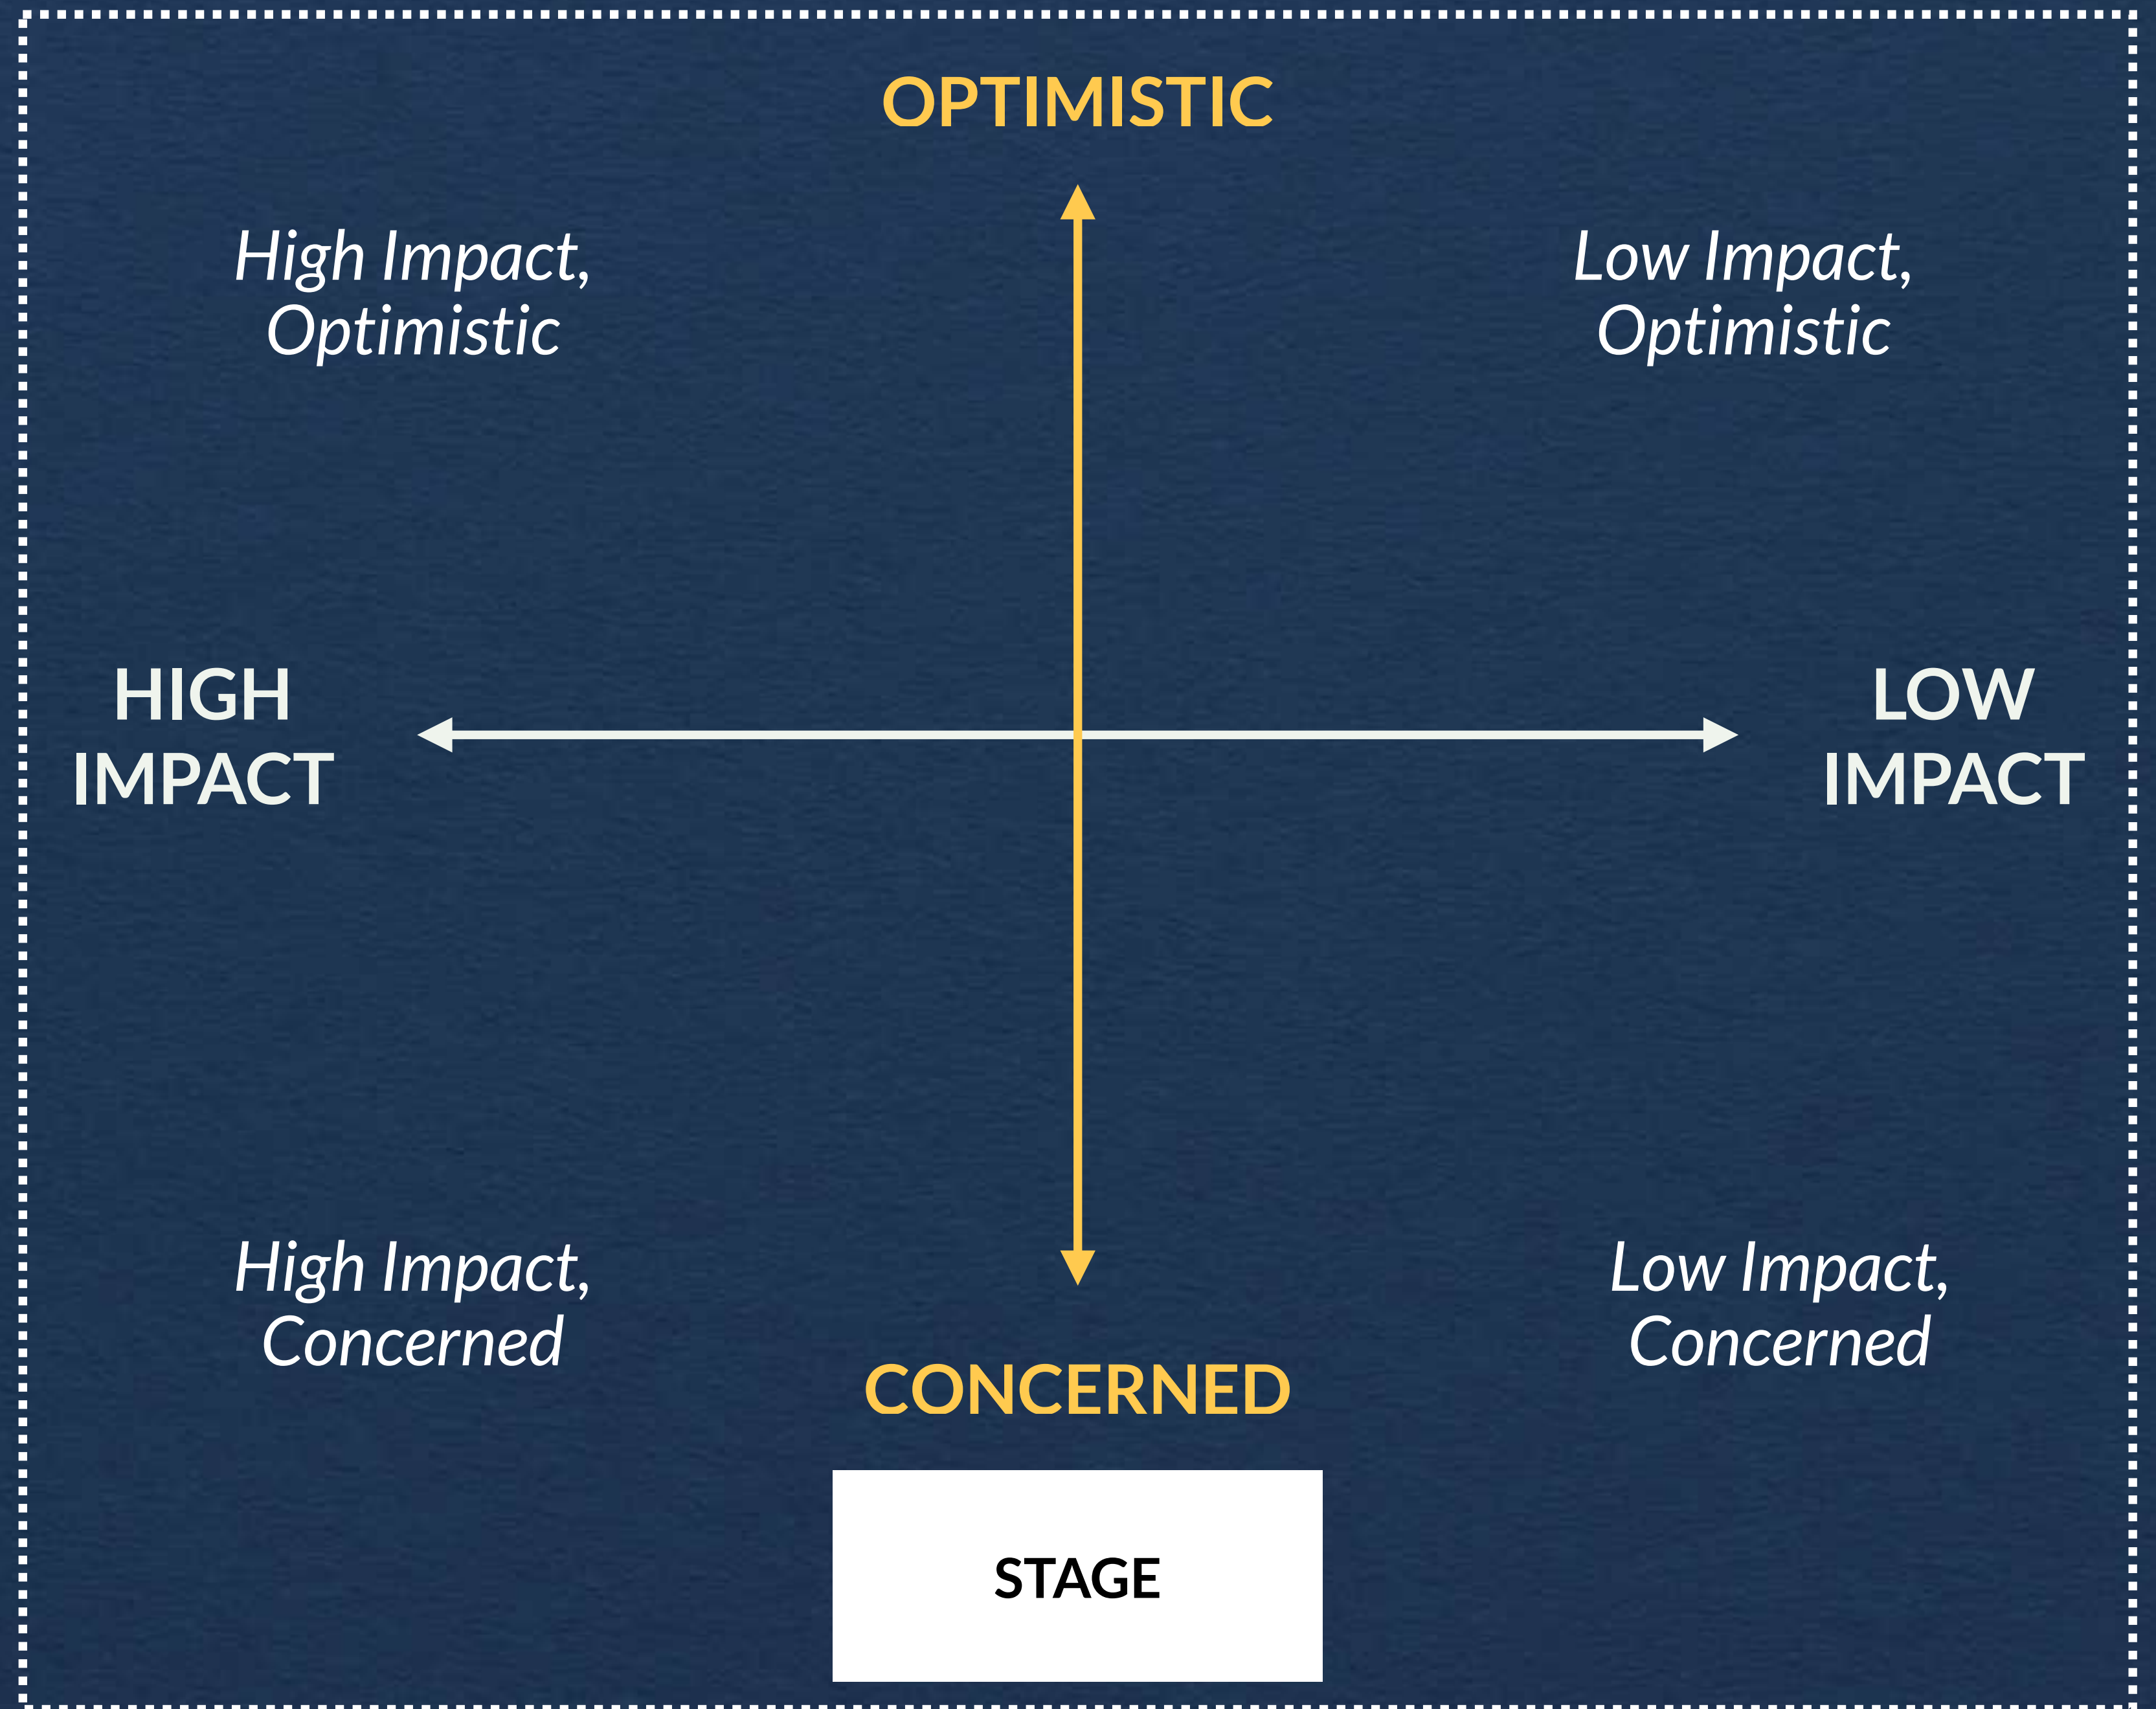


1.

How significant do you think Generative and Agentic AI's impact will be on your specific role or organization in the next 1-2 years?

2.

How optimistic or concerned are you about that impact?

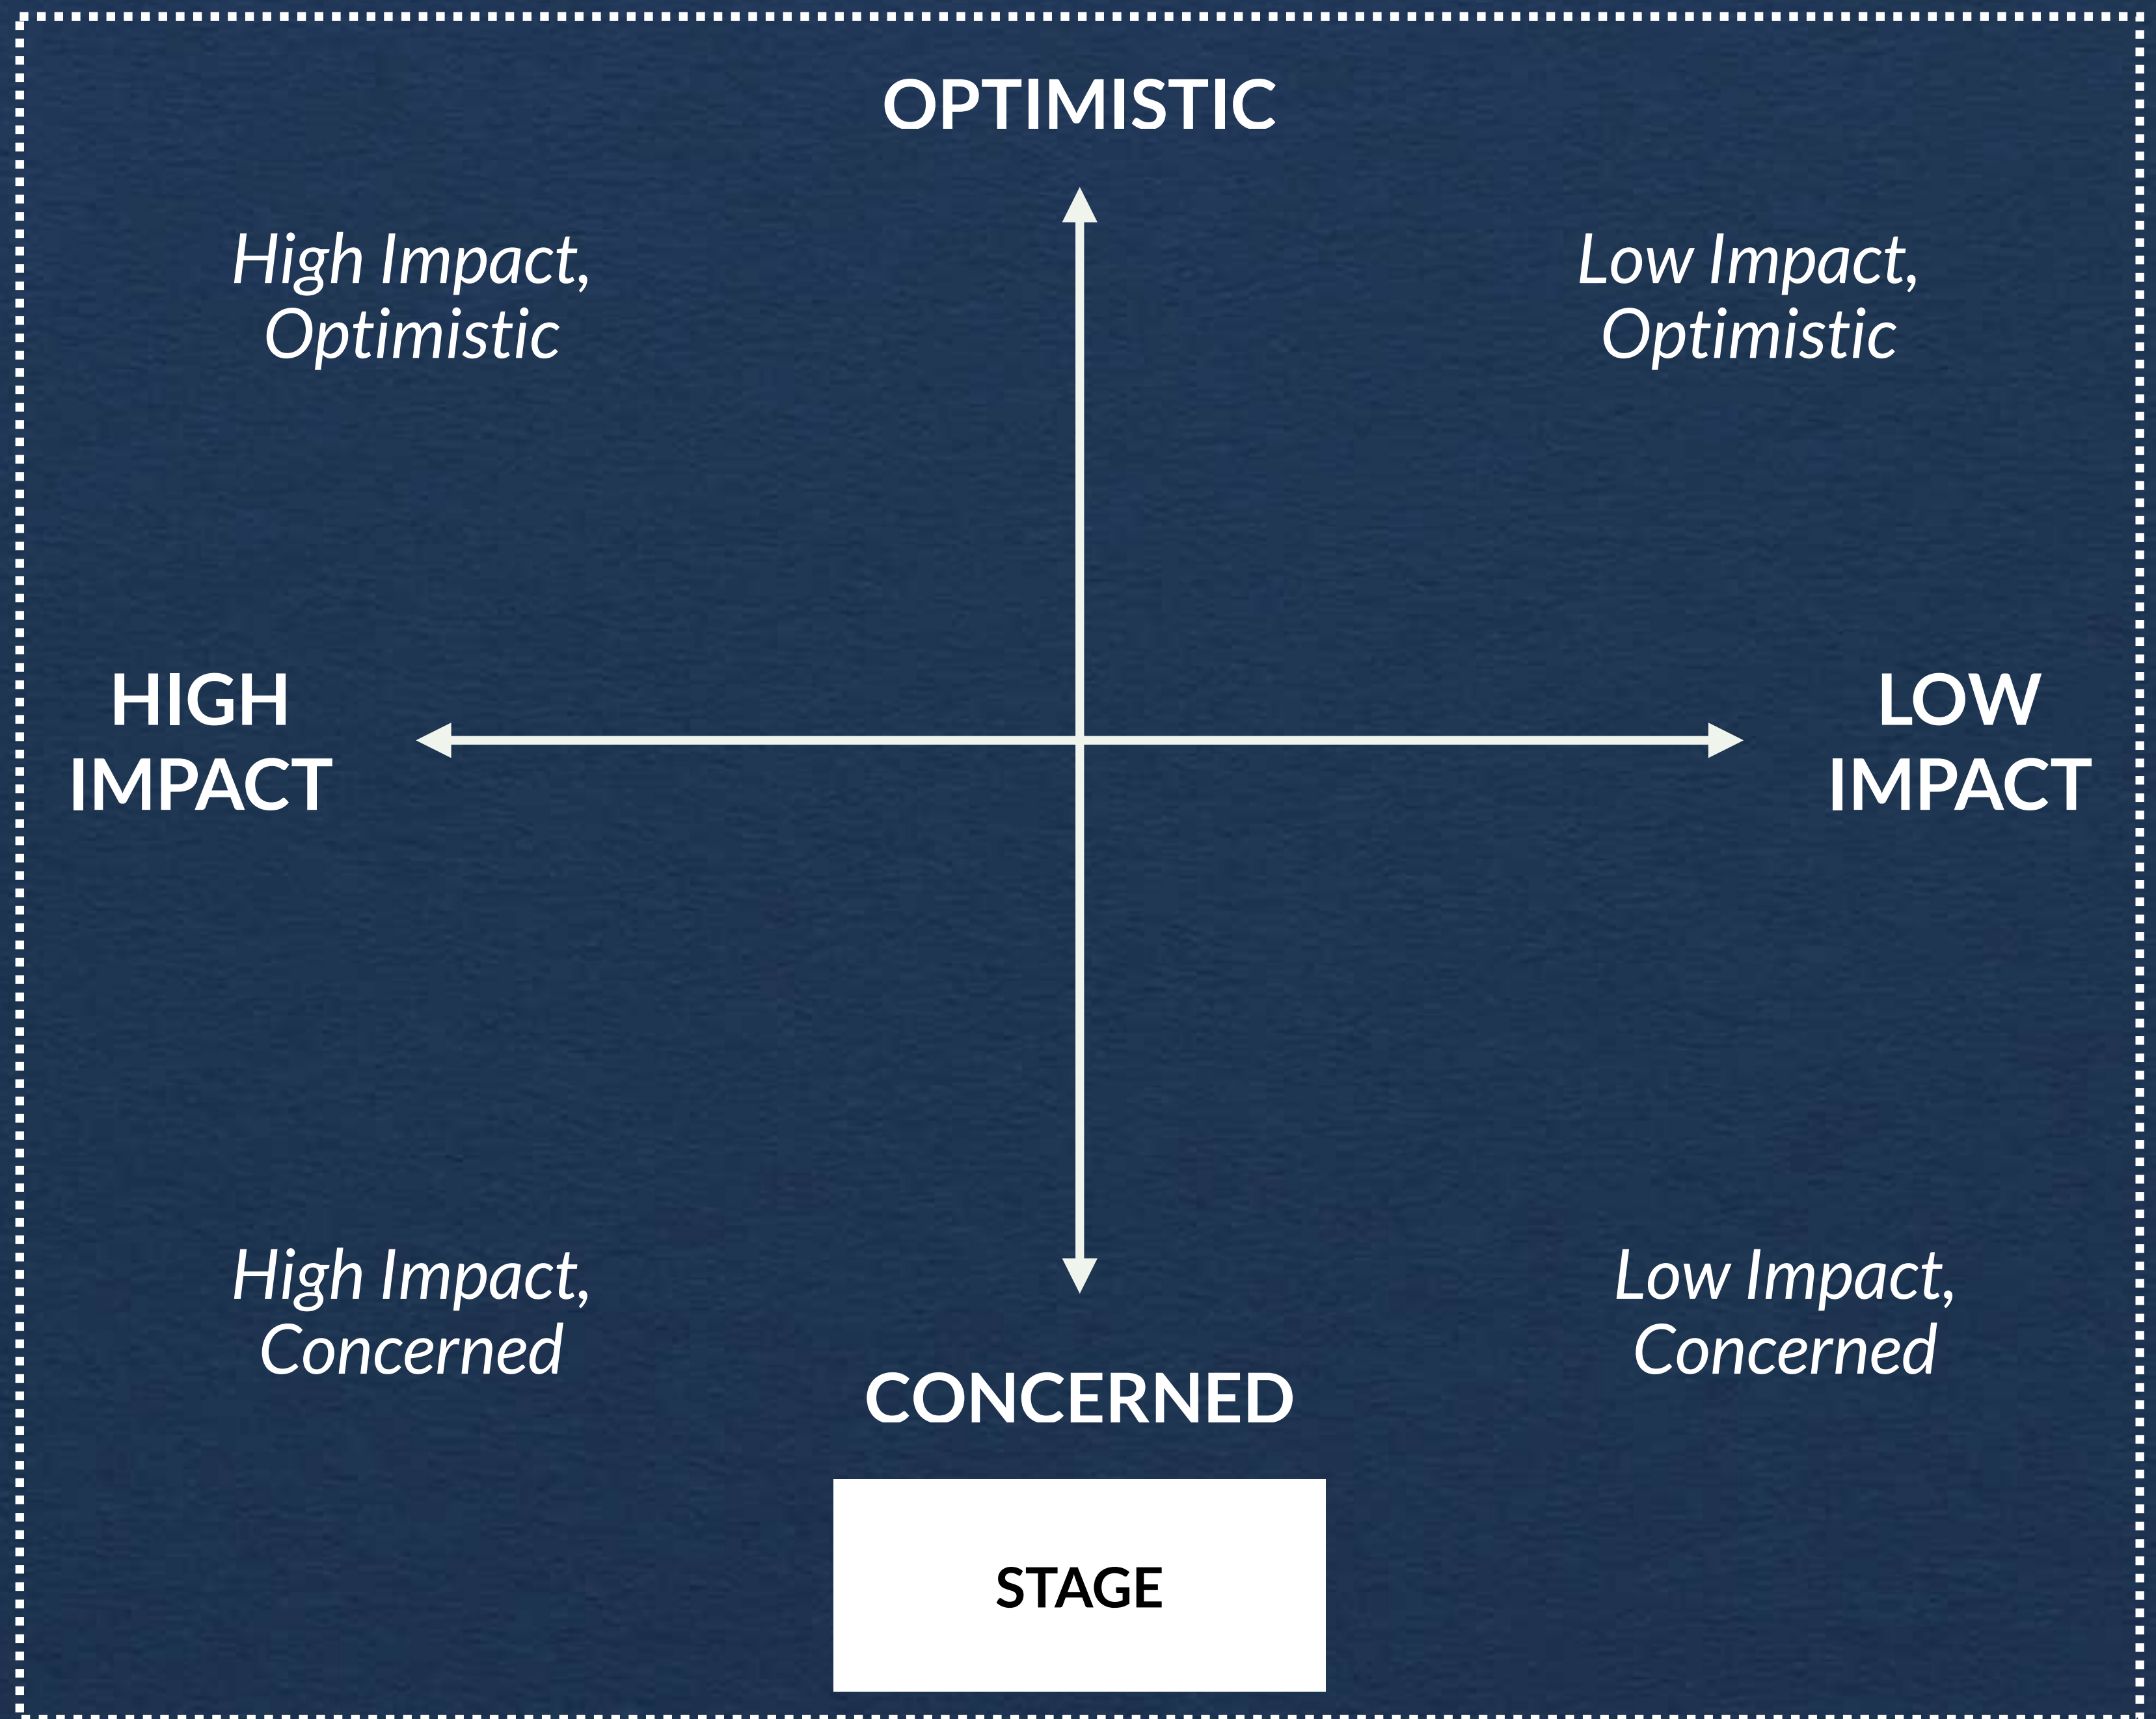


Look around the room.

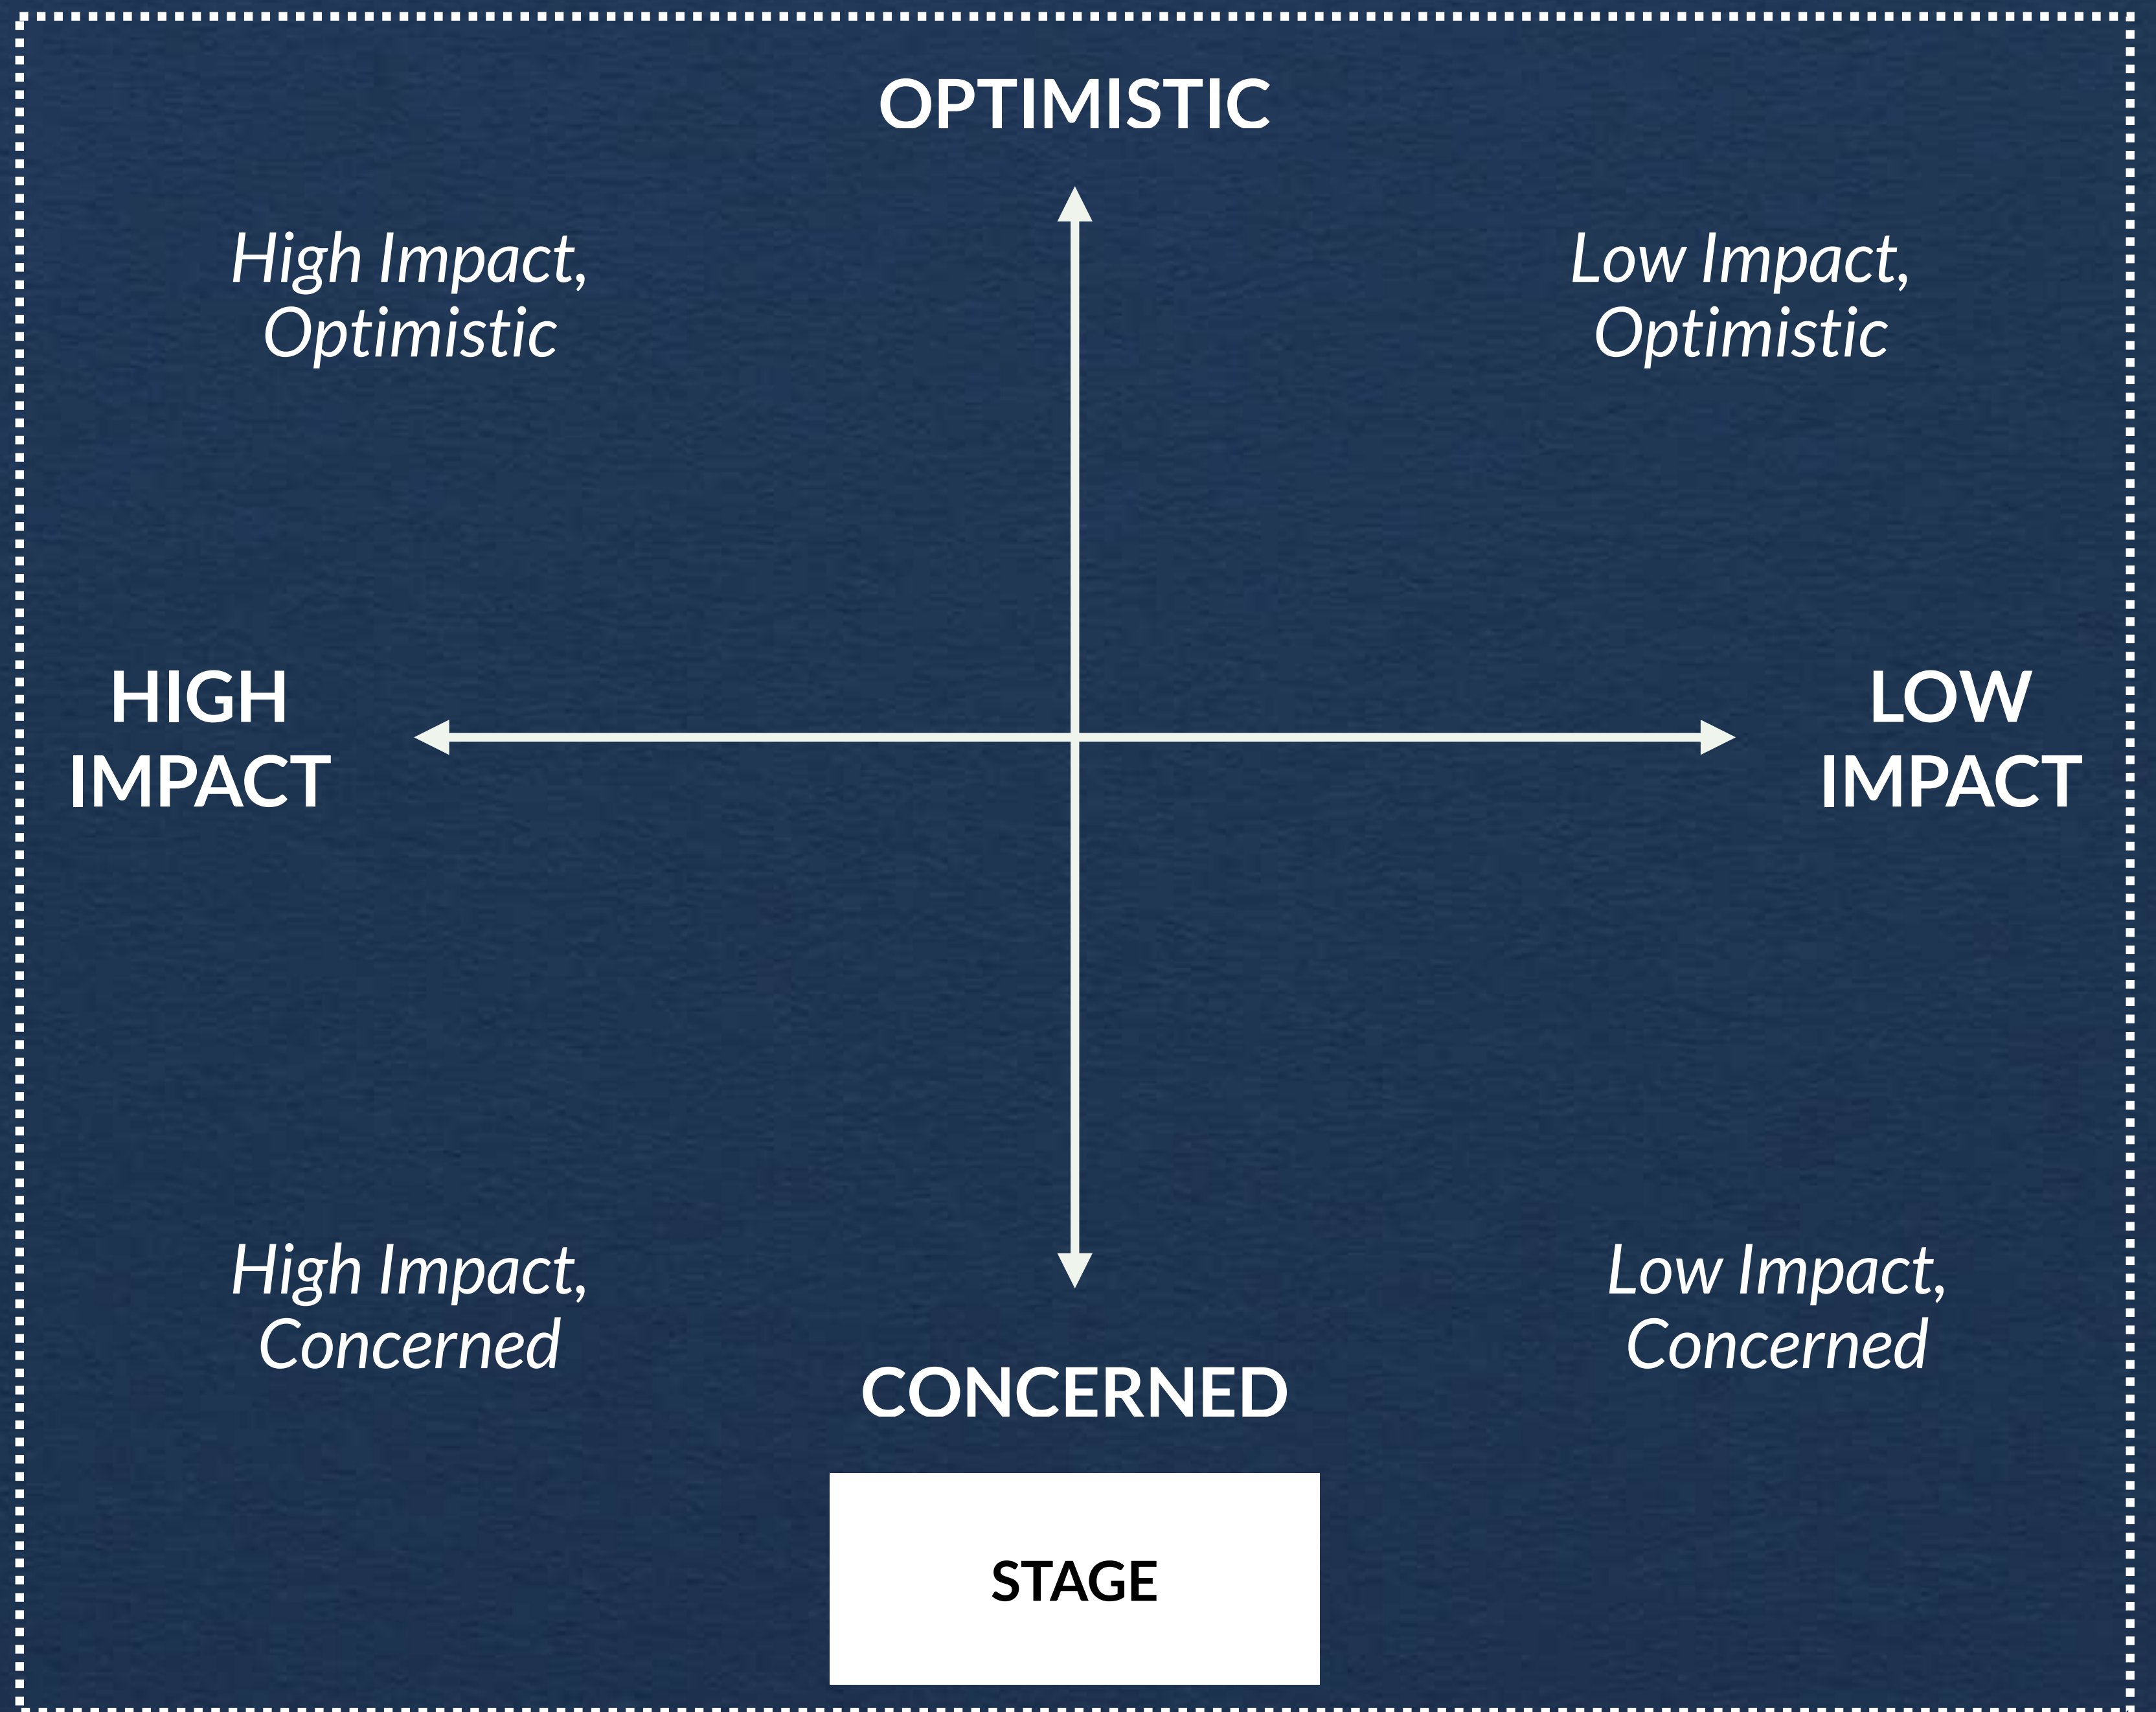
**Notice where your
colleagues are
standing.**

**Notice where
leadership is standing.**

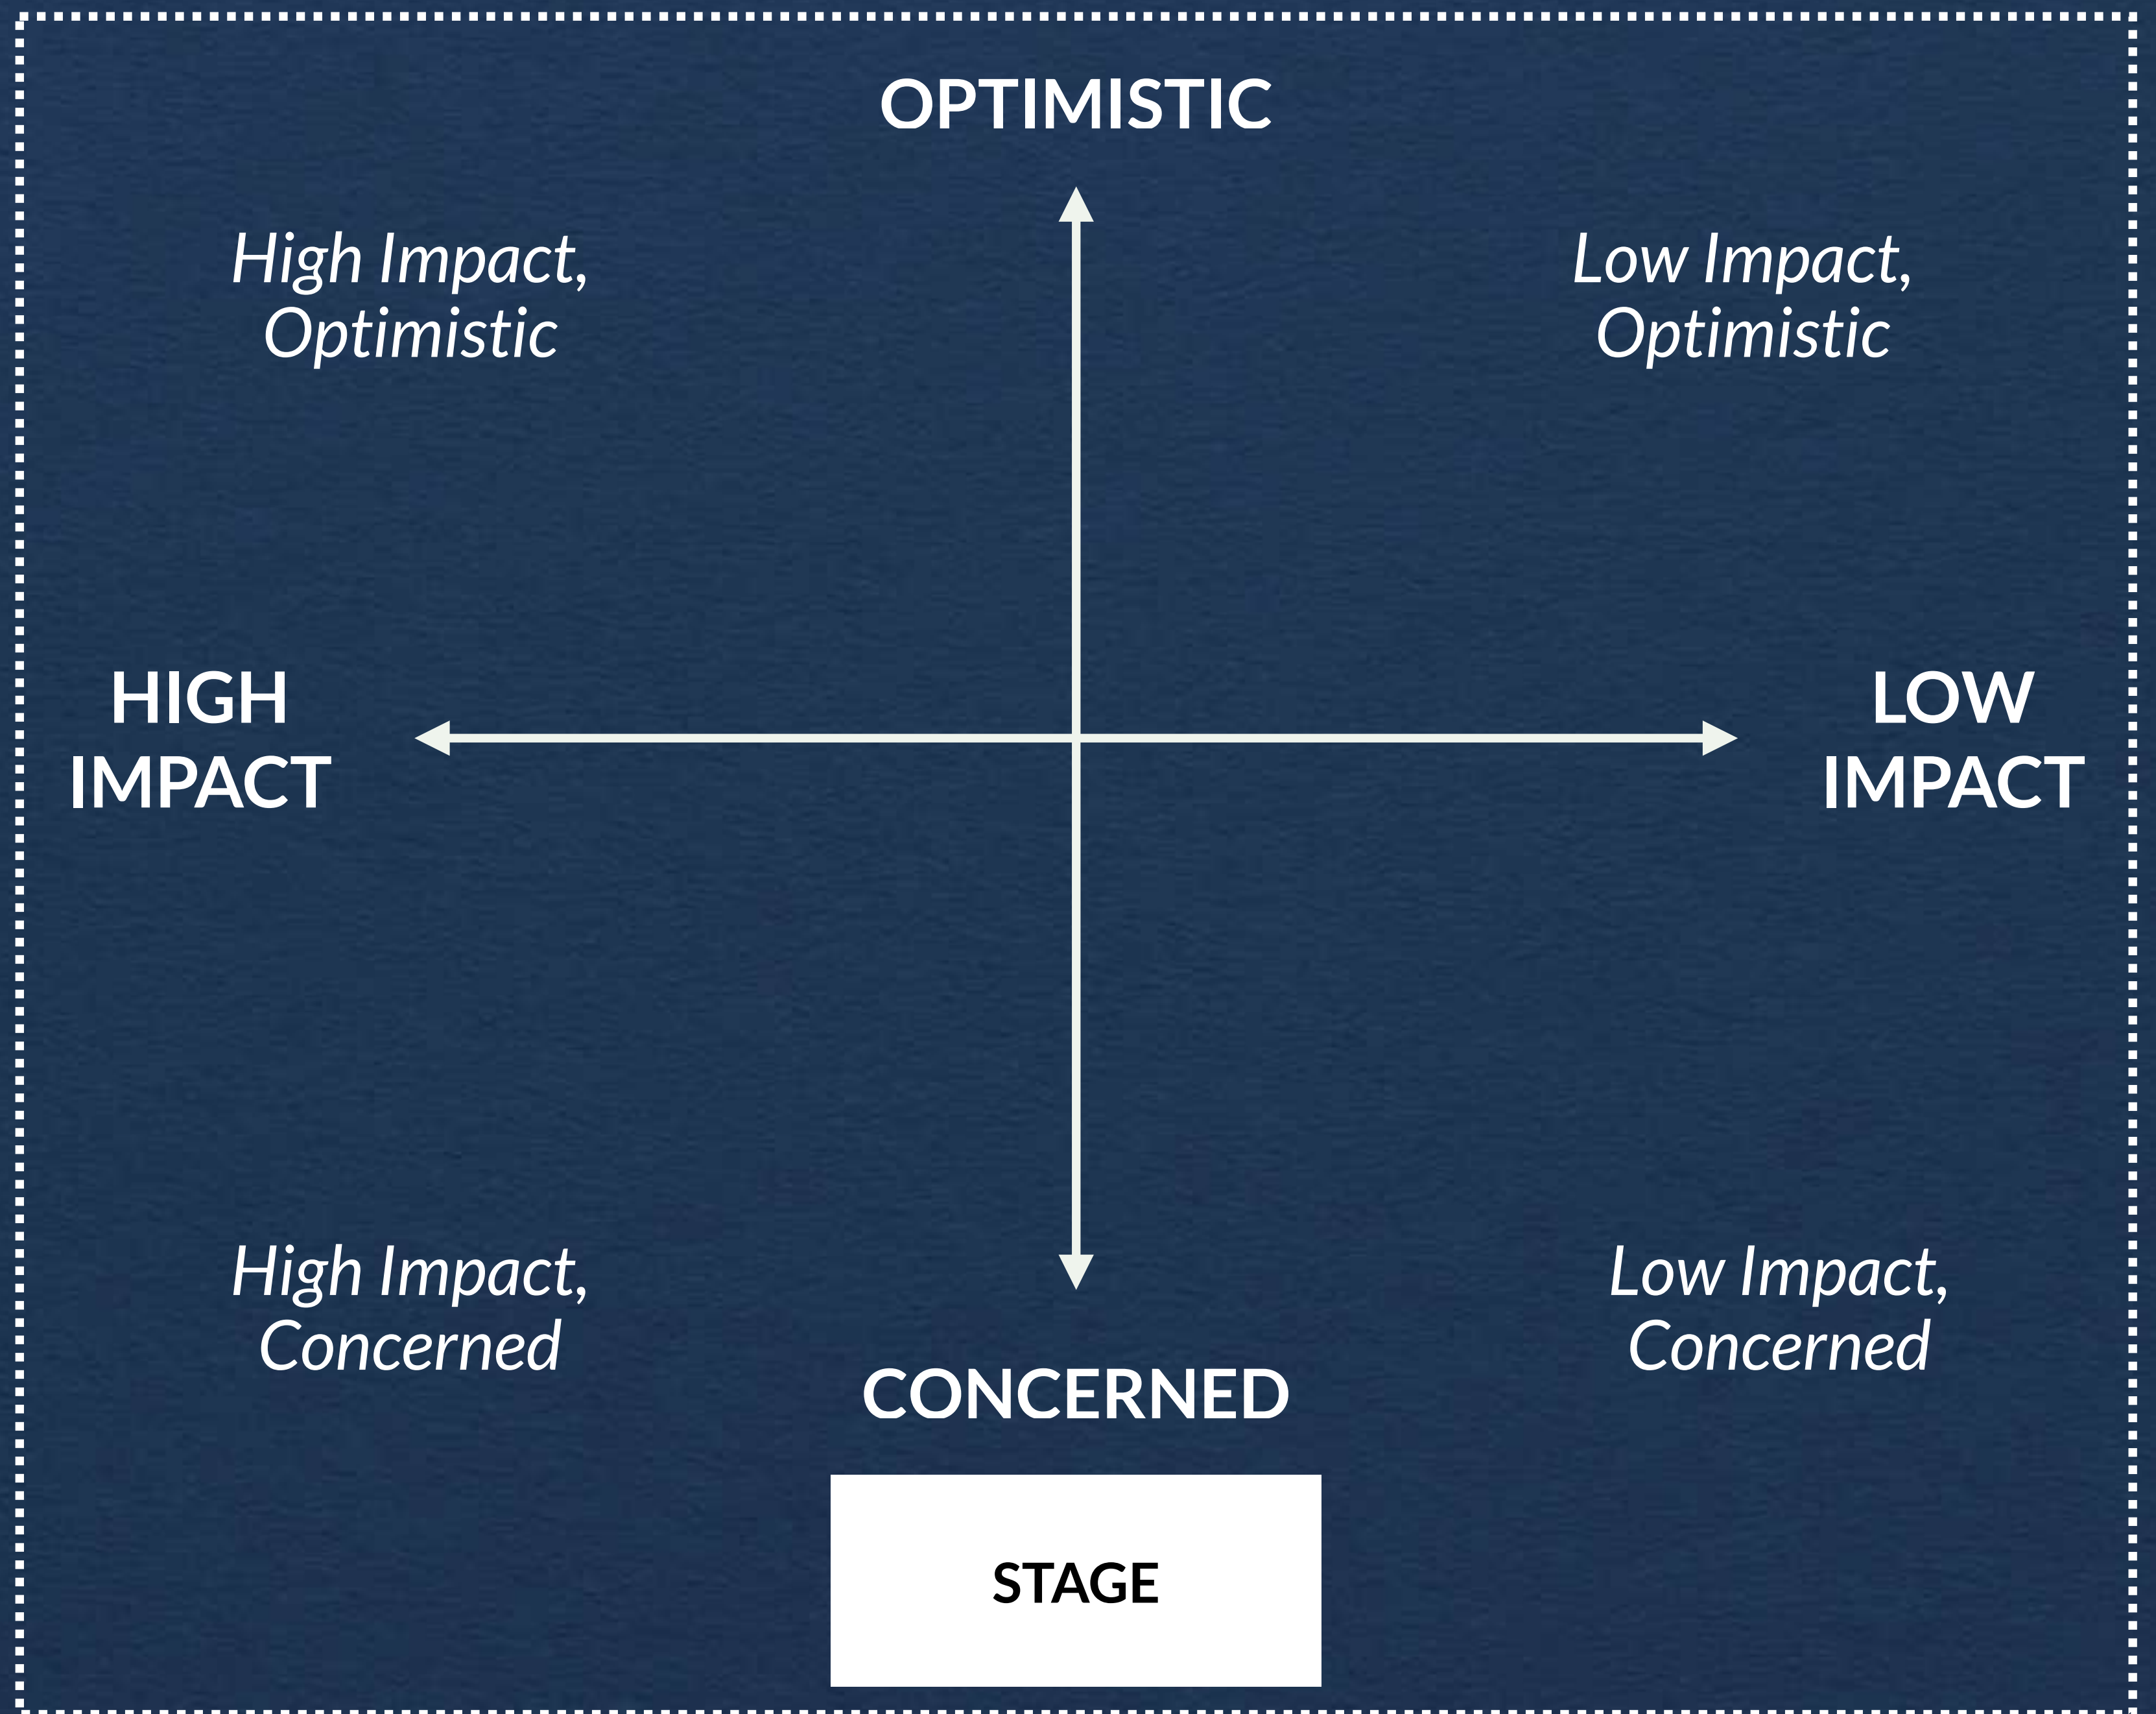
**Notice where member/
customer-facing roles
are standing.**



**Who's willing to
share your
perspective in a
sentence or two?**



This diversity of perspectives exists across your organization too.



**An effective AI strategy
begins by acknowledging
these different viewpoints.**

*The first step is creating space for open
dialogue across positions and departments.*

2. ESTABLISH PRINCIPLES

1. Now: Look around.
2. Find two people standing near you who you don't normally work with directly and form a group of three.
3. Grab a seat anywhere nearby - just make sure you can see the screen.
4. You'll stay in this groups for the next activity

HUMAN CENTEREDNESS

AI should aim to do good, actively contributing to the welfare of individuals and society.

CONTROL & ACCOUNTABILITY

Humans should have significant control over AI systems and be accountable for their outcomes and potential harms.

OBSERVABILITY & EXPLAINABILITY

AI systems should be both observable and explainable, allowing humans to understand and monitor their decision-making processes.

IMPROVEMENT & RESPONSIVENESS

AI systems should be continuously improved to enhance their performance and address any emerging issues.

In your groups, take a quick look at these four AI principles.

Together, choose the ONE principle that seems most relevant or interesting.

HUMAN CENTEREDNESS

AI should aim to do good, actively contributing to the welfare of individuals and society.

CONTROL & ACCOUNTABILITY

Humans should have significant control over AI systems and be accountable for their outcomes and potential harms.

OBSERVABILITY & EXPLAINABILITY

AI systems should be both observable and explainable, allowing humans to understand and monitor their decision-making processes.

IMPROVEMENT & RESPONSIVENESS

AI systems should be continuously improved to enhance their performance and address any emerging issues.

Then, discuss one clear 'DO' and one clear 'DON'T' for implementing that principle.

HUMAN CENTEREDNESS

AI should aim to do good, actively contributing to the welfare of individuals and society.



Ensure AI-powered chat has clear options to speak to a human.



Implement AI that makes lending decisions without human oversight.

CONTROL & ACCOUNTABILITY

Humans should have significant control over AI systems and be accountable for their outcomes and potential harms.



Create clear documentation of who is responsible for reviewing AI-generated content before it reaches members/customers.



Allow AI tools to access member/customer financial data without proper security controls.

OBSERVABILITY & EXPLAINABILITY

AI systems should be both observable and explainable, allowing humans to understand and monitor their decision-making processes.



Ensure loan officers can explain in simple terms how AI assists in their recommendations.



Implement "black box" systems for credit decisioning that can't be audited or explained.

IMPROVEMENT & RESPONSIVENESS

AI systems should be continuously improved to enhance their performance and address any emerging issues.



Establish regular review cycles to evaluate AI performance in member/customer service scenarios.



Deploy AI solutions without a feedback mechanism for frontline staff to report issues.

Discuss: 5 min

HUMAN CENTEREDNESS

AI should aim to do good, actively contributing to the welfare of individuals and society.

CONTROL & ACCOUNTABILITY

Humans should have significant control over AI systems and be accountable for their outcomes and potential harms.

OBSERVABILITY & EXPLAINABILITY

AI systems should be both observable and explainable, allowing humans to understand and monitor their decision-making processes.

IMPROVEMENT & RESPONSIVENESS

AI systems should be continuously improved to enhance their performance and address any emerging issues.

Let's hear from different groups!

What 'DO' and 'DON'T' did you come up with?

HUMAN CENTEREDNESS

AI should aim to do good, actively contributing to the welfare of individuals and society.

CONTROL & ACCOUNTABILITY

Humans should have significant control over AI systems and be accountable for their outcomes and potential harms.

OBSERVABILITY & EXPLAINABILITY

AI systems should be both observable and explainable, allowing humans to understand and monitor their decision-making processes.

IMPROVEMENT & RESPONSIVENESS

AI systems should be continuously improved to enhance their performance and address any emerging issues.

3. MAKING STRATEGIC CHOICES

Stand up again!
Answer these
questions with
your bodies.



STAGE

Strategic tension #1

CONTROL

How will we
manage the
introduction and
use of AI in our
organization?

Centralized

A dedicated AI
governance committee
approves all AI tools
and monitors usage
across all departments



Distributed

Departments and branch
managers can select and
implement pre-approved
AI tools based on local
community needs.

STAGE

Strategic tension #2

USERS

Who within our
organization will
use AI?

Everyone ← → **Specialists**

All employees can
use AI tools in their
daily work.

Only designated roles
are trained to use AI
tools and share insights
with their teams.

STAGE

Strategic tension #3

APPROACH

Who within our
organization will
use AI?

Pioneers

Be first-to-market with
AI-powered services in
your community to
differentiate from
competitors.



Followers

Observe how other
institutions implement
AI and adopt proven
approaches with lower
risk.

STAGE

Every strategic choice has trade-offs

The first step is creating space for open dialogue across positions and departments.

YOUR 3-STEP AI STRATEGY JUMPSTART:

1. Surface Diverse Perspectives:

Create psychological safety to discuss AI hopes and concerns.

2. Establish Clear Principles:

Translate values into practical AI do's and don'ts.

3. Make Strategic Choices:

Explicitly decide key trade-offs for your organization.

This is a new world.

We're only at the beginning.

And it's ours to shape.



THANK YOU.

Let's keep the
conversation going.

✉ brent@dxn.is