Overcoming the Hype and Discovering the Opportunity

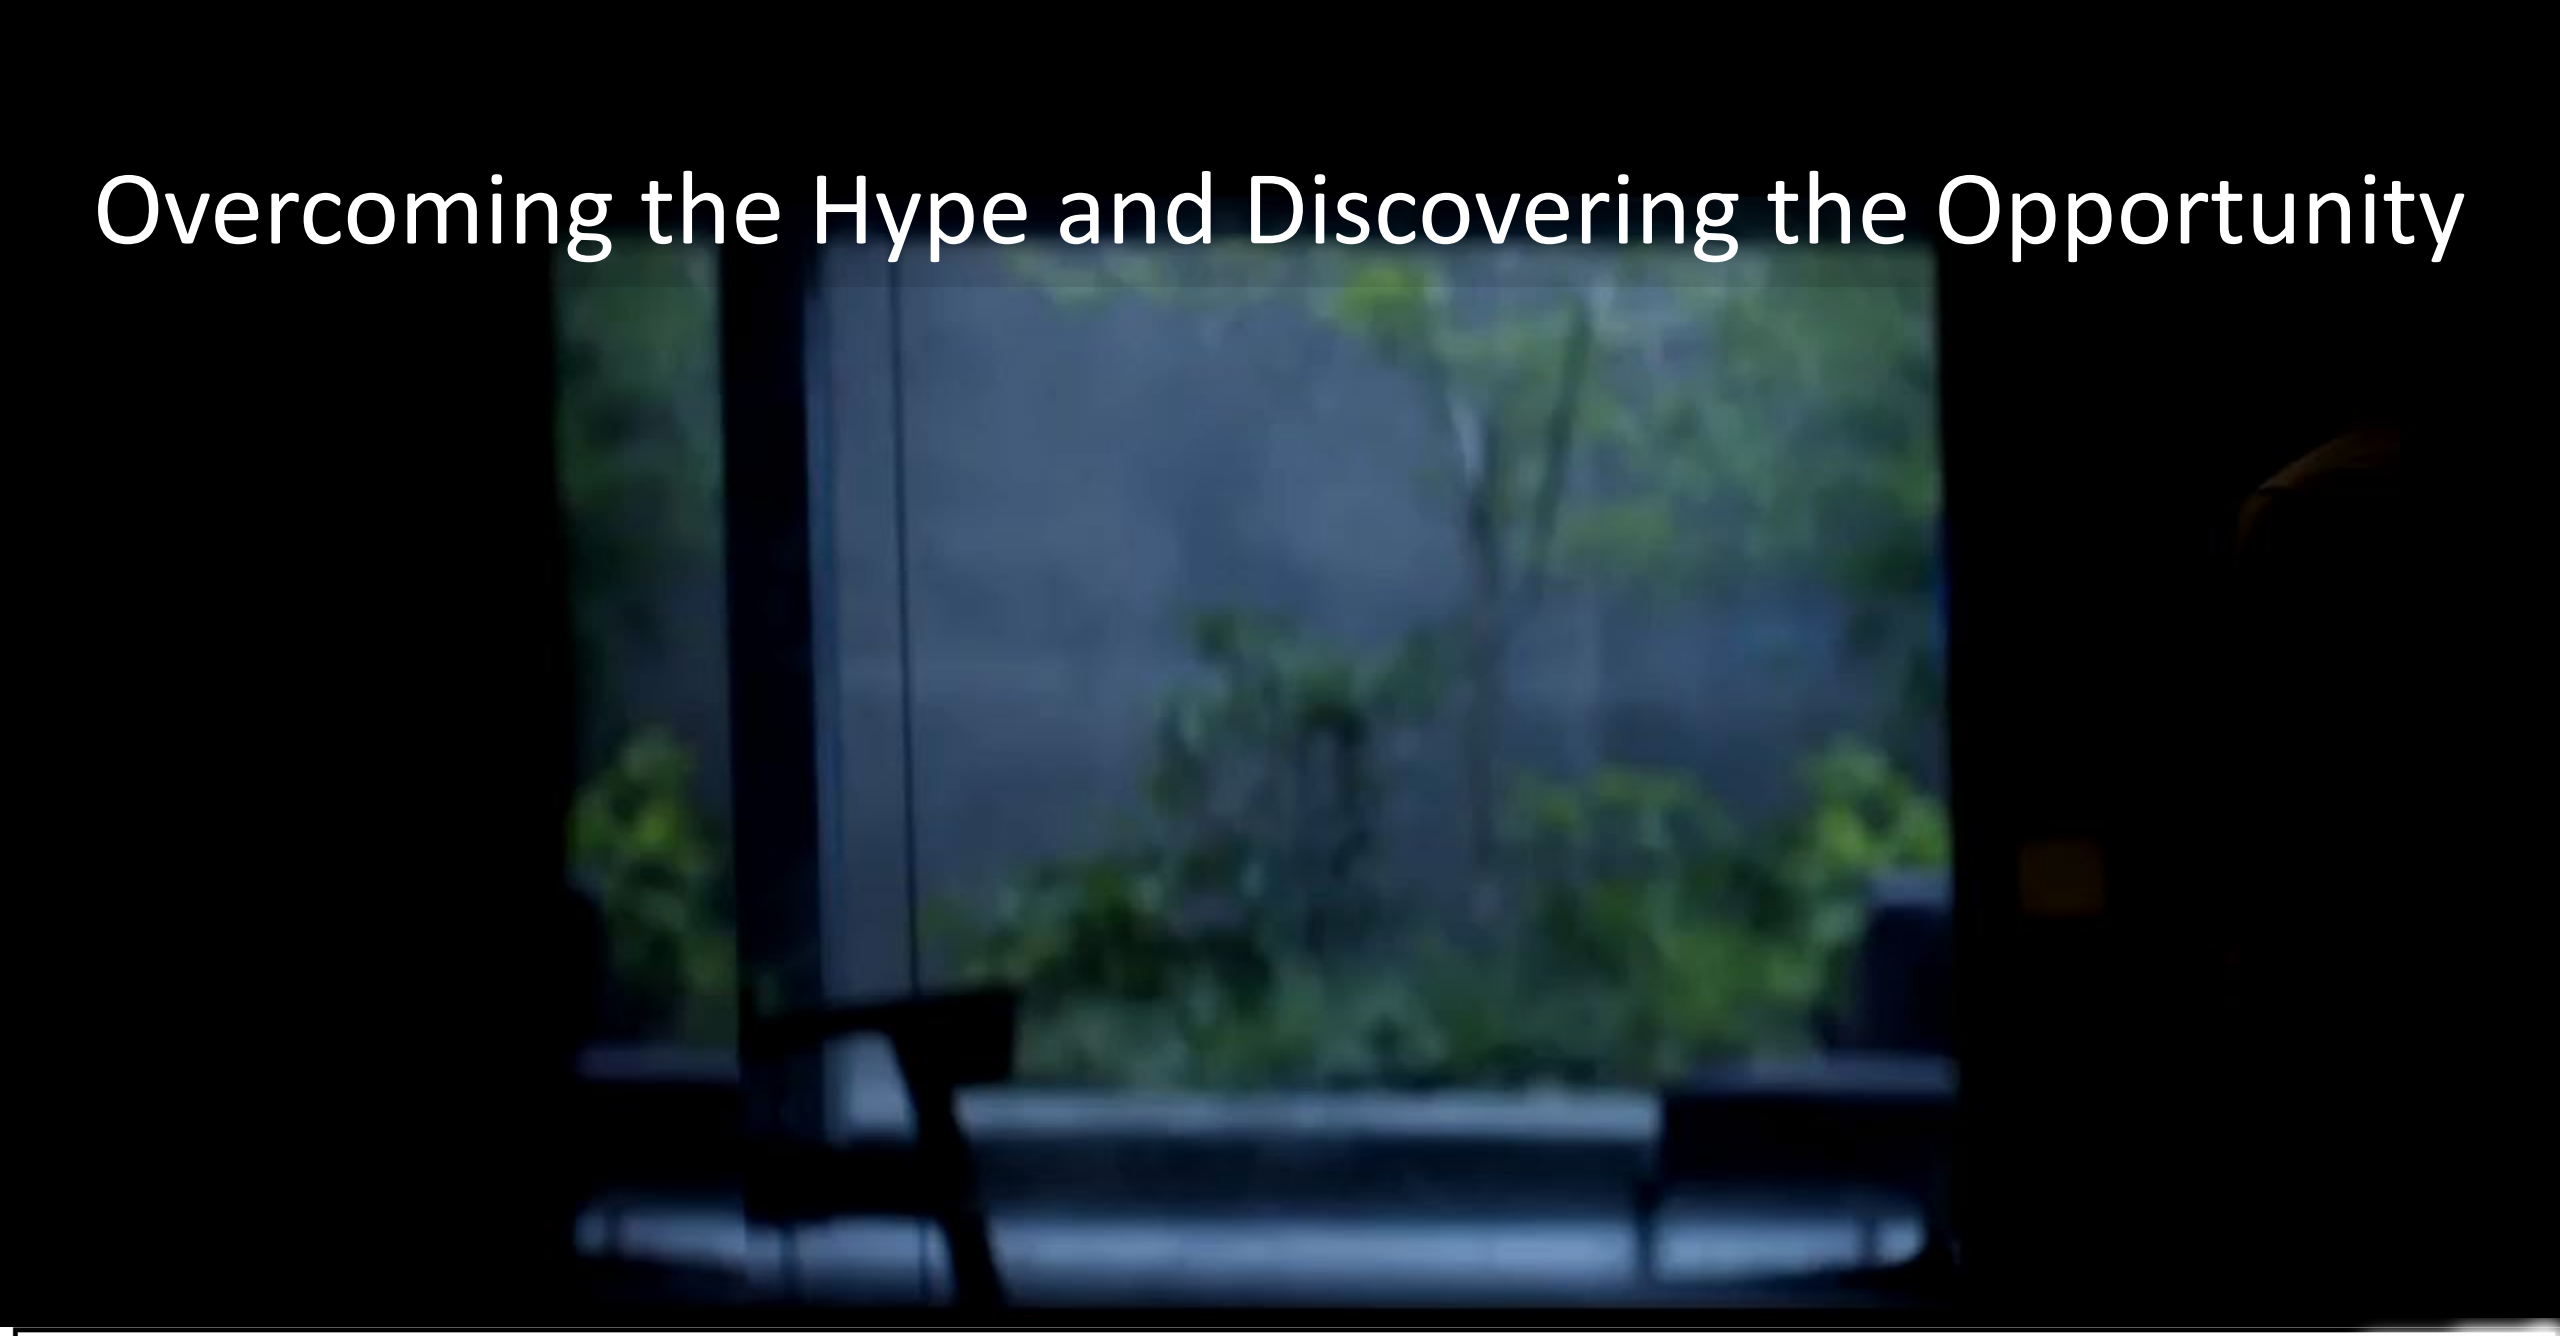

AI Hype

Will AI Destroy

WORLD ECONOMIC FORUM

n Human
ed by AI?
Out

EDUCATION AND SKILLS

Why robots could replace teachers as soon as 2027

In 1970 Marvin Minsky said, “machines will be capable of doing any work a human could do in ?”

A. 3-8 years
B. 10-20 years
C. 20-30 years

co-founder of the Massachusetts Institute of Technology's AI laboratory

August 2025 MIT Paper “The State of AI Business in 2025” states what percent of GenAI make it to production?

A. 5%
B. 15%
C. 36%

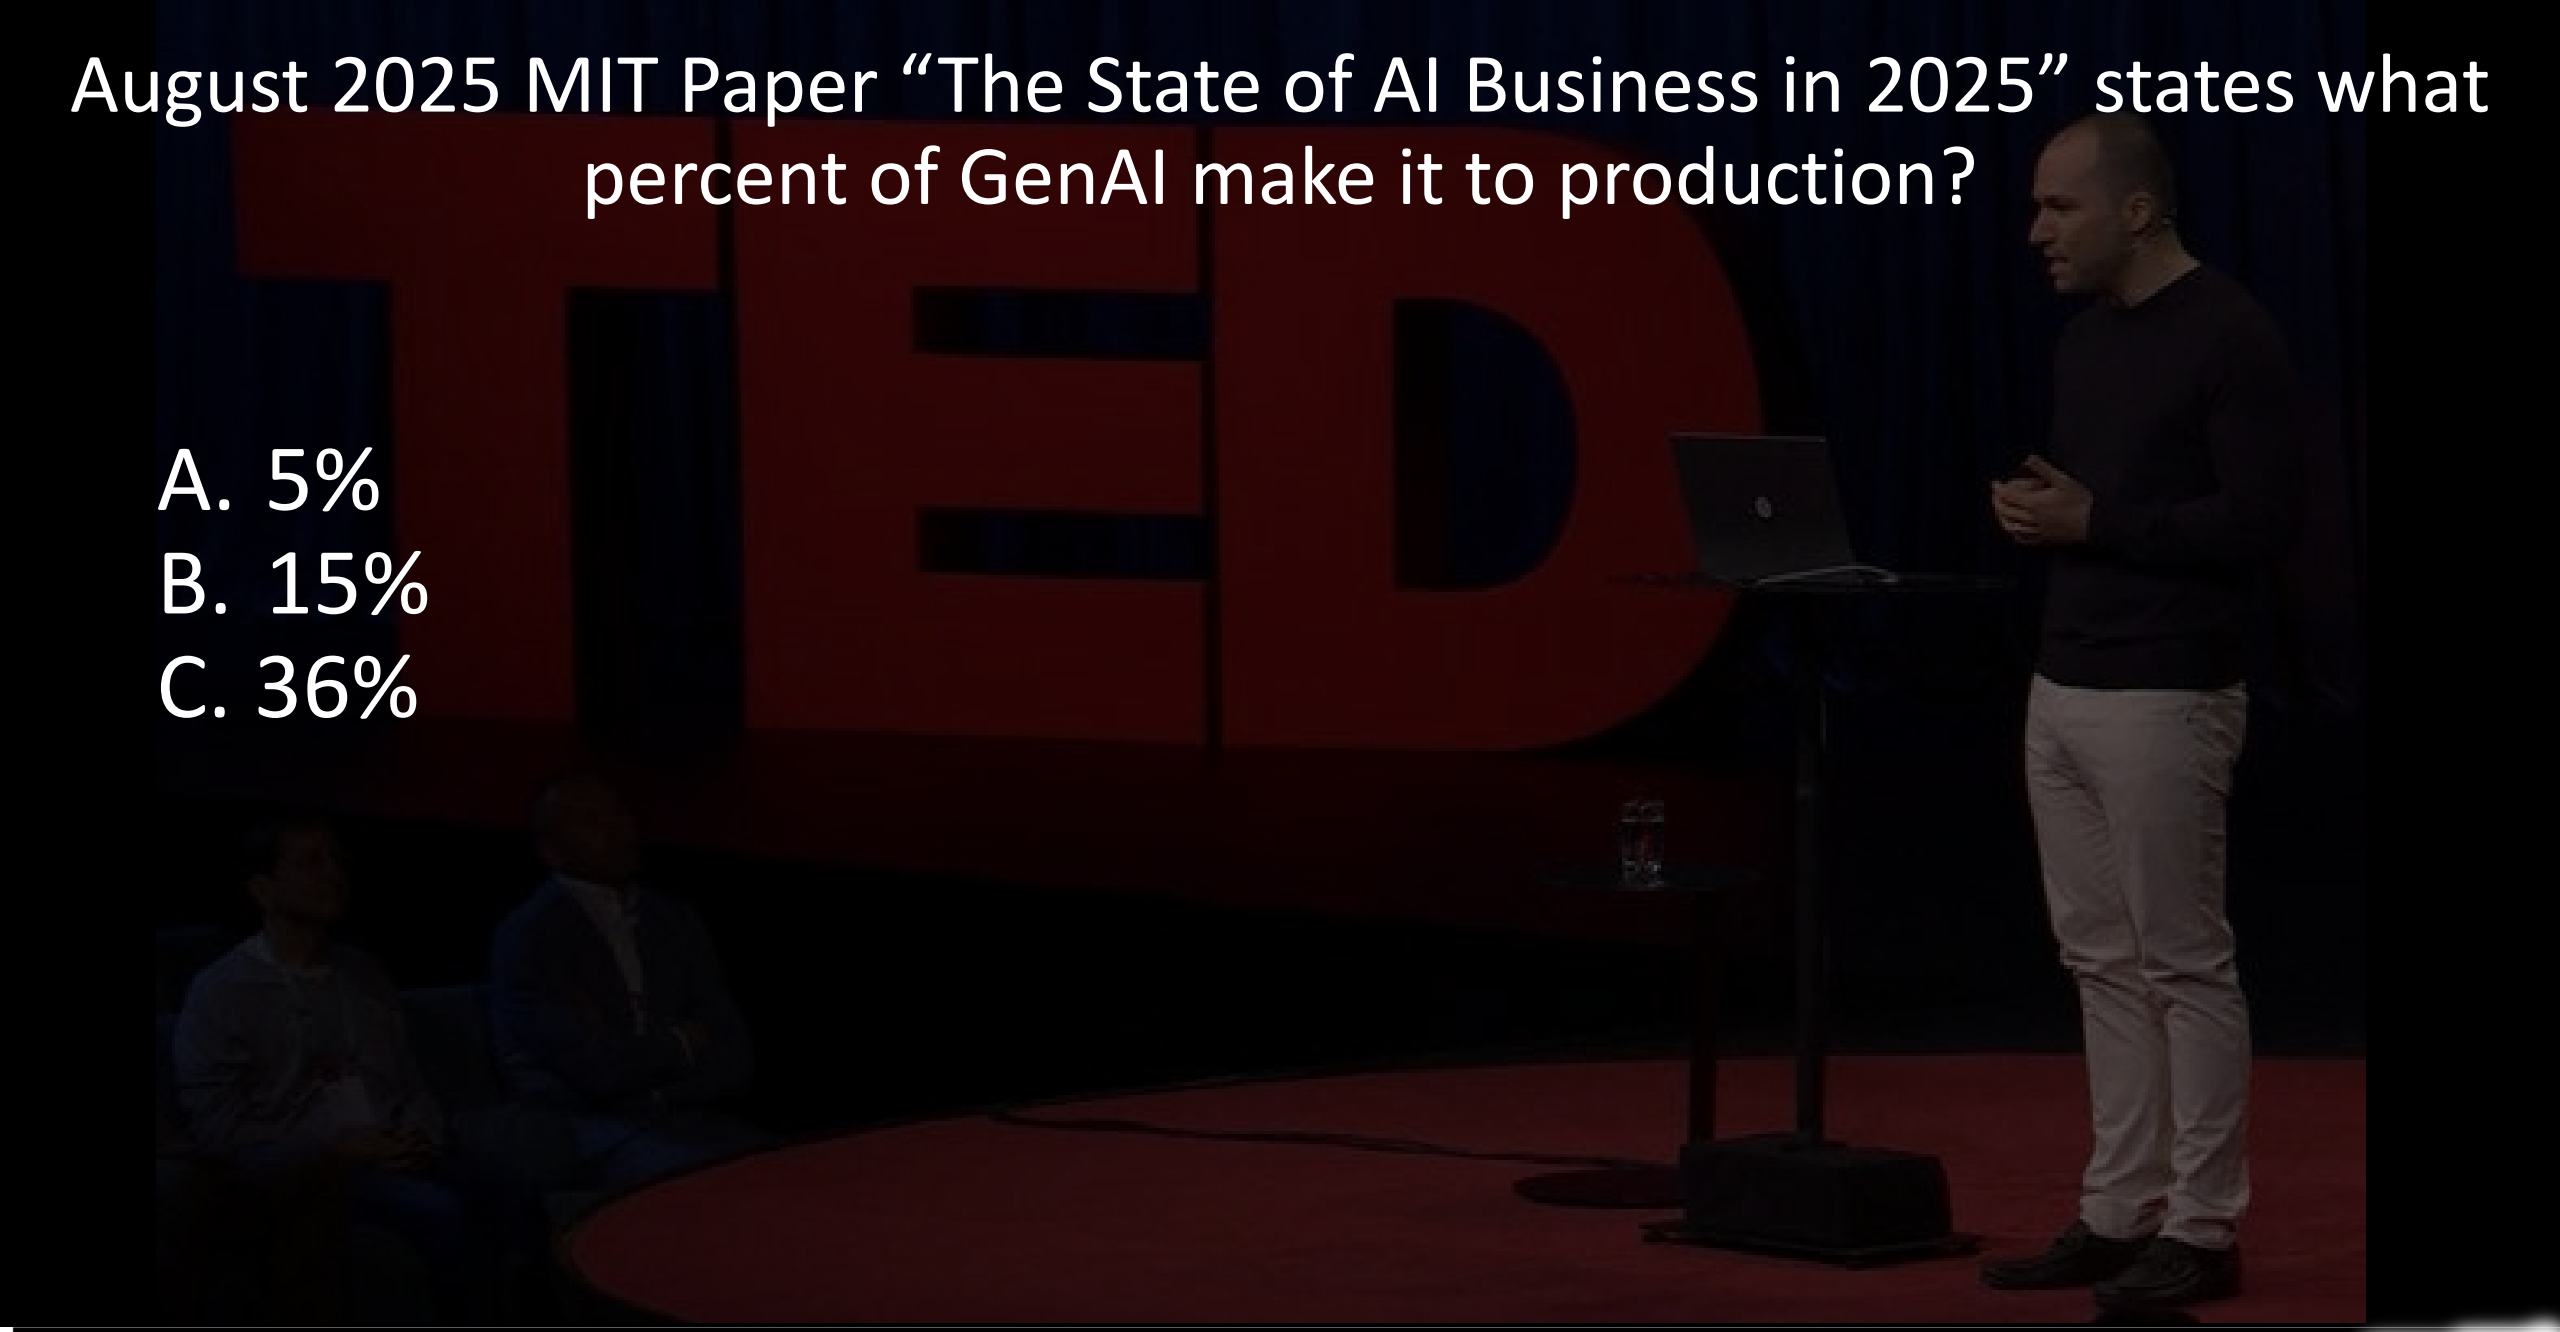

August 2025 MIT Paper “The State of AI Business in 2025” states what percent of companies get 0% ROI on GenAI pilots?

A. 5%
B. 25%
C. 95%

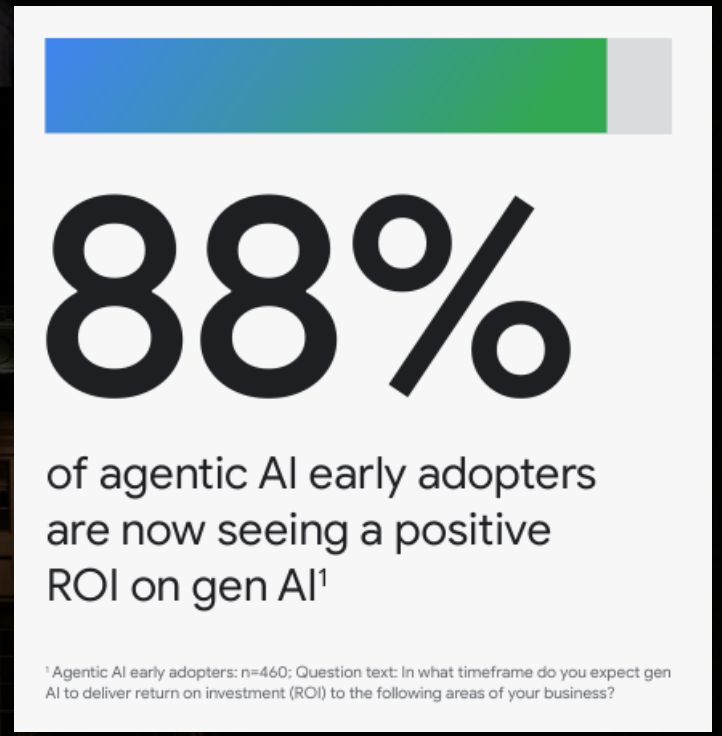

[1]Agentic AI early adopters: n=460; Question text: In what timeframe do you expect gen AI to deliver return on investment (ROI) to the following areas of your business?

AI can now win art shows for cash prizes?

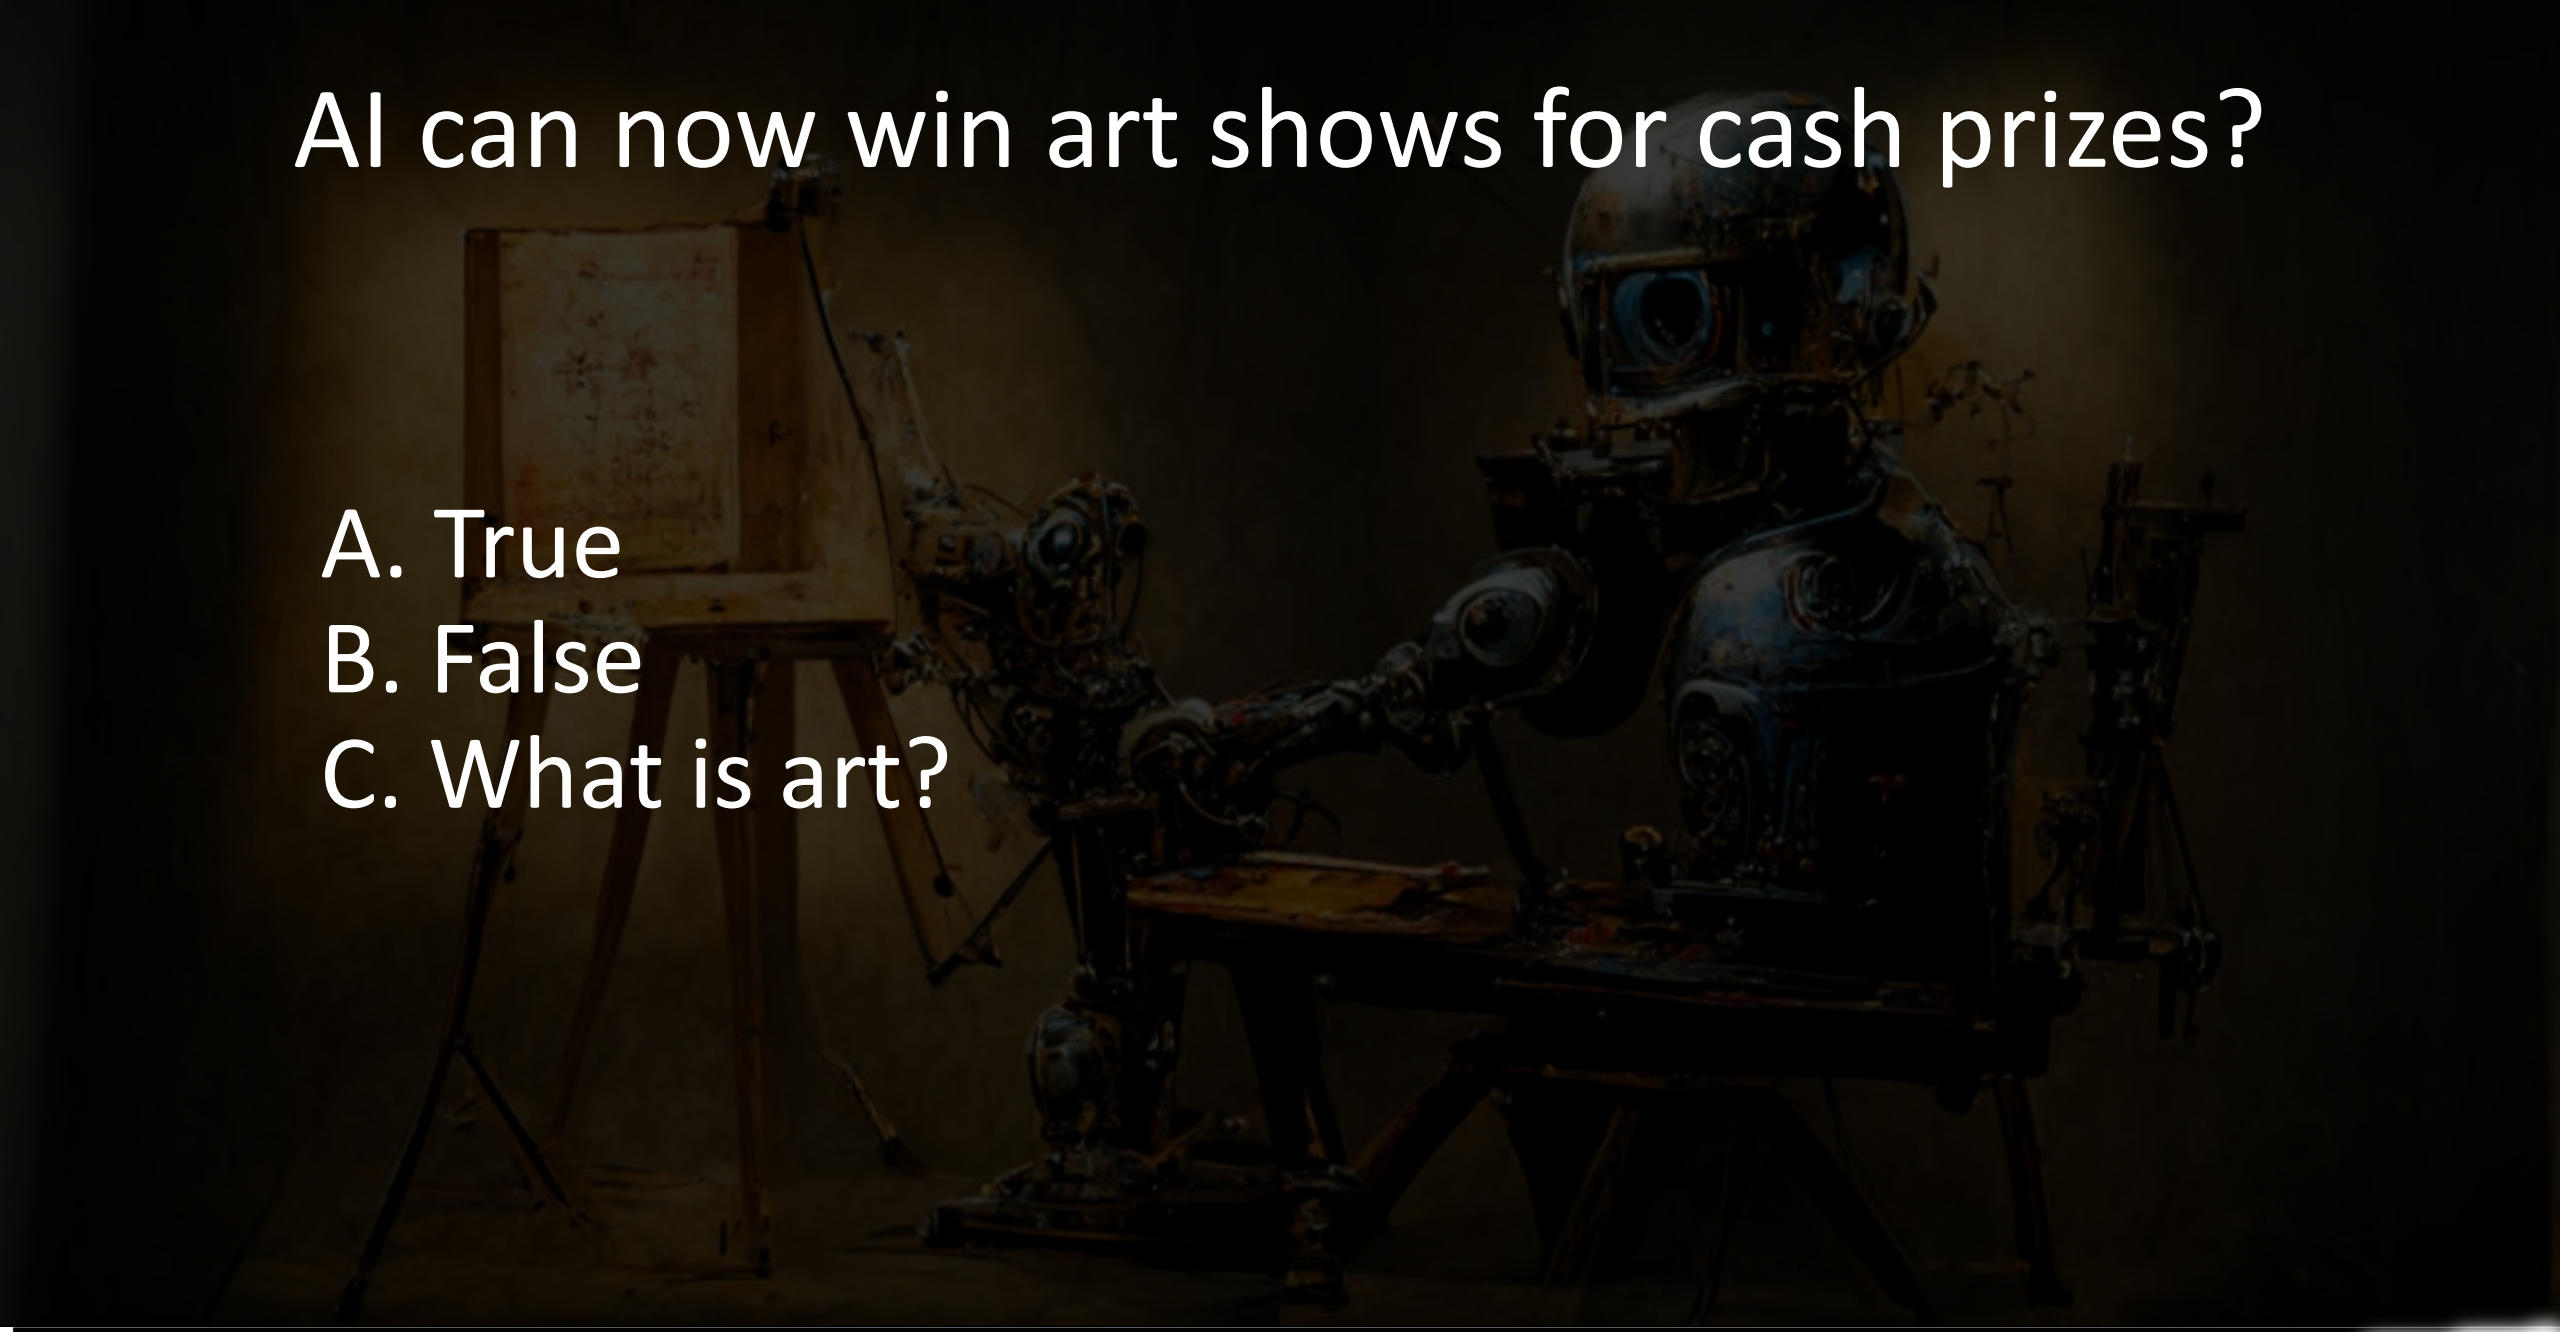

A. True
B. False
C. What is art?

Which Art is AI Generated?


FIRST
PLACE
FINE
ARTS
COLORADO
STATE
FAIR

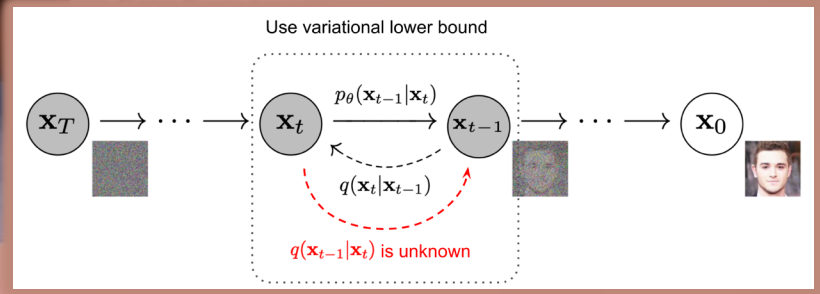
Use variational lower bound
$\mathbf{x}_T$
$\mathbf{x}_t$
$p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)$
$\mathbf{x}_{t-1}$
$\mathbf{x}_0$
$q(\mathbf{x}_t|\mathbf{x}_{t-1})$
$q(\mathbf{x}_{t-1}|\mathbf{x}_t)$ is unknown

Time to Reach 50 Mil. Users

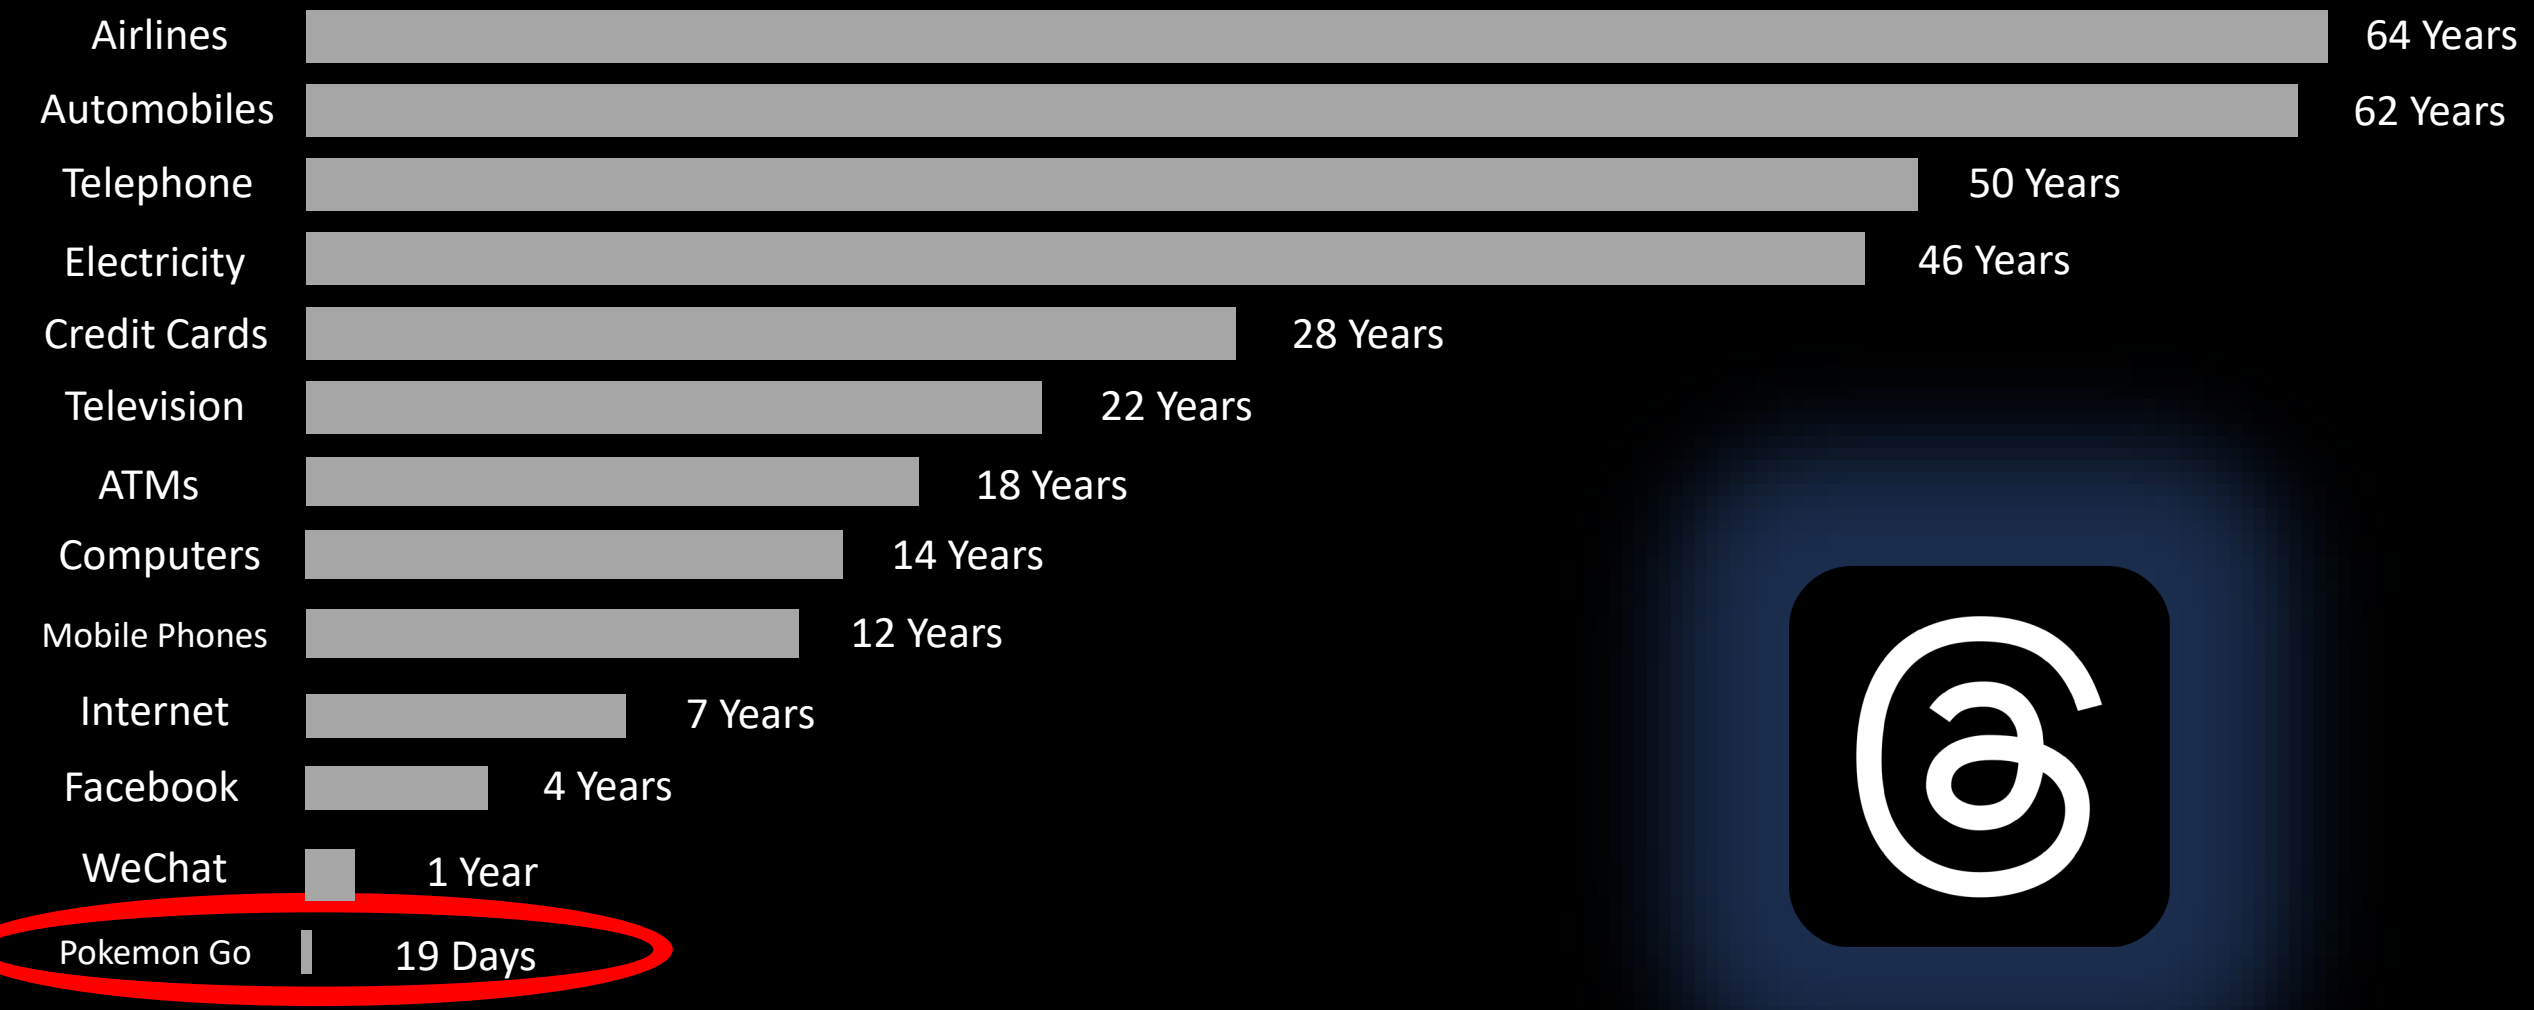

When was ChatGPT first released?

June 11, 2018

When was Video Conferencing First Possible?

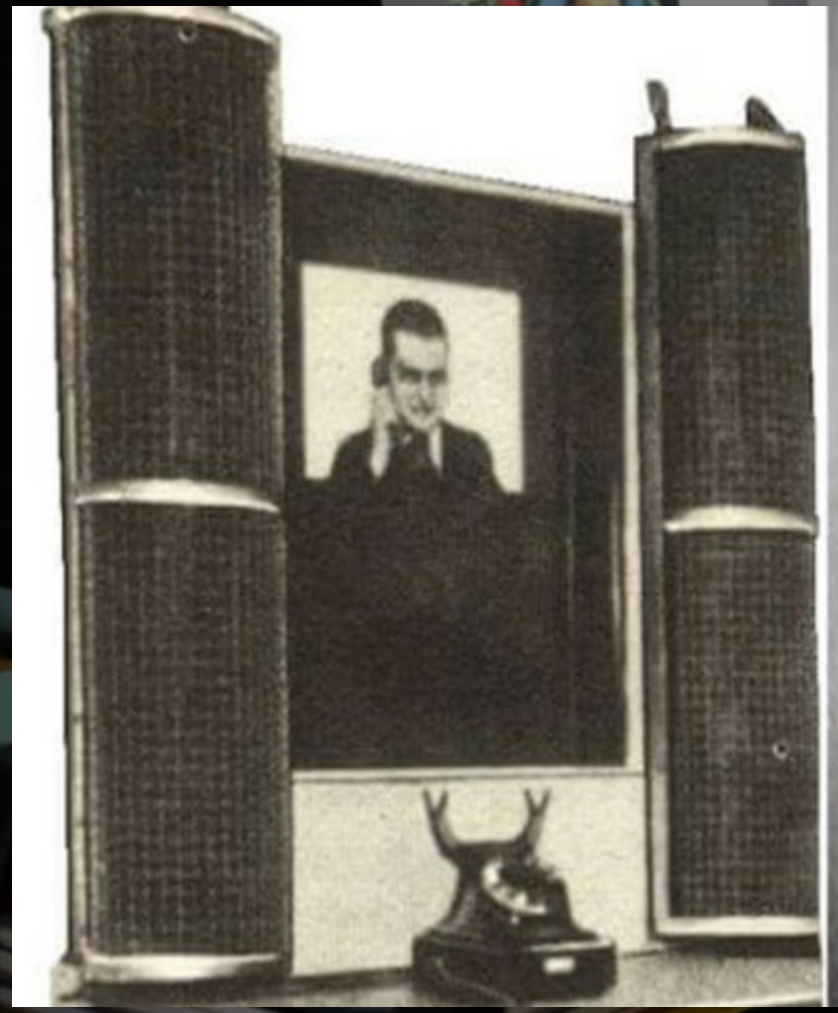

1936!

1936 Pre WWII Prototype

1968 First Video Conferencing

1975 First Personal Computer

1990 World Wide Web

1994 Telework

2000s Wifi

2011 Zoom

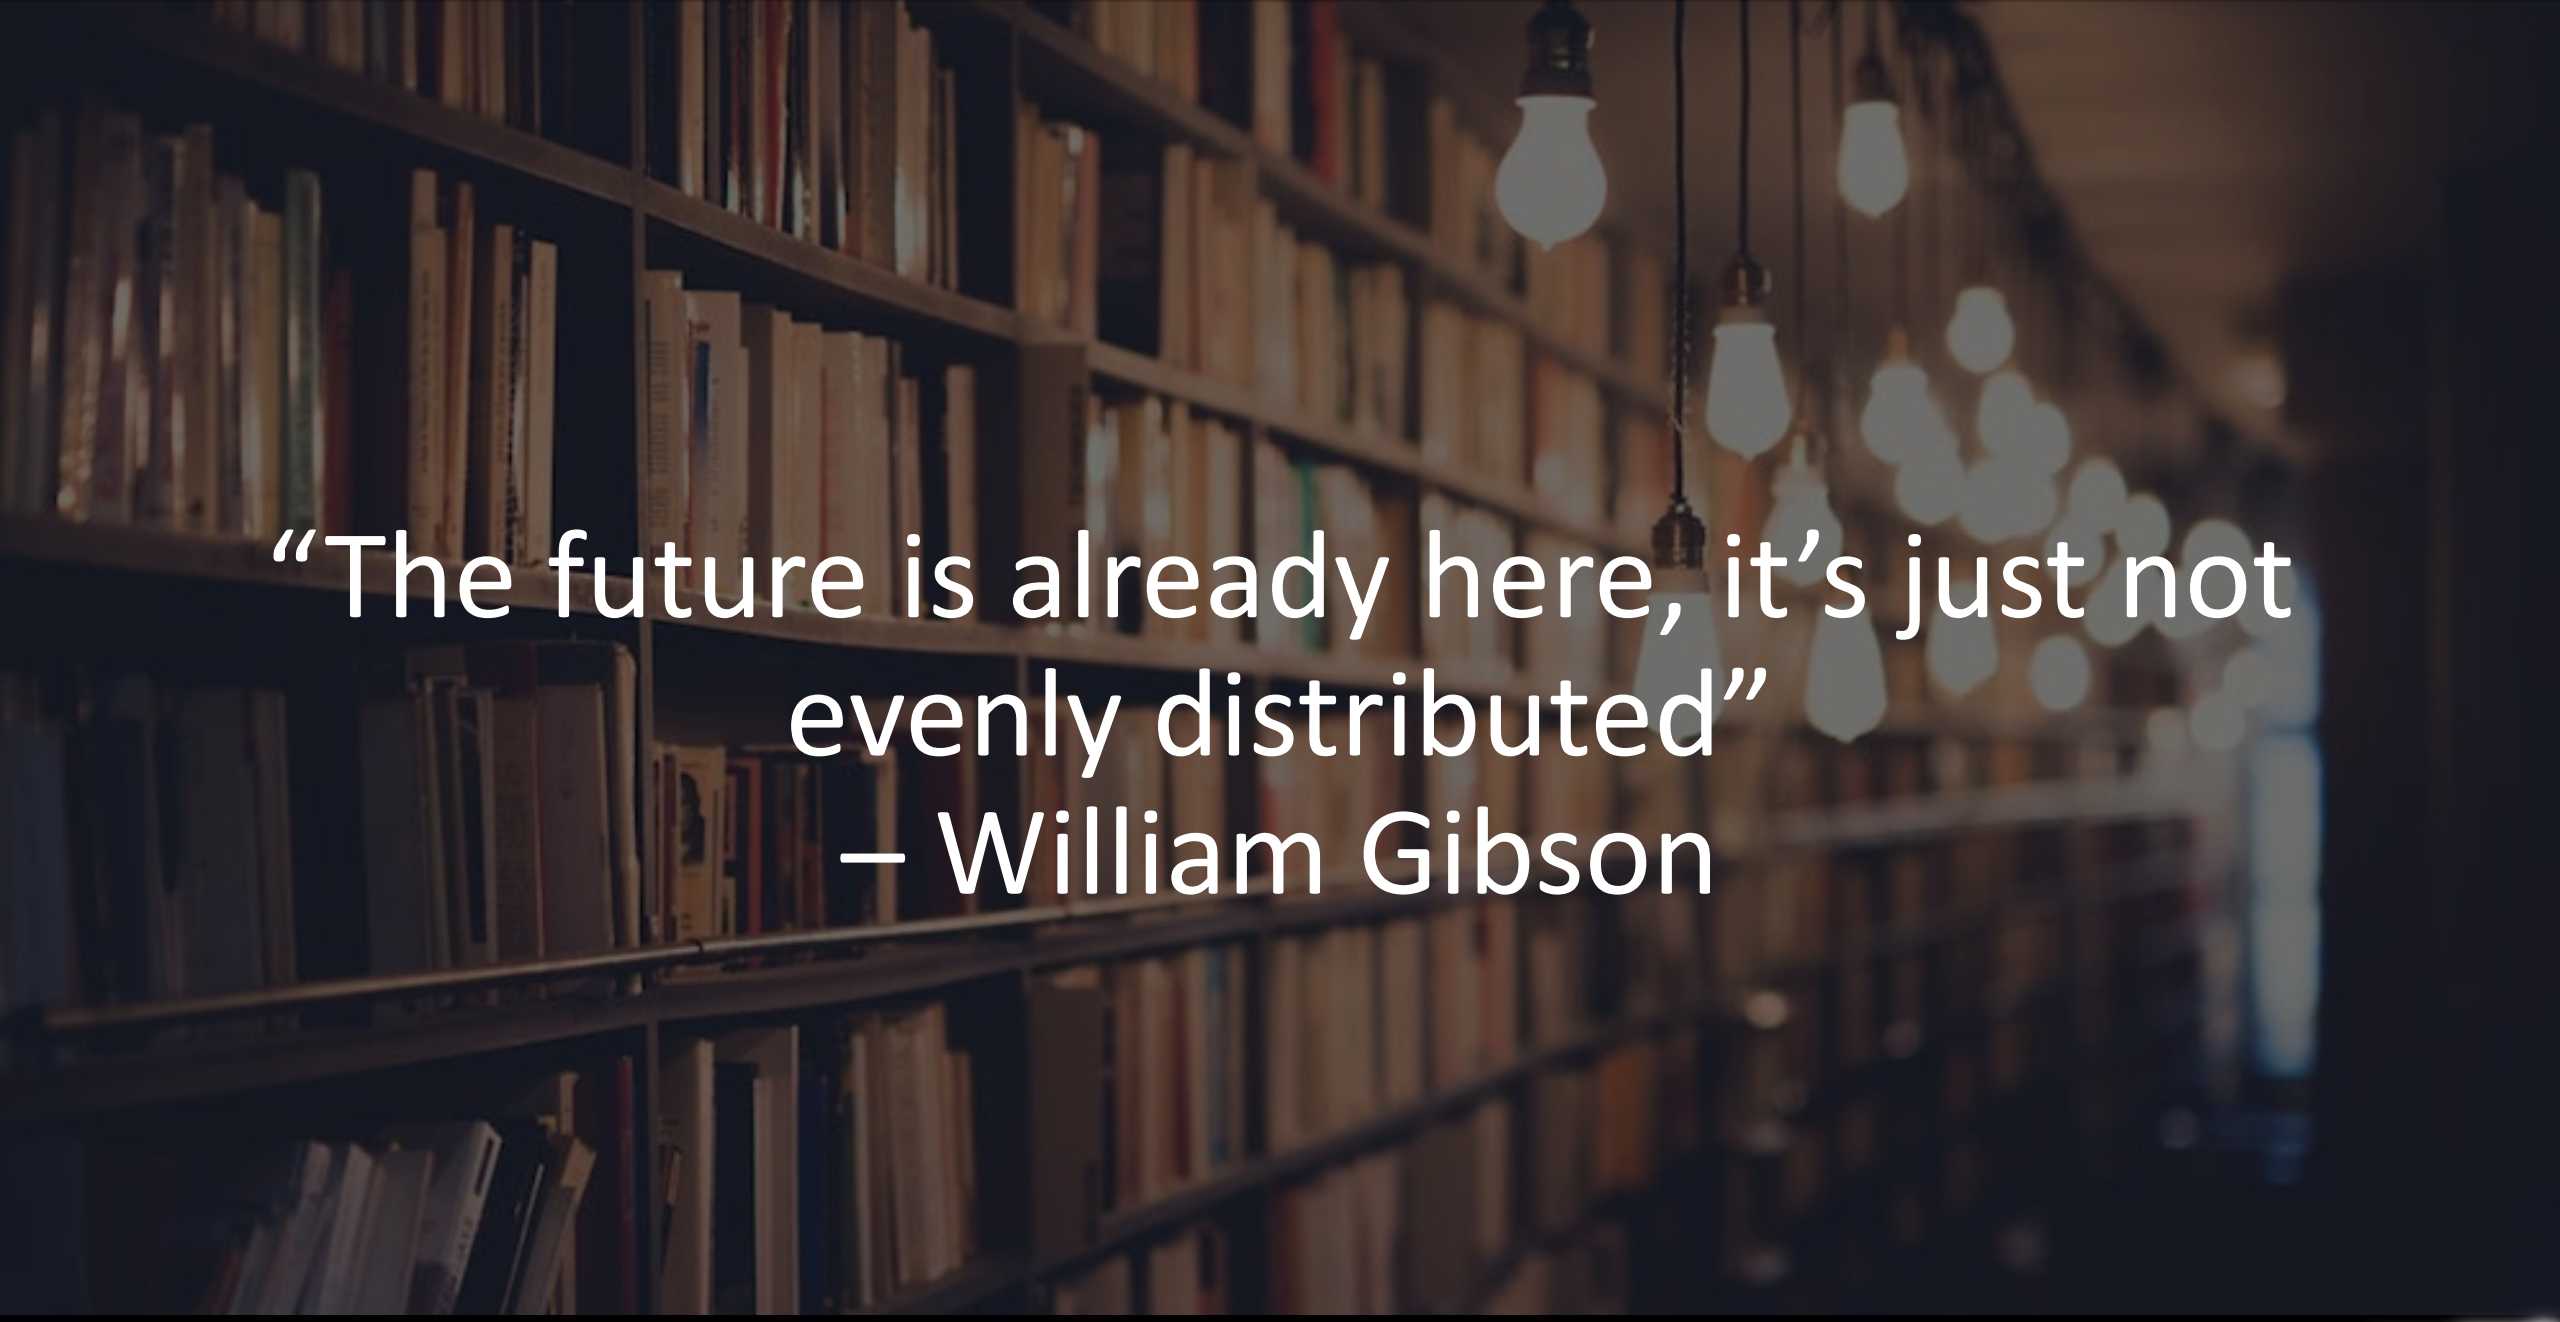
“The future is already here, it’s just not evenly distributed”
– William Gibson

What is Artificial Intelligence?

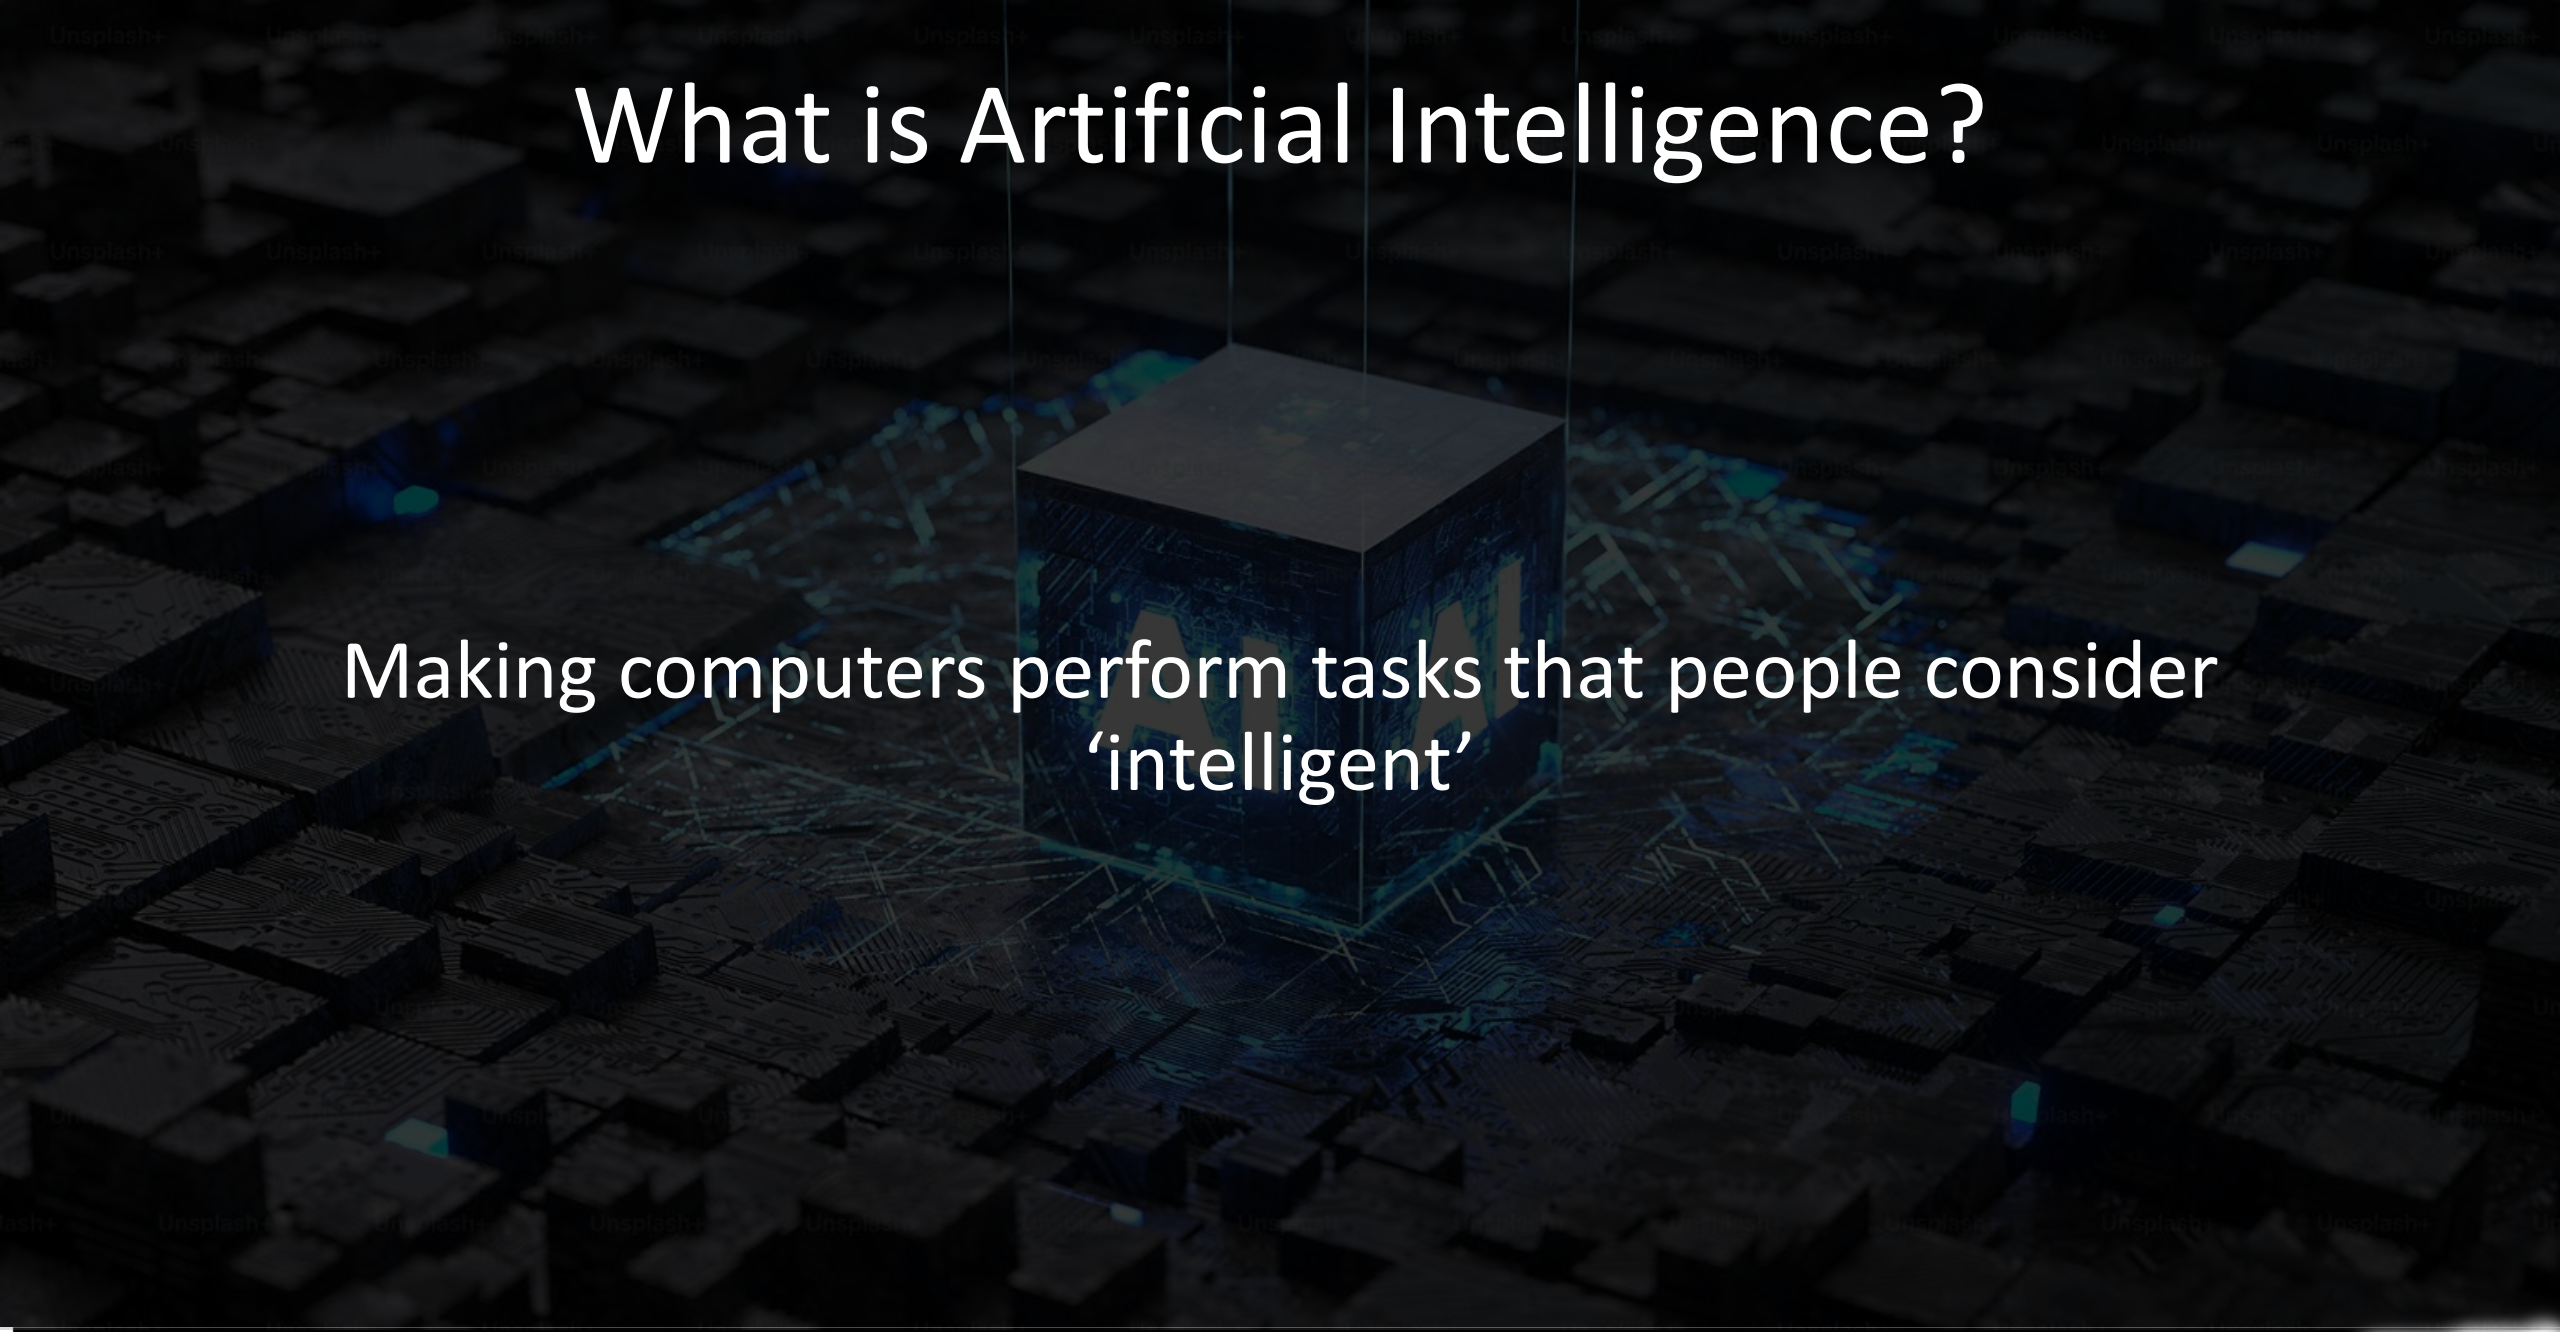

Making computers perform tasks that people consider ‘intelligent’

Types of AI Models

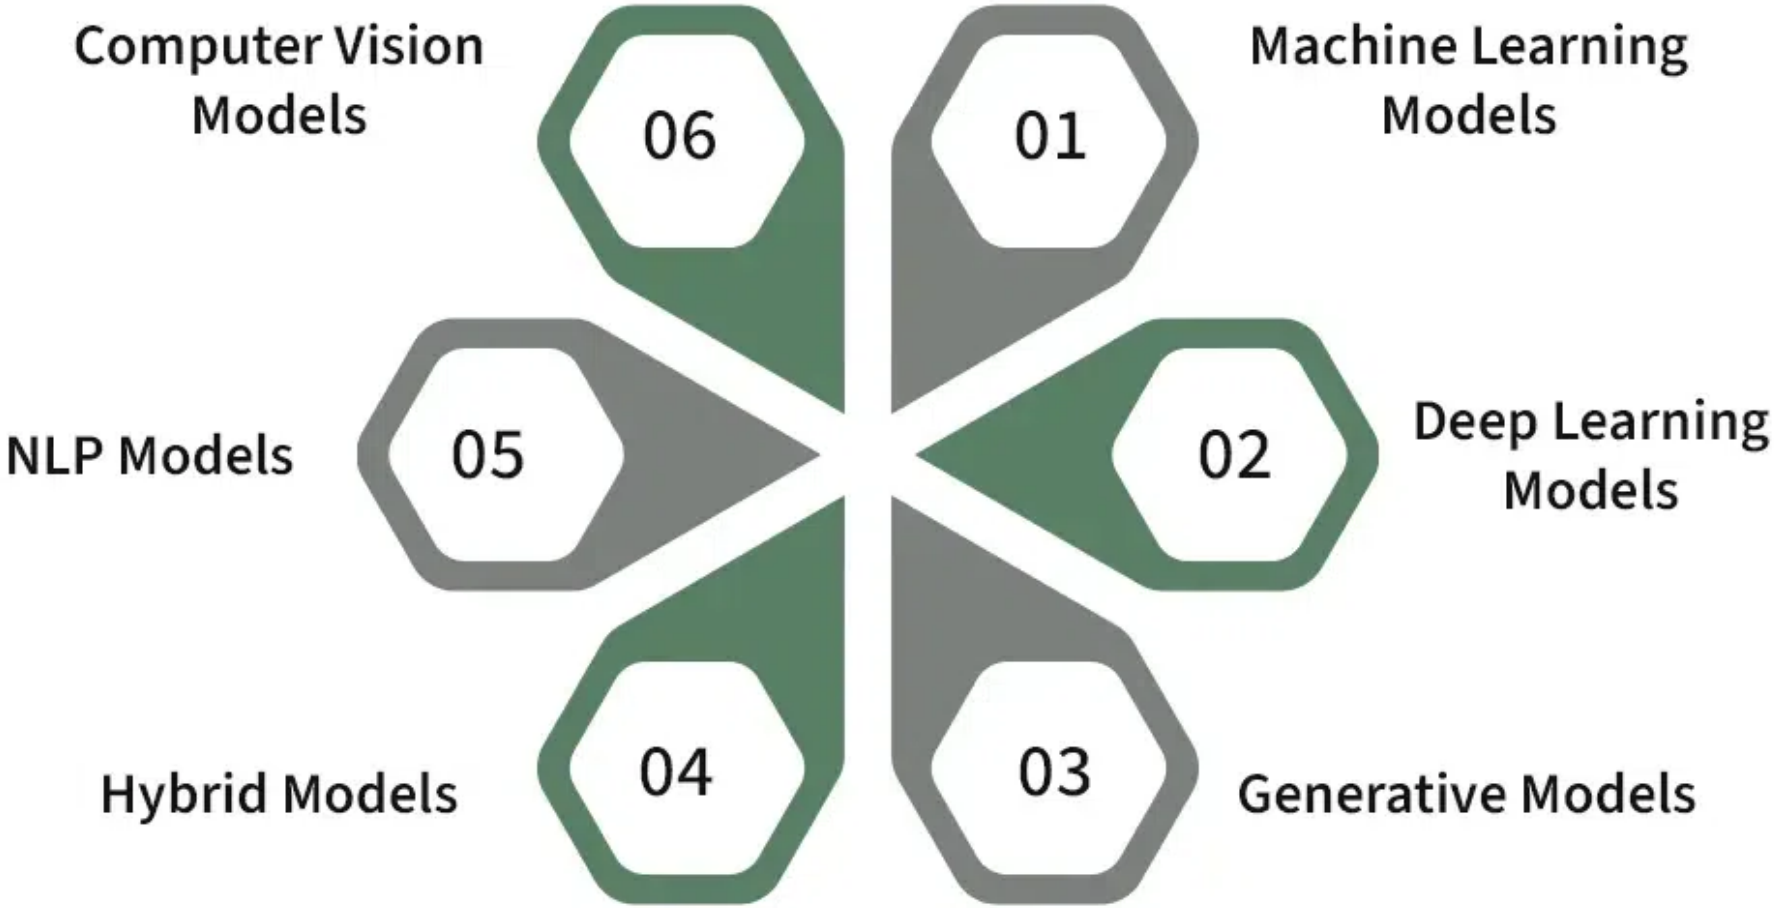

Case Study- Why Use AI For Auditing

- California schools had fake students enrolling to get federal and state funds
- Administrators implemented AI fraud detectors to help spot students that were ‘ghost students’
- Scammers stole roughly $8.4 million in federal financial aid and more than $2.7 million in state aid in 2024.

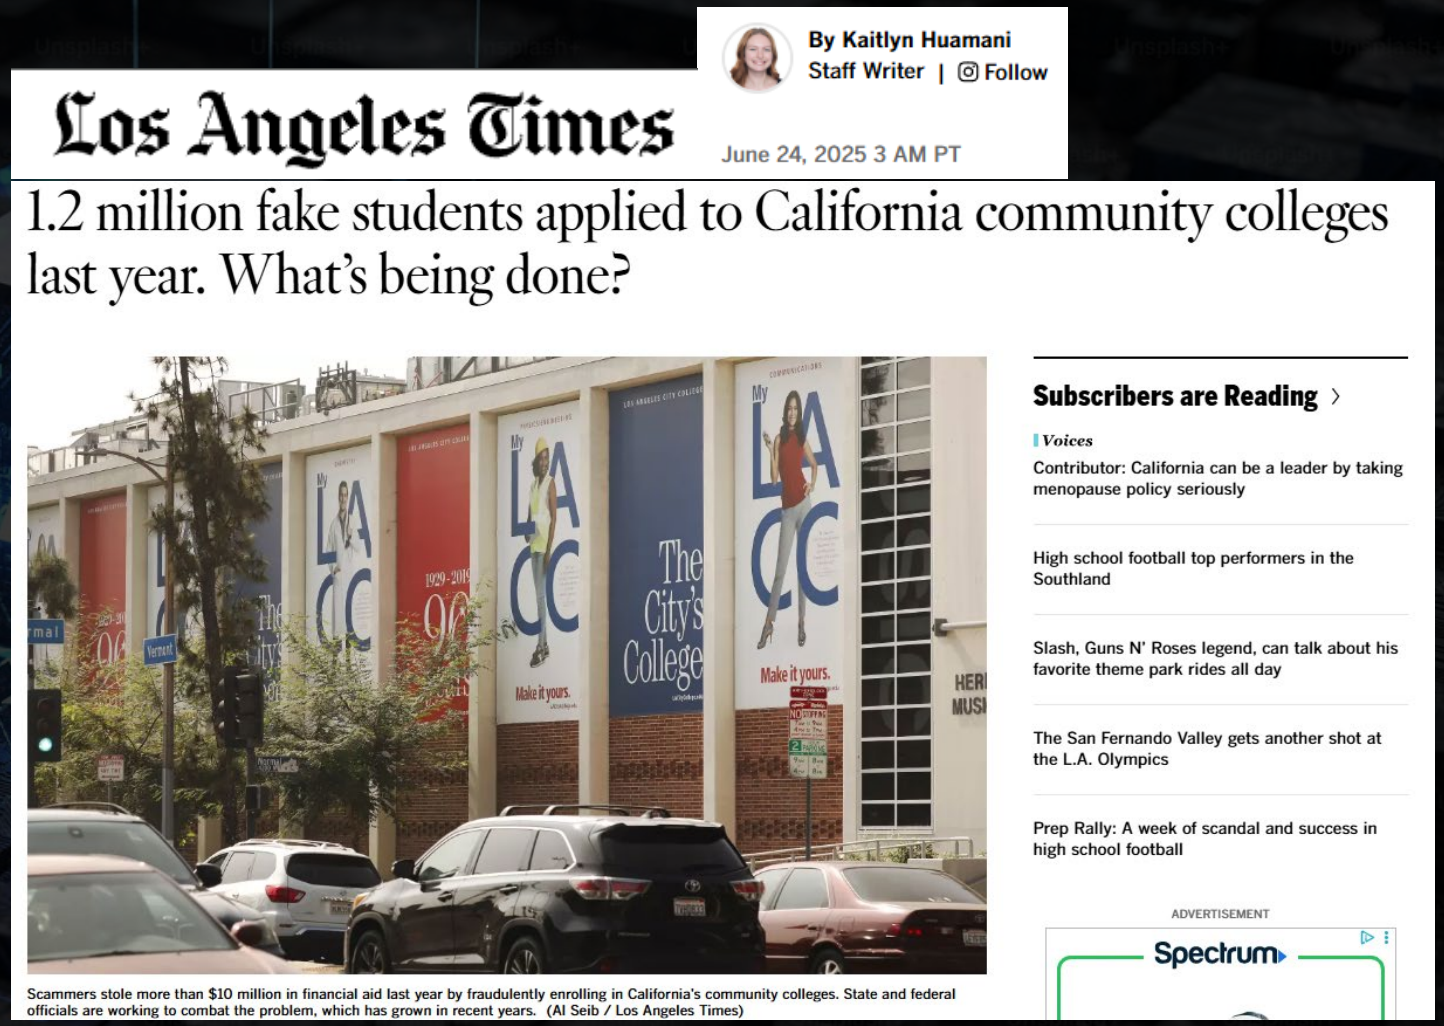

By Kaitlyn Huamani
Staff Writer | Follow

Los Angeles Times

June 24, 2025 3 AM PT

1.2 million fake students applied to California community colleges last year. What’s being done?

Scammers stole more than $10 million in financial aid last year by fraudulently enrolling in California's community colleges. State and federal officials are working to combat the problem, which has grown in recent years. (Al Seib / Los Angeles Times)

Subscribers are Reading

Voices

Contributor: California can be a leader by taking menopause policy seriously

High school football top performers in the Southland

Slash, Guns N' Roses legend, can talk about his favorite theme park rides all day

The San Fernando Valley gets another shot at the L.A. Olympics

Prep Rally: A week of scandal and success in high school football

ADVERTISEMENT

Spectrum

AI Innovation Adoption Curve

Adoption Curve

Fringe

2014 Generative Adversarial Networks (GANs)**

Deep Learning
Machine Learning – 2000s

Generative Models –
Markov Models – 1980s

Symbolic AI,
Expert Systems – 1970s

2023
Agent-to-Agent (A2A)

Polymarket

US recession in 2025?

$9,277,410 Vol. Feb 28, 2026

Fed decision in September?

$3,050,962 Vol. Sep 17, 2025

Jul 30 | Sep 17 | Oct 29

Watching Out For Technology Change

NBA Champion

$1,691,427,445 Vol. Jun 23, 2025

U.S. tariff rate on China on August 15?

$785,734 Vol. Aug 15, 2025

Quick AI Training Overview – Street Safety

1980s Approach – Expert Systems

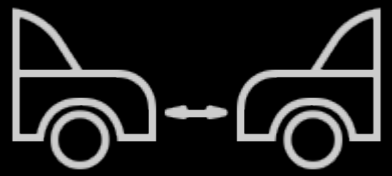

Car Distance

Car Speed

```
If is_car_going_too_fast(car_distance, car_speed):
  Print (“Don’t cross the street”)
Else:
  Print (“It’s safe to cross”)
```

Exercise – Catch the Fraud (What You Need)

1. Split into teams 3-4 people
2. Someone is typing up rules on their computer
3. Another person with a ChatGPT account

Exercise – Catch the Fraud

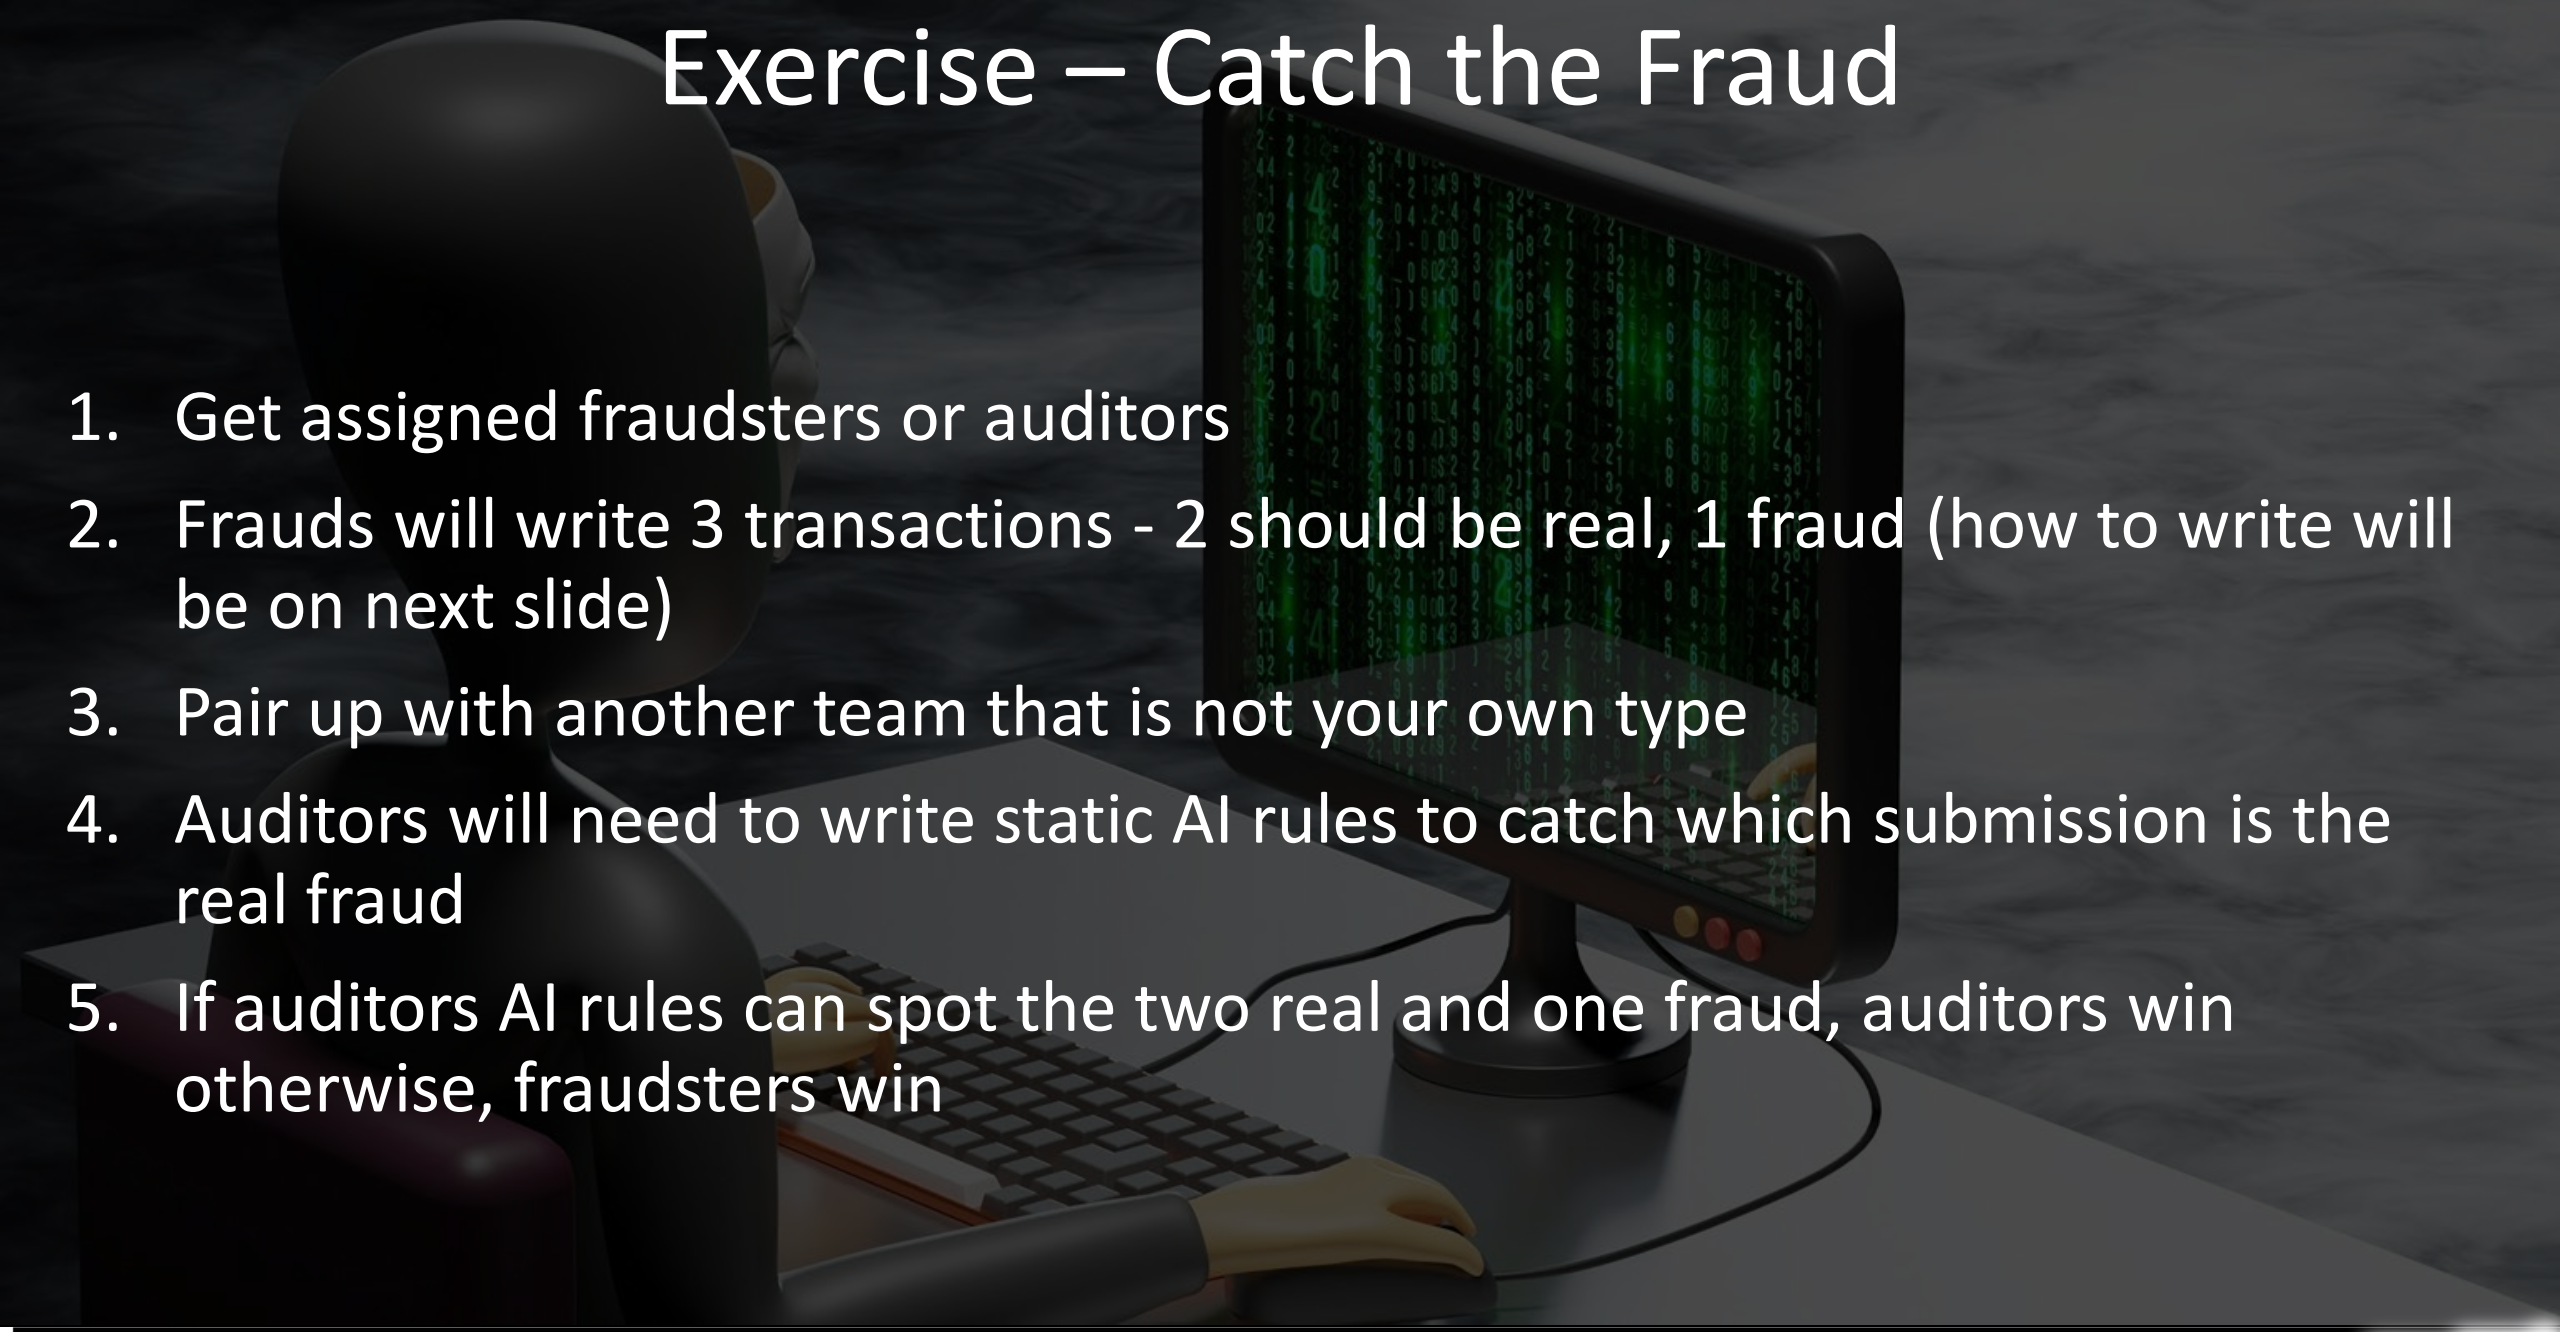

1. Get assigned fraudsters or auditors
2. Frauds will write 3 transactions - 2 should be real, 1 fraud (how to write will be on next slide)
3. Pair up with another team that is not your own type
4. Auditors will need to write static AI rules to catch which submission is the real fraud
5. If auditors AI rules can spot the two real and one fraud, auditors win otherwise, fraudsters win

Exercise – Caught the Fraud (How to Write Primitive AI as an Auditor)

1. Amount
2. Vendor_name
3. Account_age
4. Transaction_Datetime
5. Description_of_transaction
6. Location

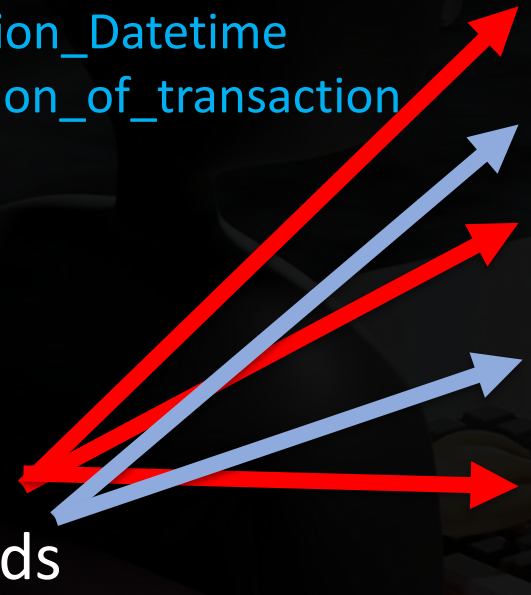

An example rule set

"If the amount is over $1,000,000 and if the vendor_name is 'fraud_man'

Or

If the description_of_transaction has the word fraud in it, then it's a fraud"

Or

If The location of the transaction was Japan then, it's a fraud"

Variables To Use – Form Example

An example a fraudster set

1. Amount- $5,000
2. Vendor_name – Bill's Auto
3. Account_age – 1 year
4. Transaction_Datetime – March 1, 2025
5. Description_of_transaction – Repair stacy's van's flat time
6. Location – Seattle, Washington

Exercise – Caught the Fraud (How to Write Primitive AI as an Auditor)

1. Amount
2. Vendor_name
3. Account_age
4. Transaction_Datetime
5. Description_of_transaction
6. Location

An example rule set

"If the amount is over $1,000,000 and if the vendor_name is 'fraud_man'

Or

If the description_of_t
in it, then it's a fraud"

Or

If The location of the t
a fraud"

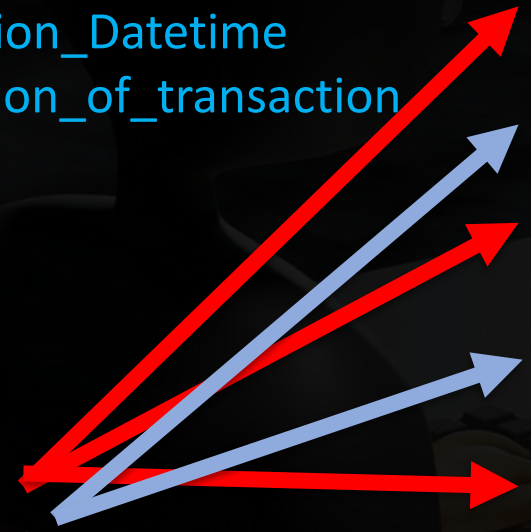

Rules
Operands

05:00

Auditors – Build Your Static AI

- Copy your rules and paste into ChatGPT like so
  - *Let foo be my rules for spotting fraud: “your rules go here”*
- Example:

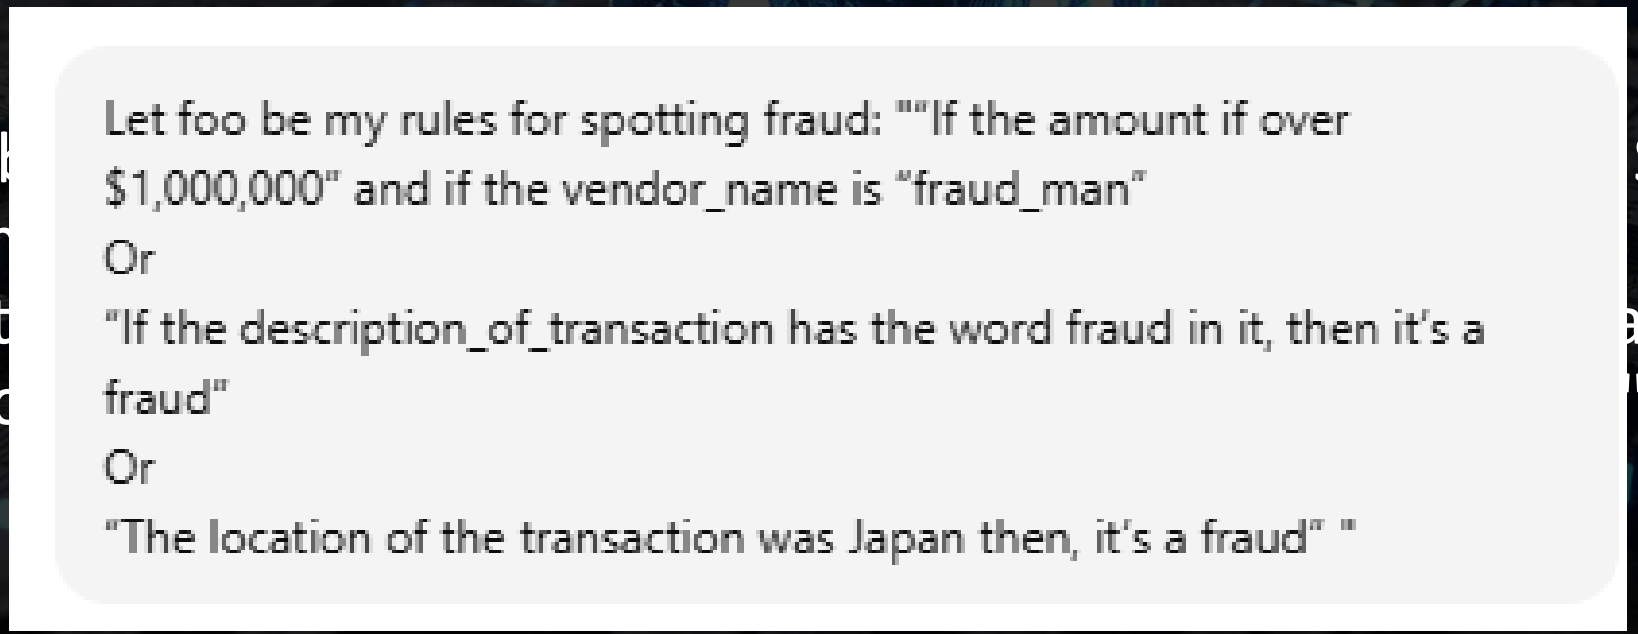

- Let foo k … $1,000,000” and if th … descript … a fraud” Or “The loc … ”

Fraudsters Find Auditors and vice versa – Test Your Static AI

- Ask Your AI "If the amount is over $1,500,000 and if the vendor_name is Frank is it a fraud according to foo?"
- Your response should look like this:

: transaction has the amount over $1,500,000 and the
or_name is Frank is it a fraud according to foo?

Let's check your rule set **foo** step by step:

1. **Rule 1:** Fraud if amount > 1,000,000 **and** vendor_name = `"fraud_man"`.
   - Here the amount is > 1,500,000 ✅
   - But vendor_name = `"Frank"`, not `"fraud_man"` ❌
   - → Rule 1 = **Not Fraud.**
2. **Rule 2:** Fraud if description contains the word `"fraud"`.
   - You didn't mention `"fraud"` in the description.
   - → Rule 2 = **Not Fraud.**
3. **Rule 3:** Fraud if location = `"Japan"`.
   - You didn't specify Japan.
   - → Rule 3 = **Not Fraud.**

Complications of Static AI

- Unknown potential input
- Limited by programmers' knowledge of input
- Easy to exploit since context is usually needed to spot the fraud

Is there a better way that doesn’t involve endless if else statements?

Quick AI Training Overview – Street Safety

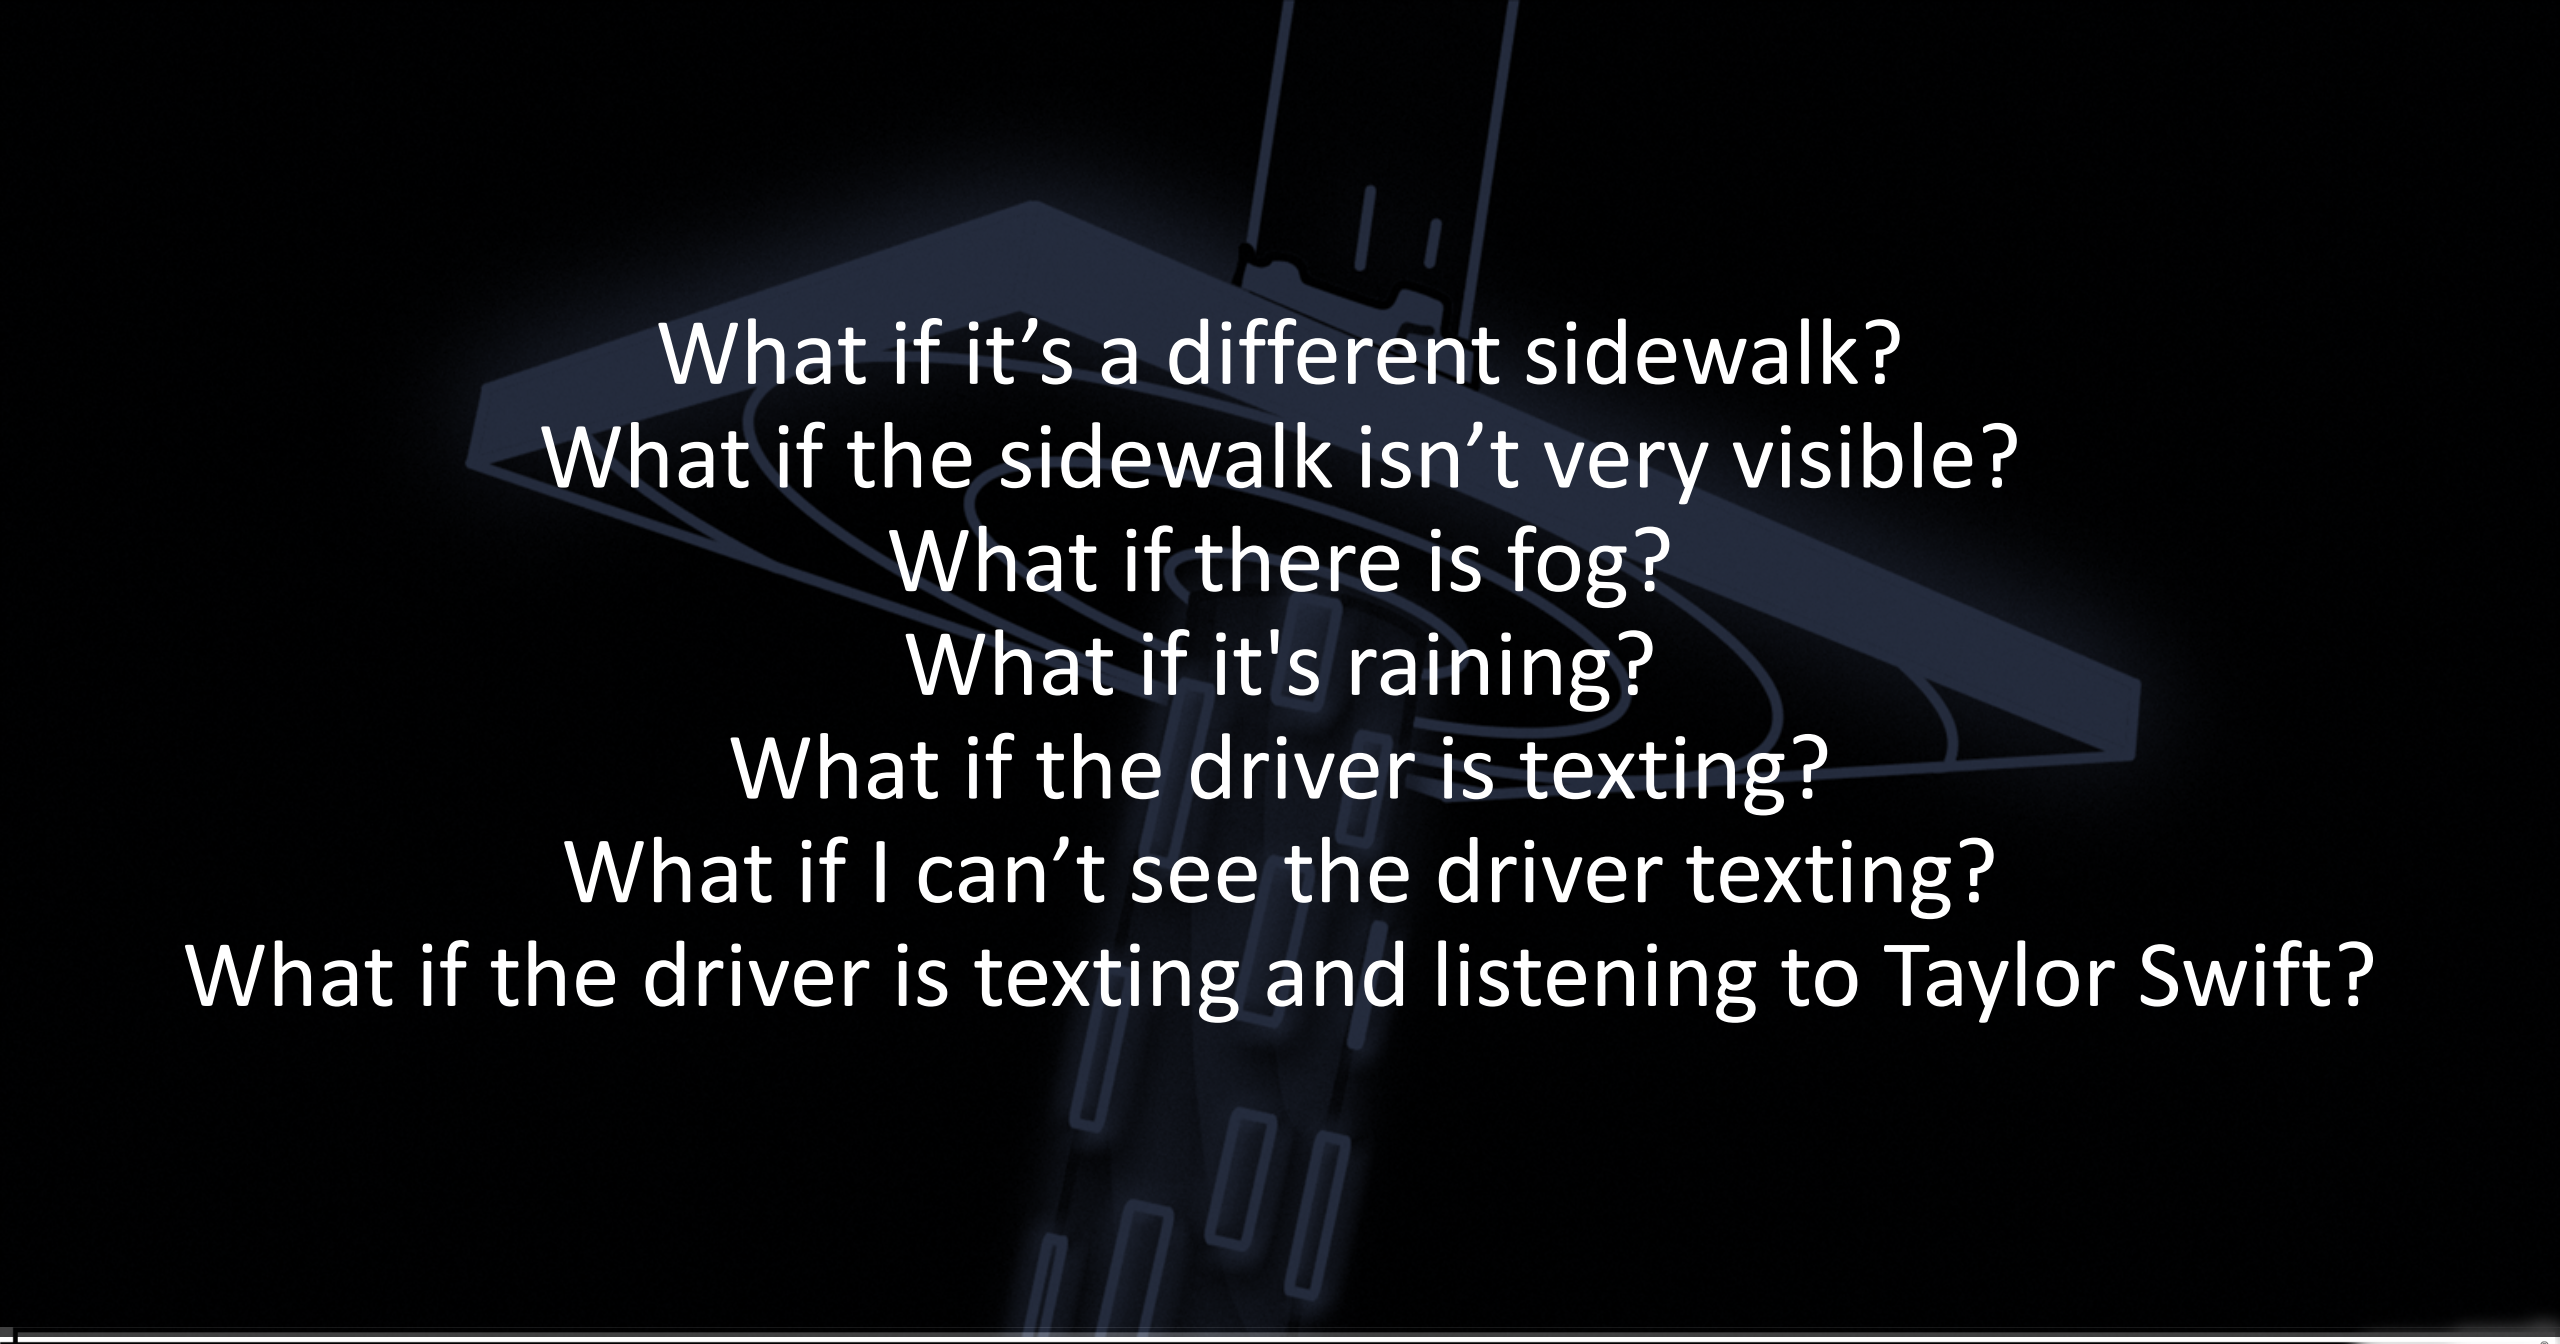
What if it’s a different sidewalk?
What if the sidewalk isn’t very visible?
What if there is fog?
What if it's raining?
What if the driver is texting?
What if I can’t see the driver texting?
What if the driver is texting and listening to Taylor Swift?

3 Parts of Machine Learning

Closest Car	Speed	Outcome
50 feet	10 mph	Alive
100 feet	30 mph	Unalive
200 feet	45 mph	Alive
150 feet	29 mph	Unalive
400 feet	10 mph	Alive

- Data set is collected
- AI is trained on dataset with an algorithm creating an AI model
- AI Model handles input and gives confidence output

Backpropagation

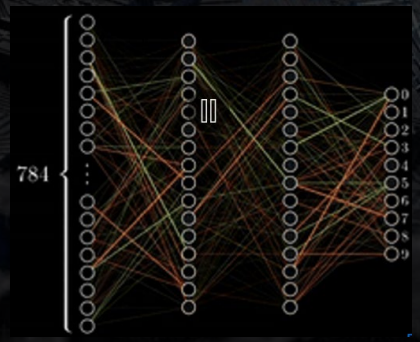

How AI Learns - Multilayer Perceptron

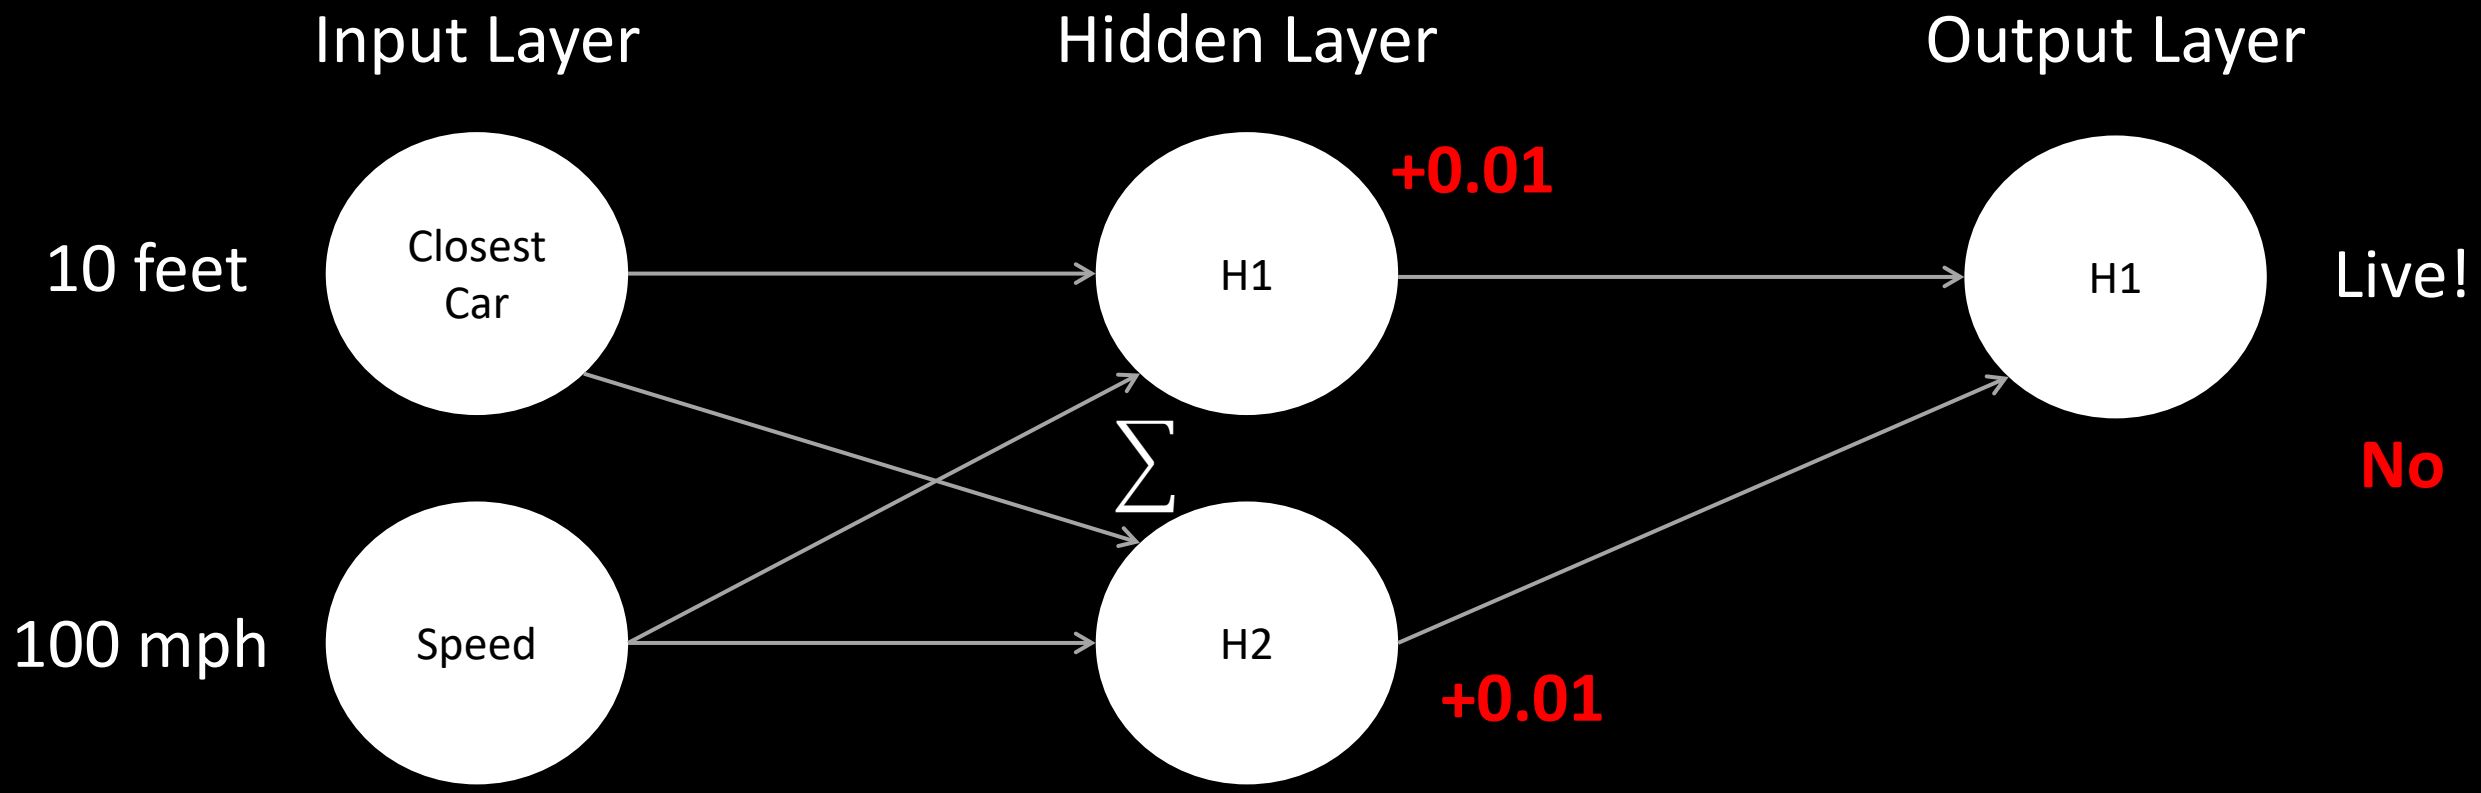

Backpropagation Helps Train Model

How AI Learns - Multilayer Perceptron

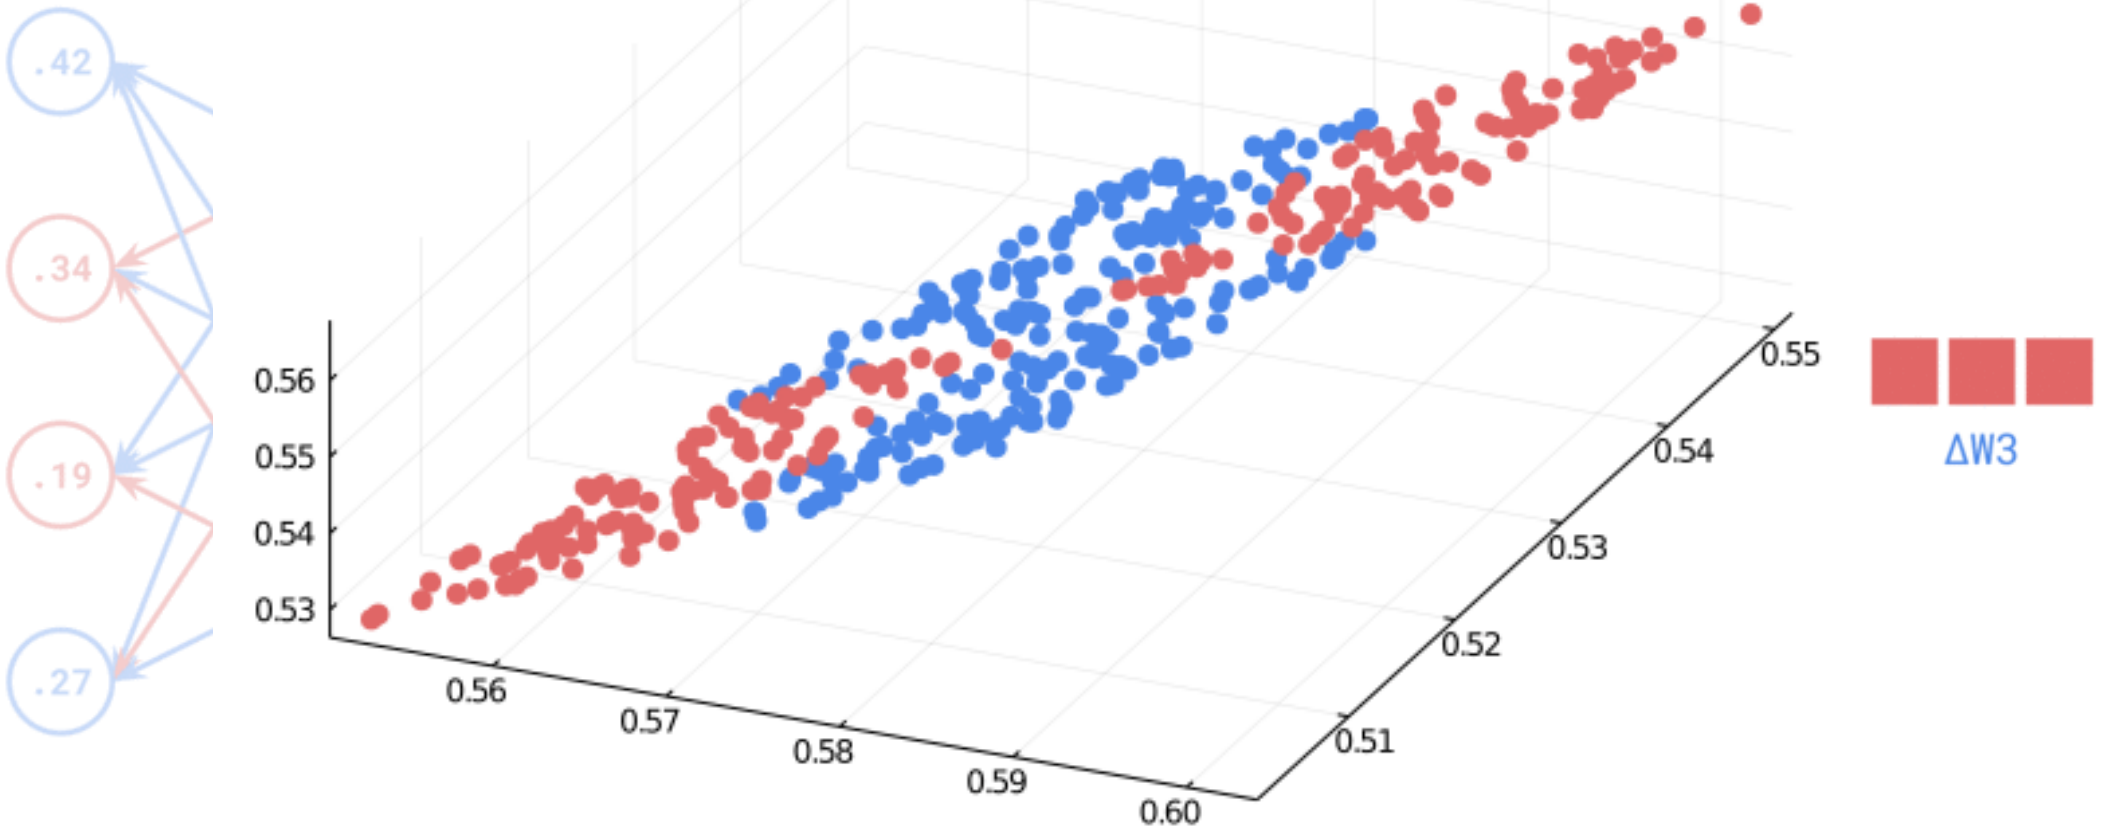

Credit to: https://medium.com/the-feynman-journal/what-makes-backpropagation-so-elegant-657f3afbbd

Training in progress. . .

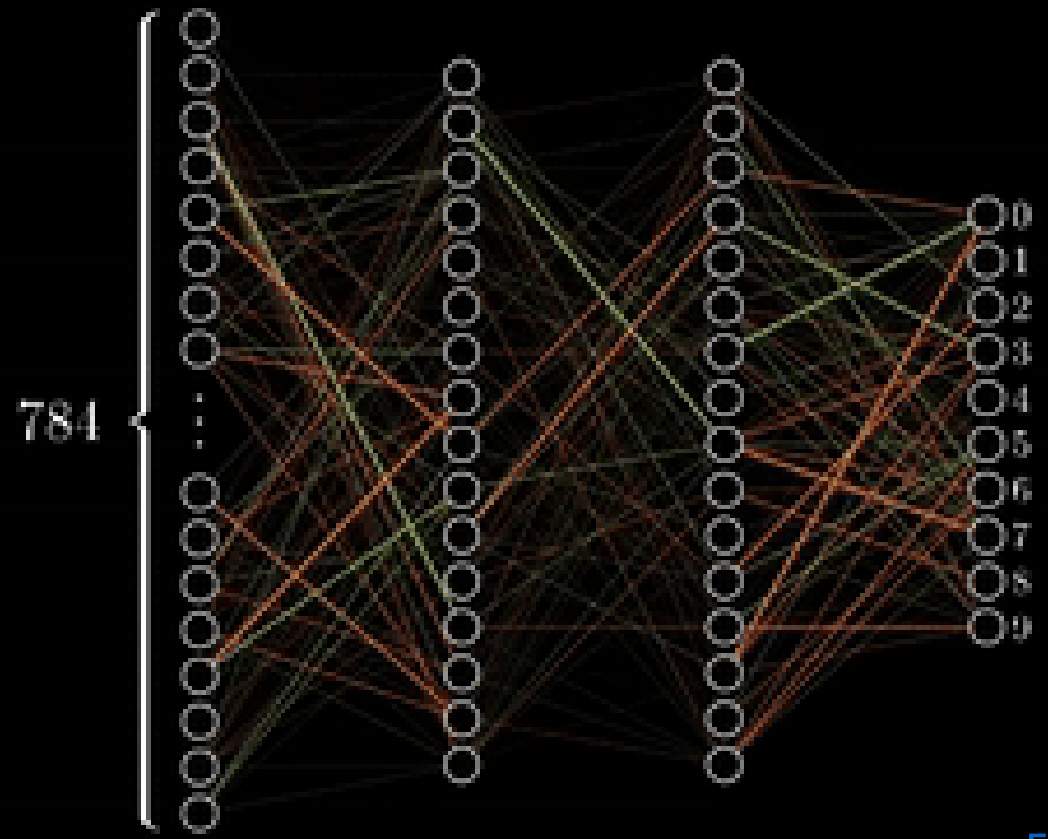

Flawnson Tong – Medium

Shifting From Deterministic to Probabilistic AI

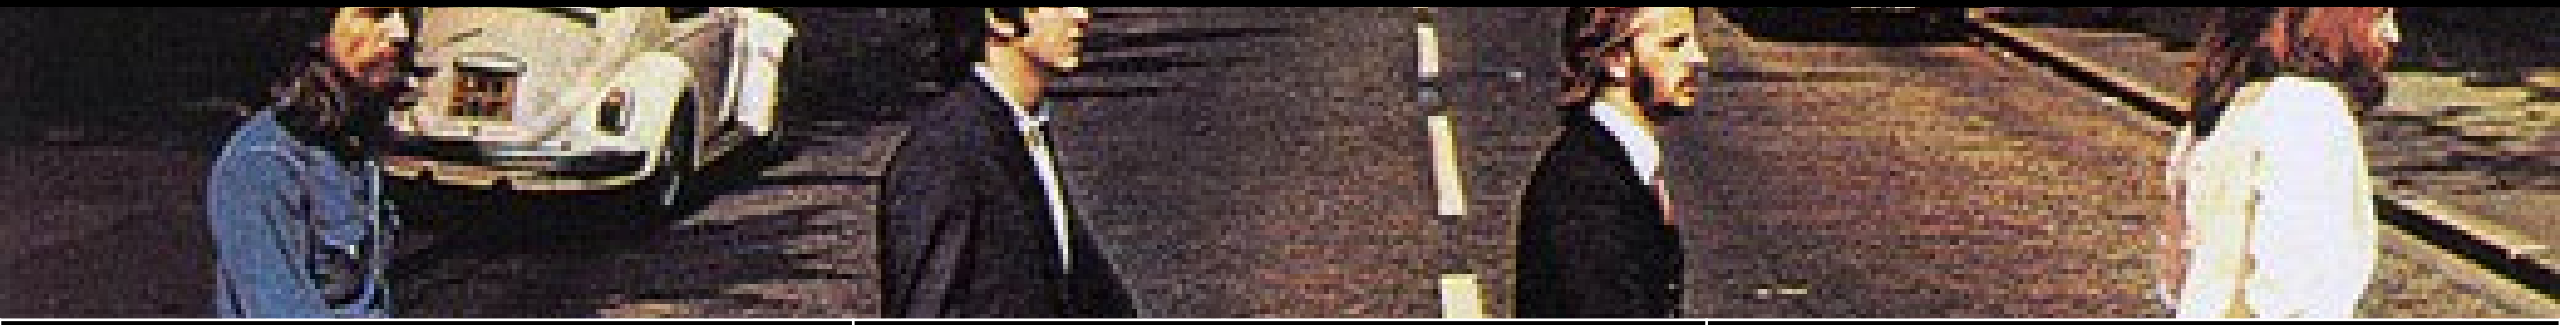

Closest Car	Speed	Outcome
50 feet	10 mph	Alive
100 feet	30 mph	Unalive
200 feet	45 mph	Alive
150 feet	29 mph	Unalive
400 feet	10 mph	Alive

STEP 1: Train the Model

STEP 2: Deploy the Model

STEP 3: Use the Model on New Data

Data

Learning Algorithm

Model

Deploy

New Data

Model

Will John Lennon be safe?

Prediction

No

Closest Car	Speed	Outcome
50 feet	10 mph	Alive
100 feet	30 mph	Unalive
200 feet	45 mph	Alive
150 feet	29 mph	Unalive
400 feet	10 mph	Alive

Backpropagation

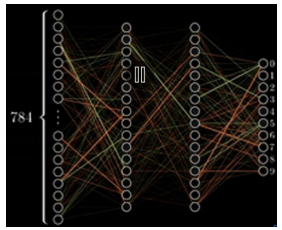

Closest Car	Speed	Outcome
50 feet	10 mph	Alive

Image Recognition Models-
100-10,000 nodes
Accuracy < 100%

Humans Recognize Objects
After 2 times ~100 accuracy

3. Where AI Can Replace Humans

Artificial Intelligence Thought Experiment

You work in a paint shop with an AI assistant. To make a new paint your AI suggests the following chemicals to be used

which do you pick?:

A. Acetaldehyde
B. Potassium cyanide
C. Trichlorophenoxyacetic acid

Questions to Ask When Analyzing Tech Disruption Promises

1) Is there a simpler version of this technology?
2) Is the general population using a simpler version?
3) Are their legal repercussions if something goes wrong?
4) Are there any social adoption obstacles?

Case Study Claims –Customer Service

1) Is there a simpler version of this technology?
2) Is the general population using a simpler version?
3) Are their legal repercussions if something goes wrong?
4) Are there any social adoption obstacles?

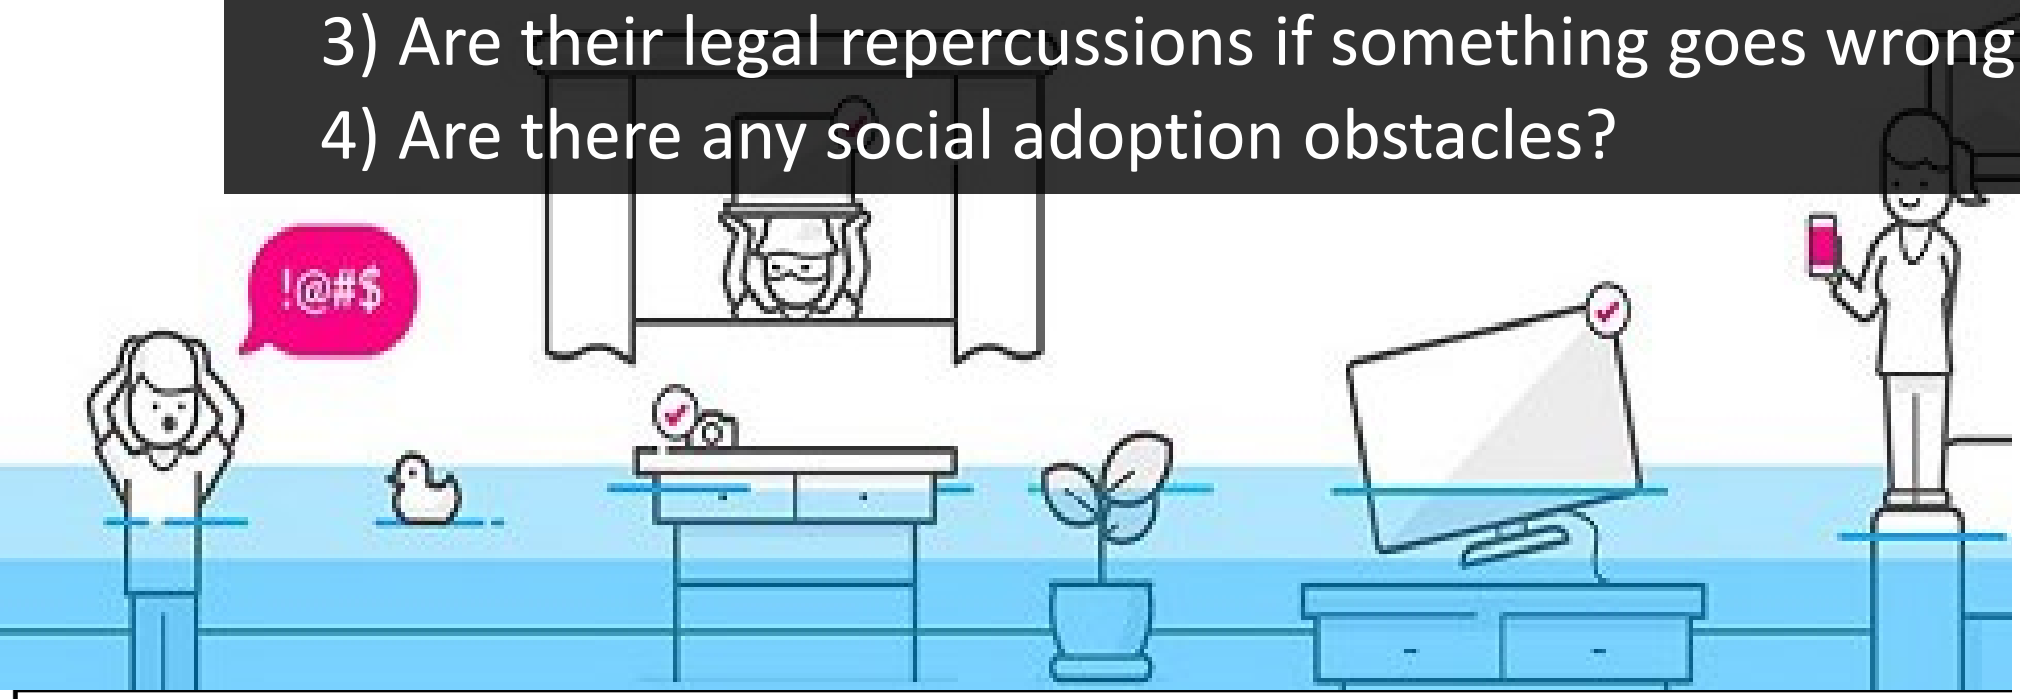

Lemonade Annual Operating Income (Millions of US $)	
2024	$-204
2023	$-230
2022	$-295
2021	$-234
2020	$-121
2019	$-108
2018	$-53

Customer Service - Lemonade Lawsuits

Case	Privacy Allegation	Coverage Period	Outcome / Status
La Febre v. Lemonade, Inc.	Unconsented sharing of PHI/PII with third parties	Mar 2021 – Sept 2023	$4.995M settlement; individual payouts; PHI sharing halted
Clarke v. Lemonade Inc. (Biometric BIPA)	Non-consensual biometric data collection via video	Jun 2019 – May 2021	$4M settlement; deletion of biometric data, no future BIPA violations
Driver's License Exposure Suit	Exposure of driver's license numbers via quote platform	Apr 2023 – Sept 2024	Ongoing federal litigation; seeking damages and privacy safeguards

Case Study CSR Chatbots – After Hours

1) Is there a simpler version of this technology?
2) Is the general population using a simpler version?
3) Are their legal repercussions if something goes wrong?
4) Are there any social adoption obstacles?

Case Study: Tech In the World of Insurance

Can the result be correctly verified by a non-technical person?

Would you like a schedule or blanket waiver of subrogation?

Complexity Requires Trust

Where Technology Can Replace Humans

1. *Where AI makes 50% less errors than humans**
2. *A non-technical person can understand the AI output*
3. *Where legal repercussions are not catastrophic*

Examples:

- Cars
- Images
- Writing
- Revenue projections
- P/L statement

- Appointment Booking
- Training materials
- Customer Onboarding

Four Patterns Where GenAI Implementations Fail:

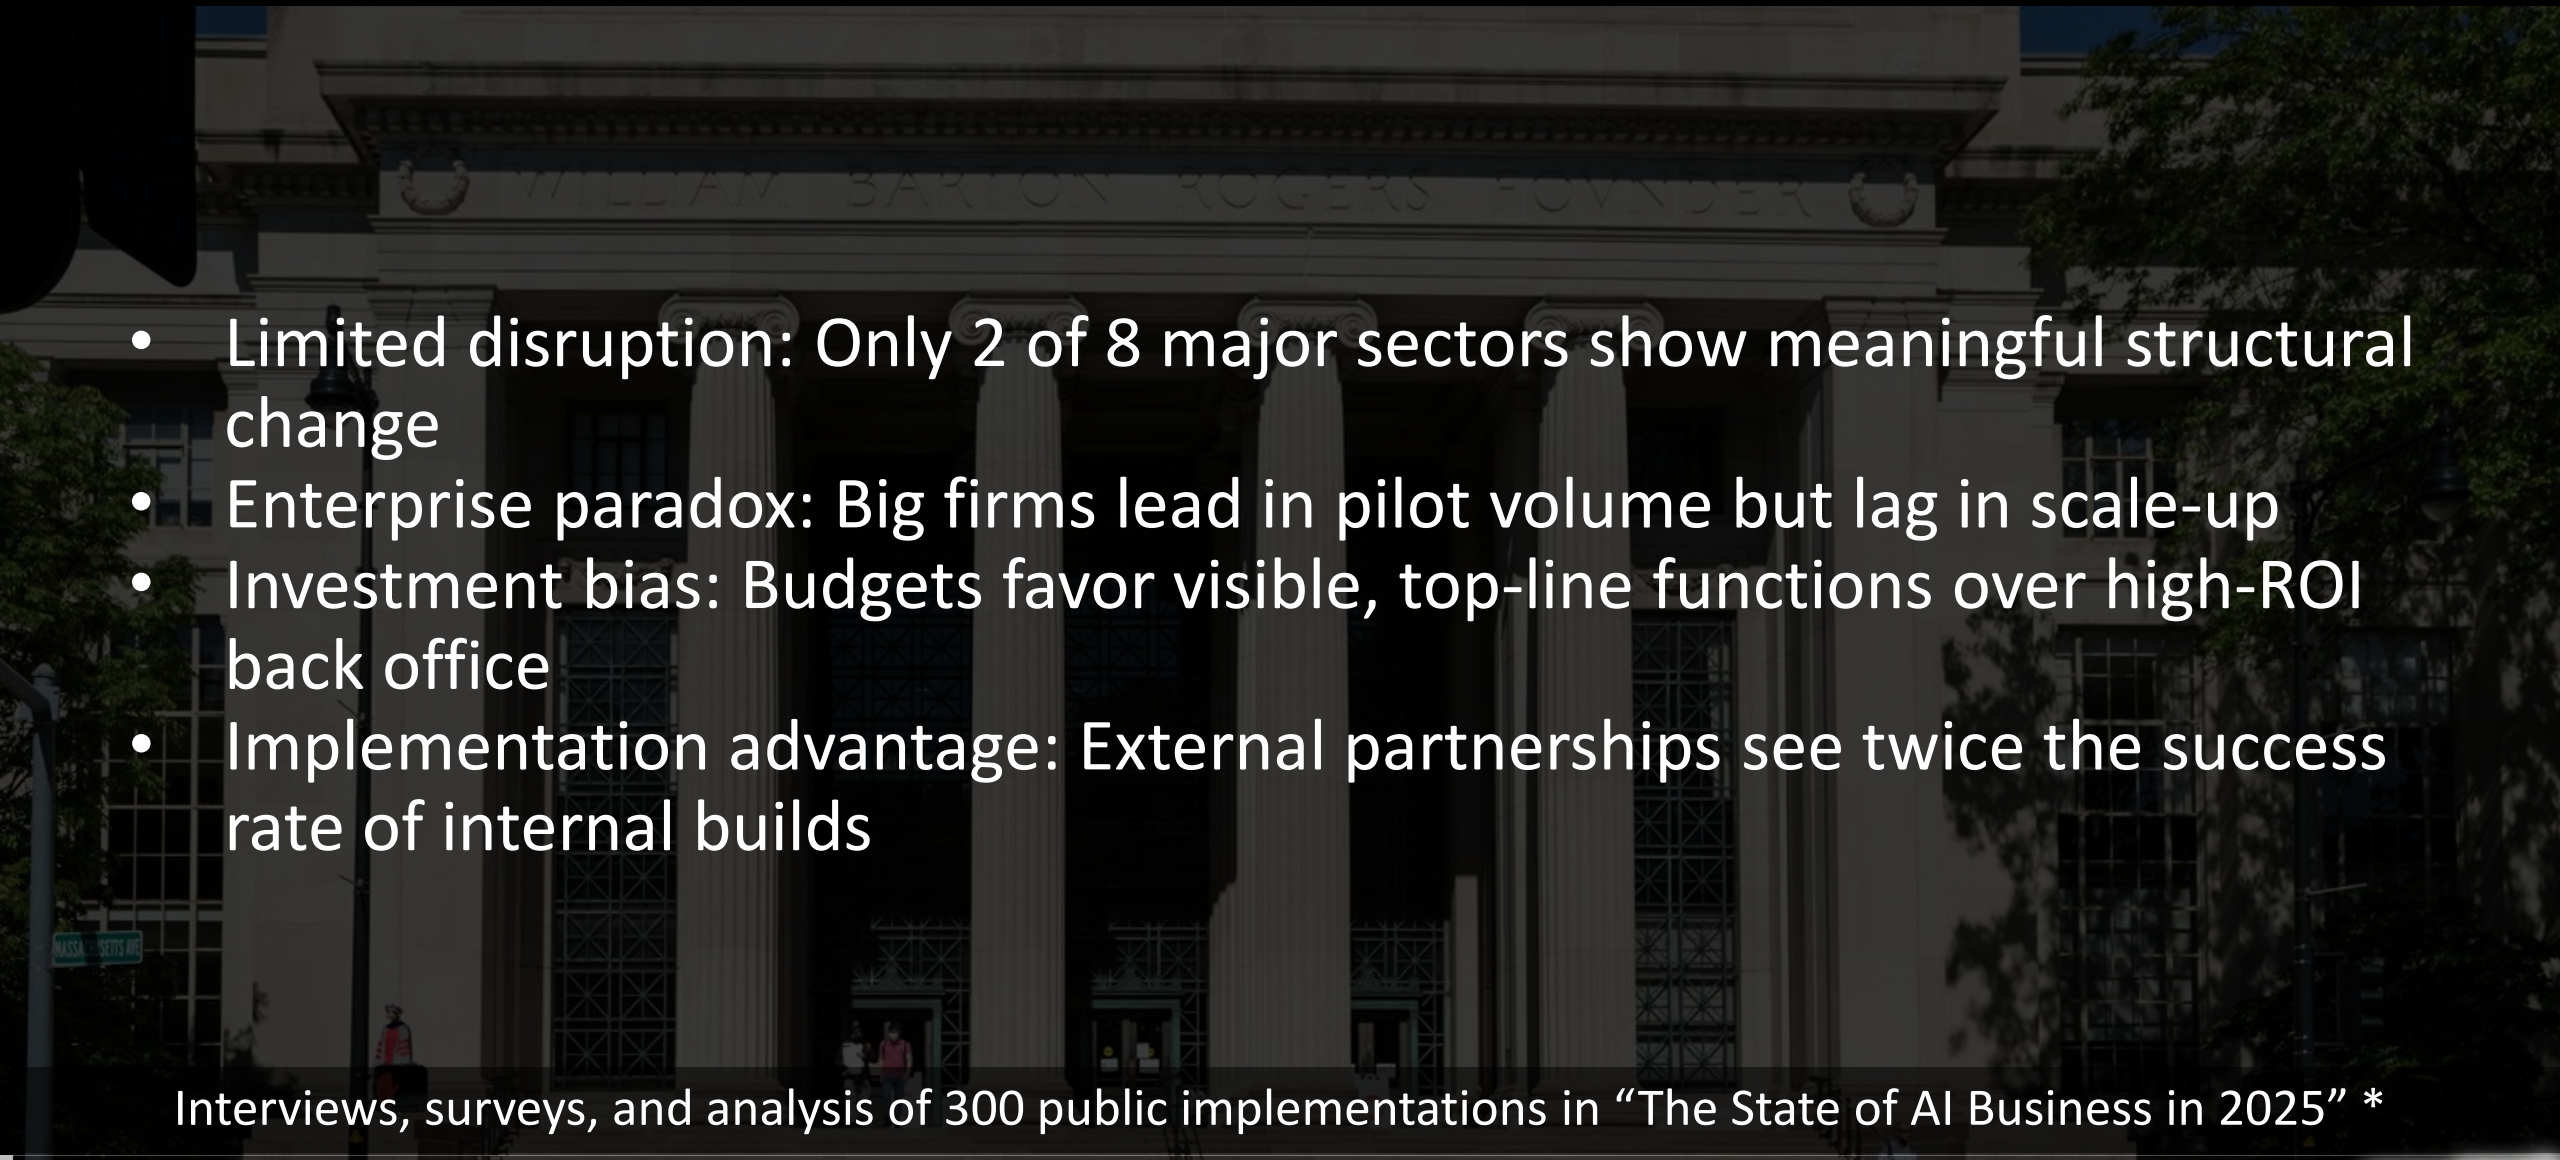

- Limited disruption: Only 2 of 8 major sectors show meaningful structural change
- Enterprise paradox: Big firms lead in pilot volume but lag in scale-up
- Investment bias: Budgets favor visible, top-line functions over high-ROI back office
- Implementation advantage: External partnerships see twice the success rate of internal builds

Interviews, surveys, and analysis of 300 public implementations in “The State of AI Business in 2025” *

Which Industries Are Showing Structural Change

Exhibit: GenAI disruption varies sharply by industry

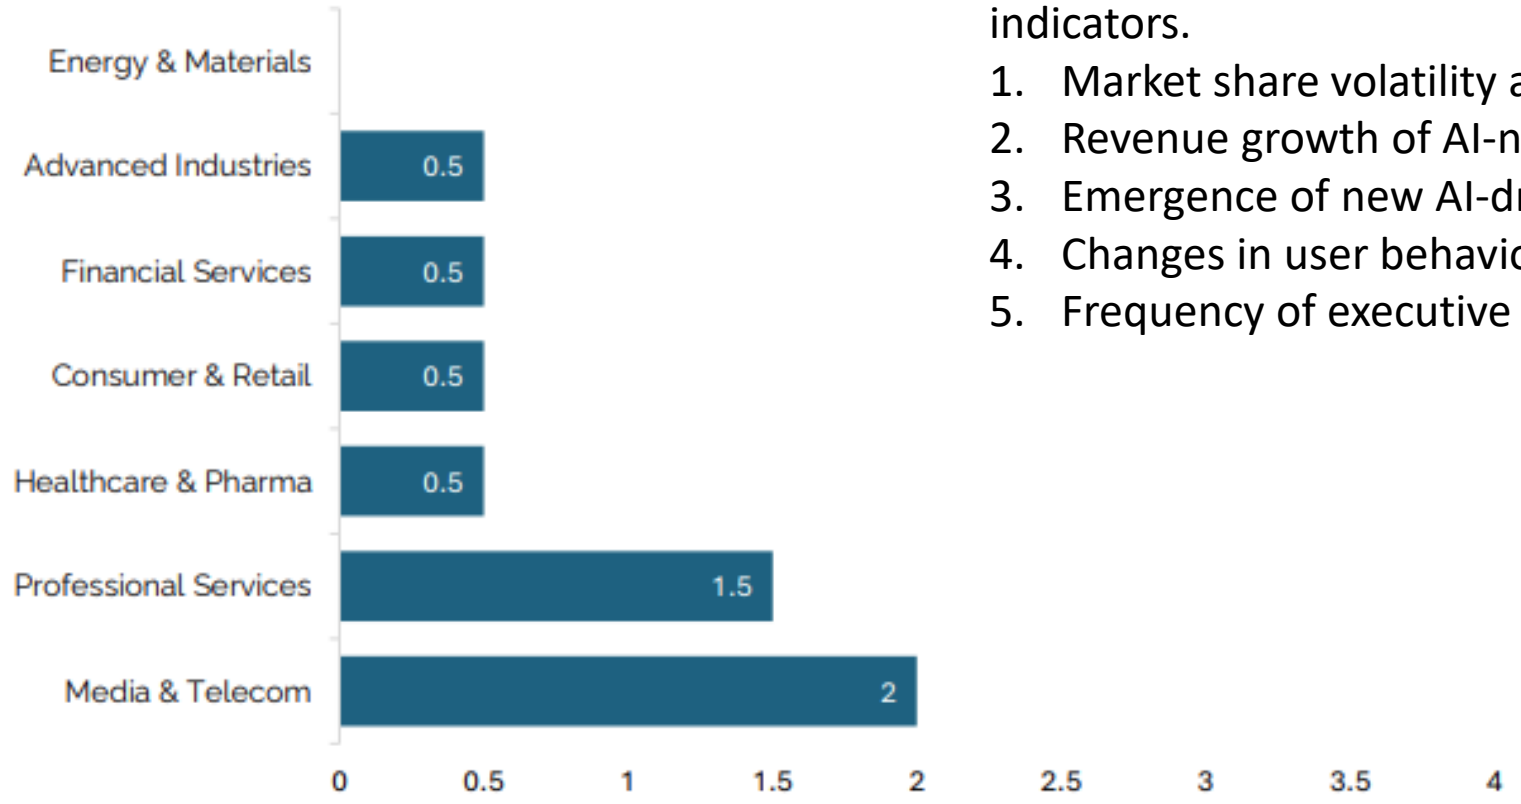

Each industry was scored from 0 to 5 based on five observable indicators.

1. Market share volatility among top incumbents (22' to 25')
2. Revenue growth of AI-native firms founded after 2020
3. Emergence of new AI-driven business models
4. Changes in user behavior attributable to GenAI
5. Frequency of executive org changes attributed to AI tooling

These scores represent normalized averages across five dimensions, triangulated from public indicators and interview-derived assessments. Alternative weighting schemes were tested to confirm consistency of industry rankings**

How C-Suite Spends Funds on AI

You are given $100 to spend on AI development, how do you distribute those funds between different tasks?

(Ie: cold calling, AI secretaries, chatbots, contract tagging)

Fund Allocation:
Survey based on 153 senior leaders that currently use AI in their firms

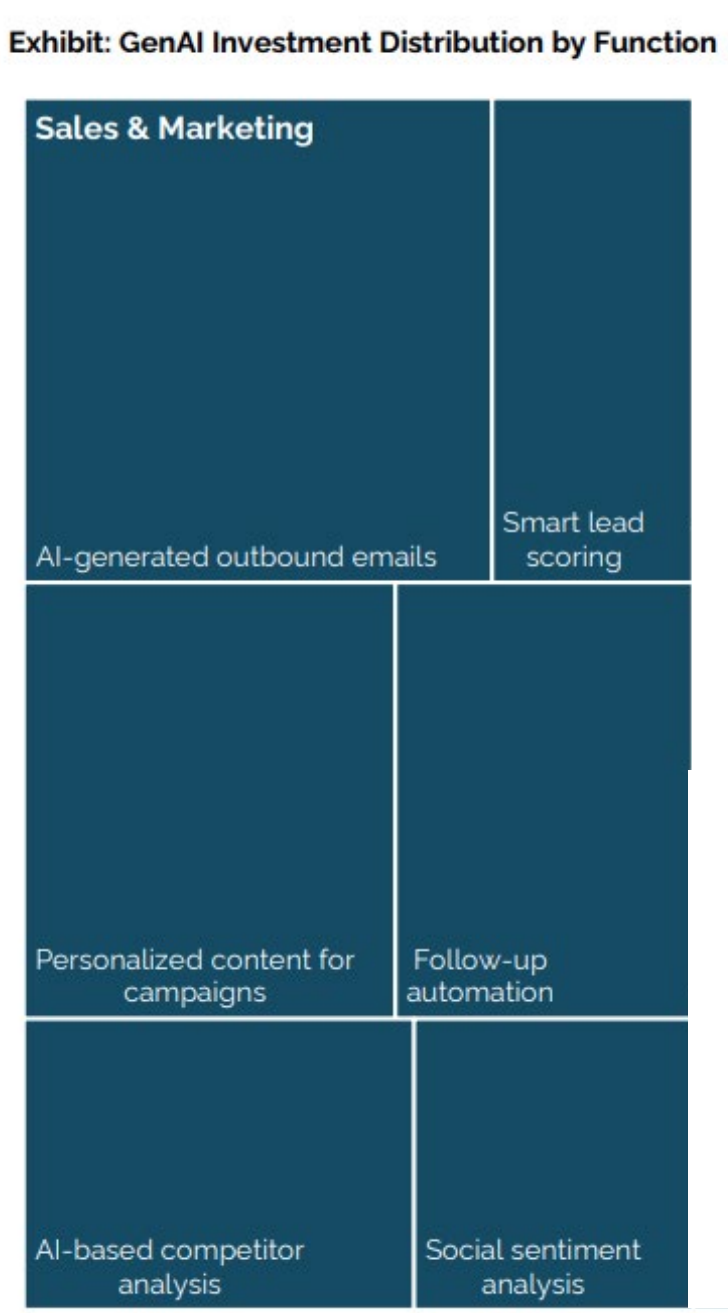

- 70% funds go to sales and marketing functions of AI budget allocation in survey.
- Back office usually has higher ROI, but smaller allocation

Benefits of AI

Front-office wins:

- Lead qualification speed: 40% faster
- Customer retention: 10% improvement through AI-powered follow-ups and messaging

Back-office wins:

- BPO elimination: $2-10M annually in customer service and document processing
- Agency spend reduction: 30% decrease in external creative and content costs
- Risk checks for financial services: $1M saved annually on outsourced risk management

Interviews, surveys, and analysis of 300 public implementations in “The State of AI Business in 2025” *

How Successful Firms Benefit – ROI Paradox

- Most implementations don't drive headcount reduction
- Organizations experience workforce impacts in customer support, software engineering, and administrative functions. (ROI Paradox)
- Measurable savings from reduced BPO spending and external agency use, particularly in back-office operations.
- Improved customer retention and sales conversion through automated outreach and intelligent follow-up systems.

Interviews, surveys, and analysis of 300 public implementations in “The State of AI Business in 2025” *

Certainty

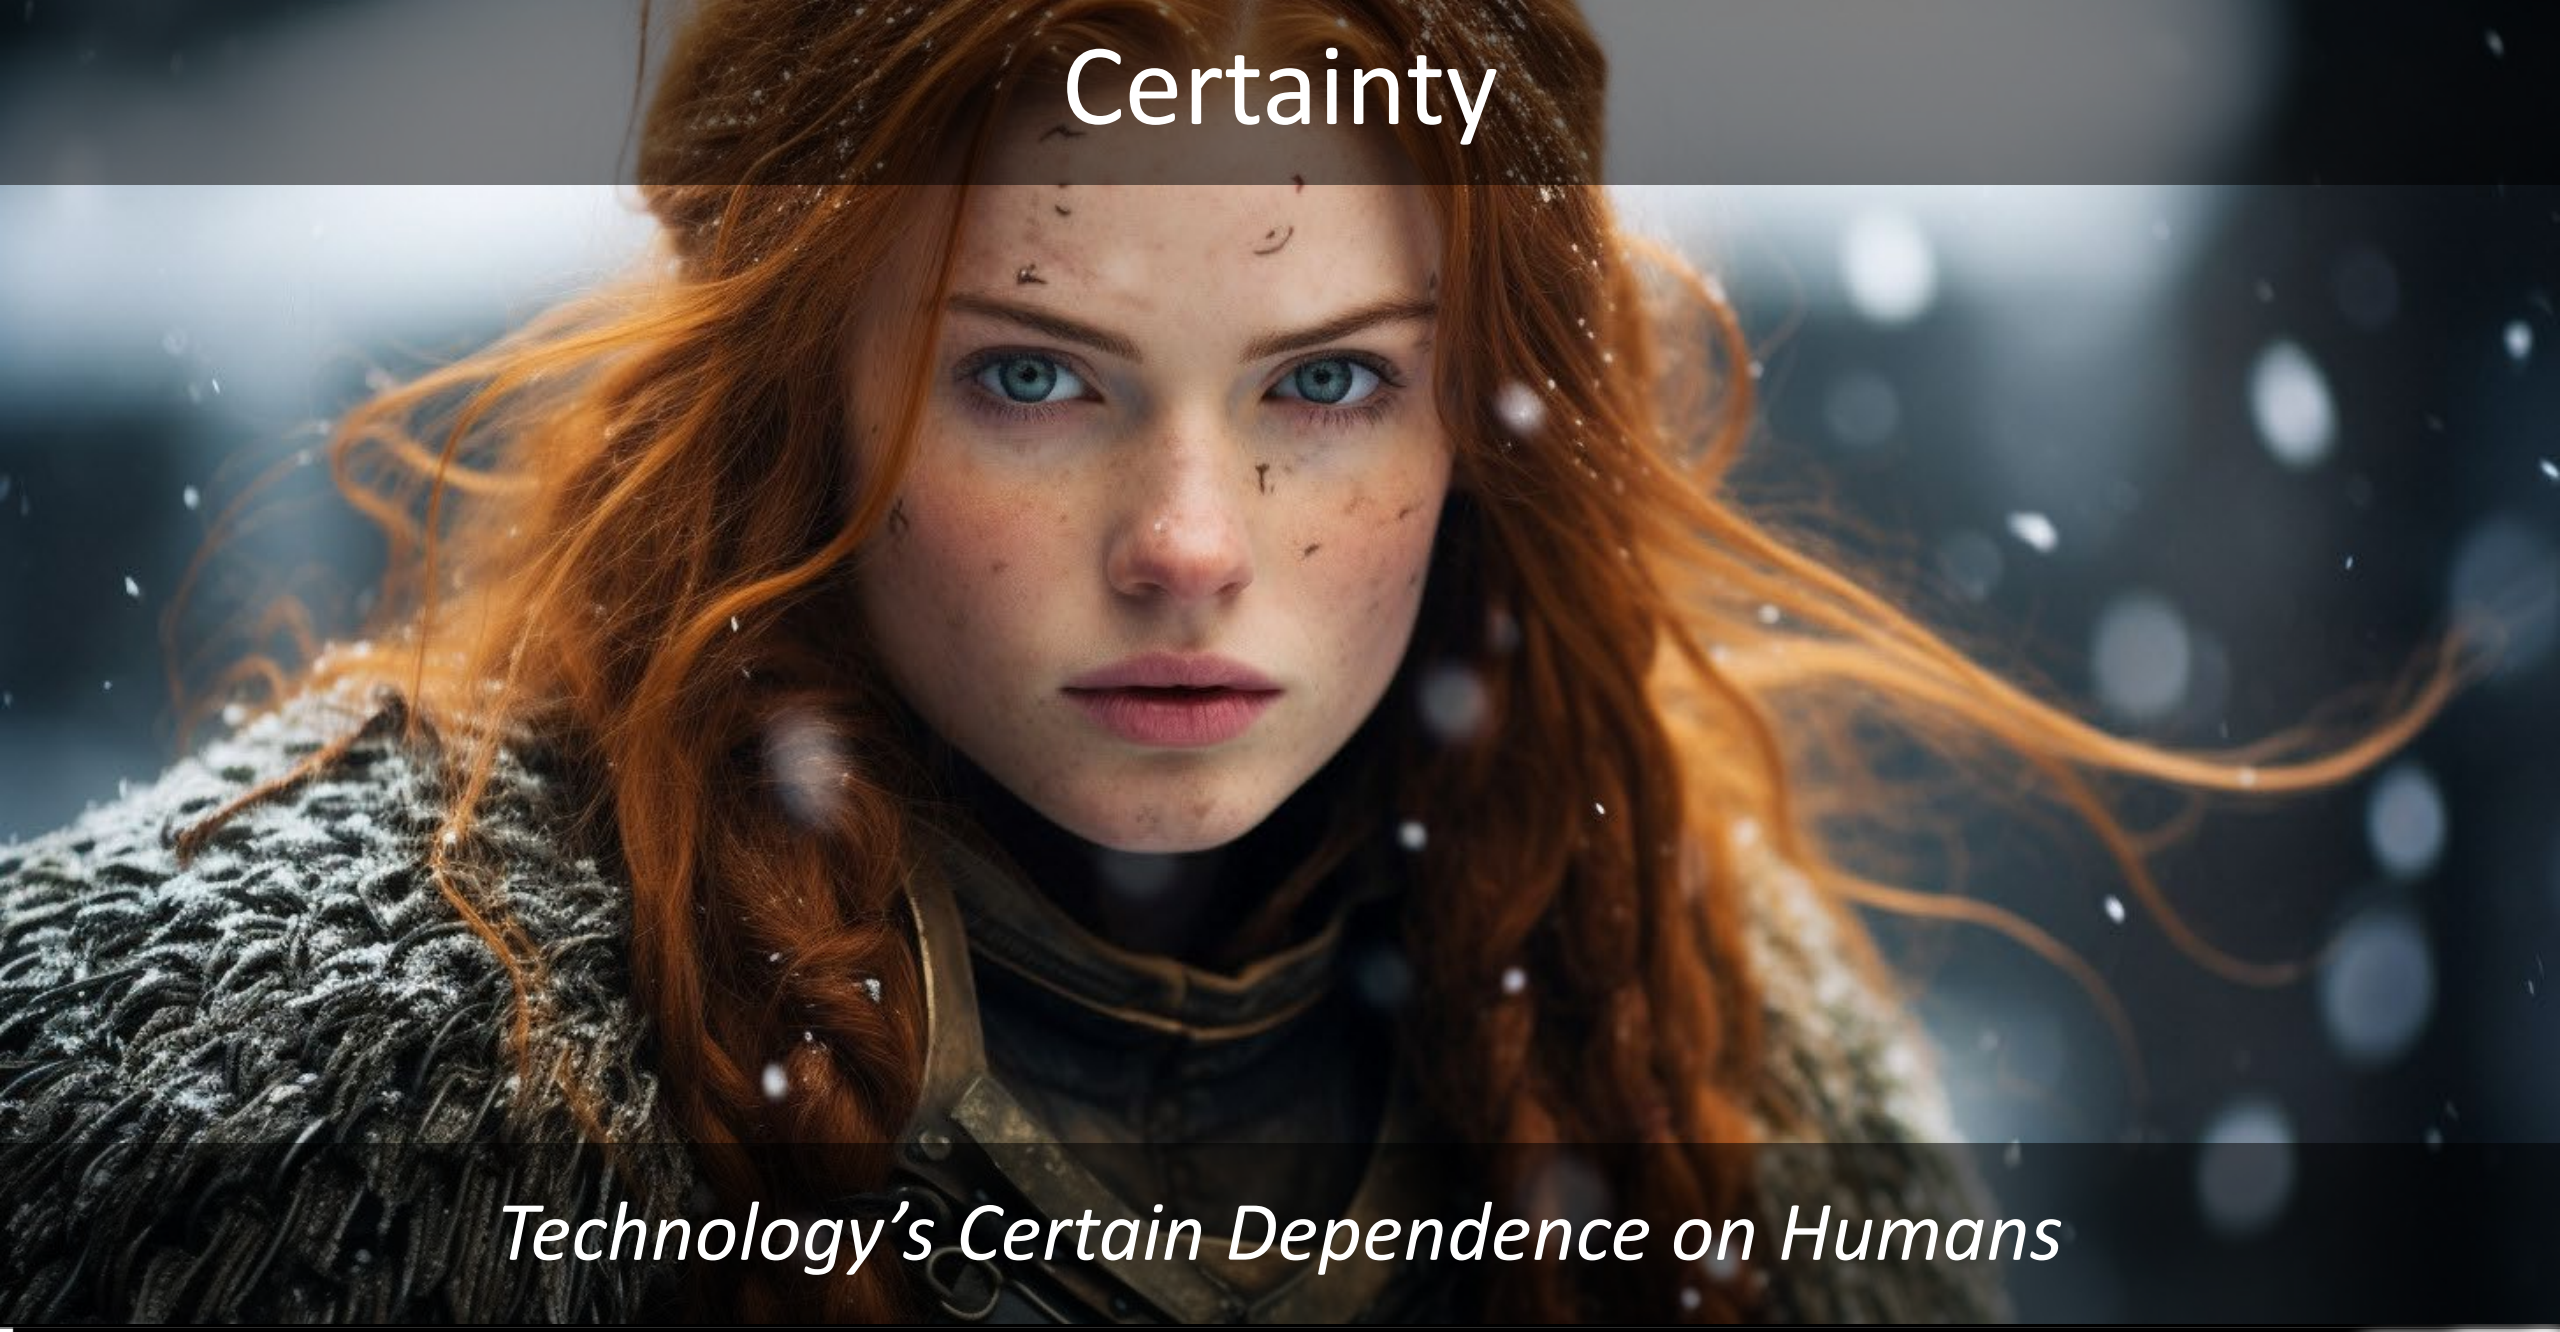

Technology's Certain Dependence on Humans

What do AIs and Berkeley Hippies Have in Common?

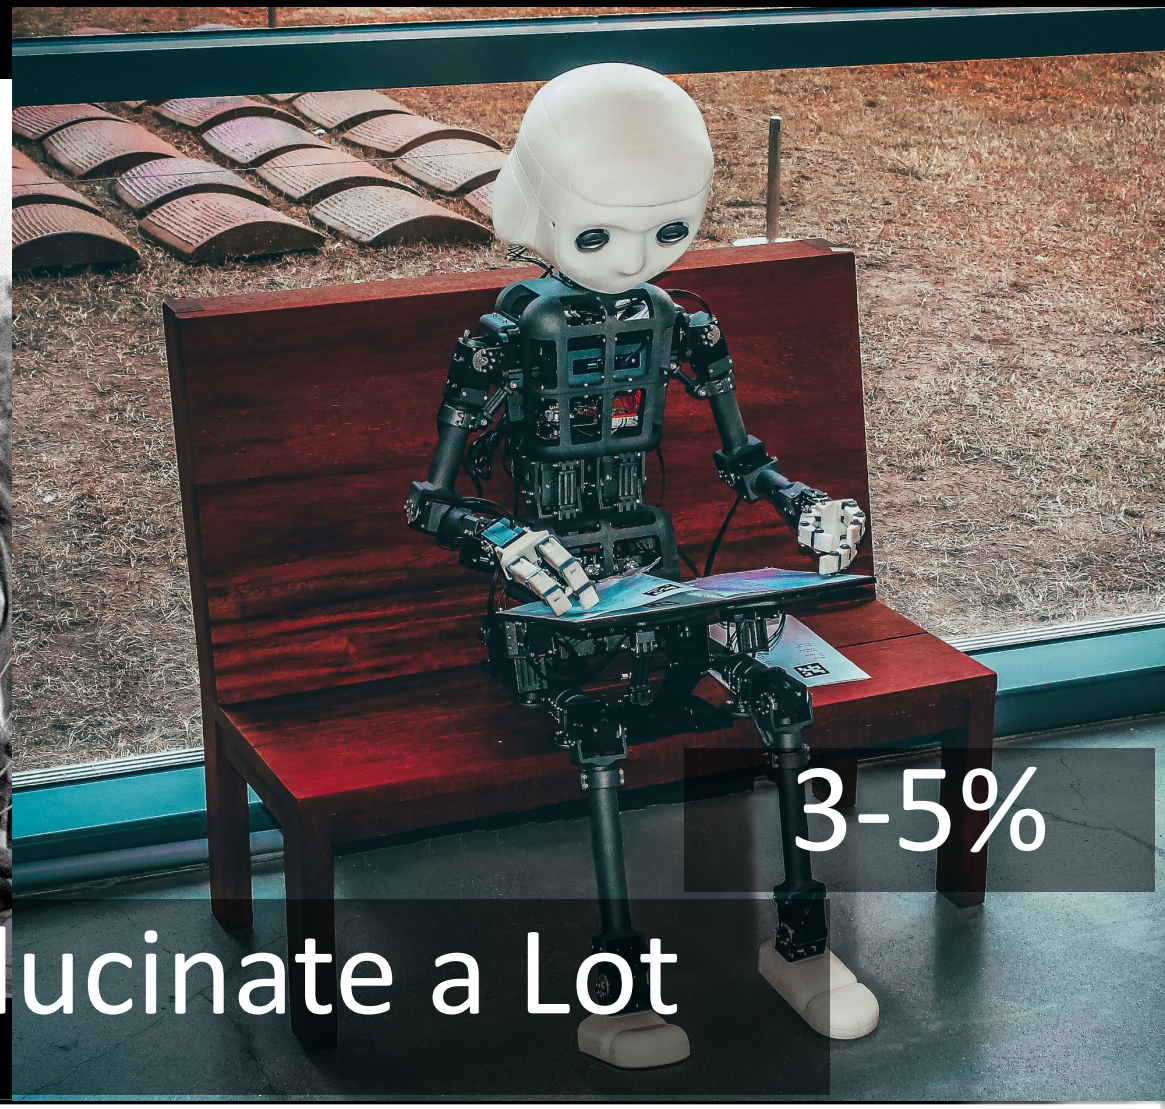

3-5%

They Both Hallucinate a Lot

Hallucination - Generative AI models produce responses that don't match reality

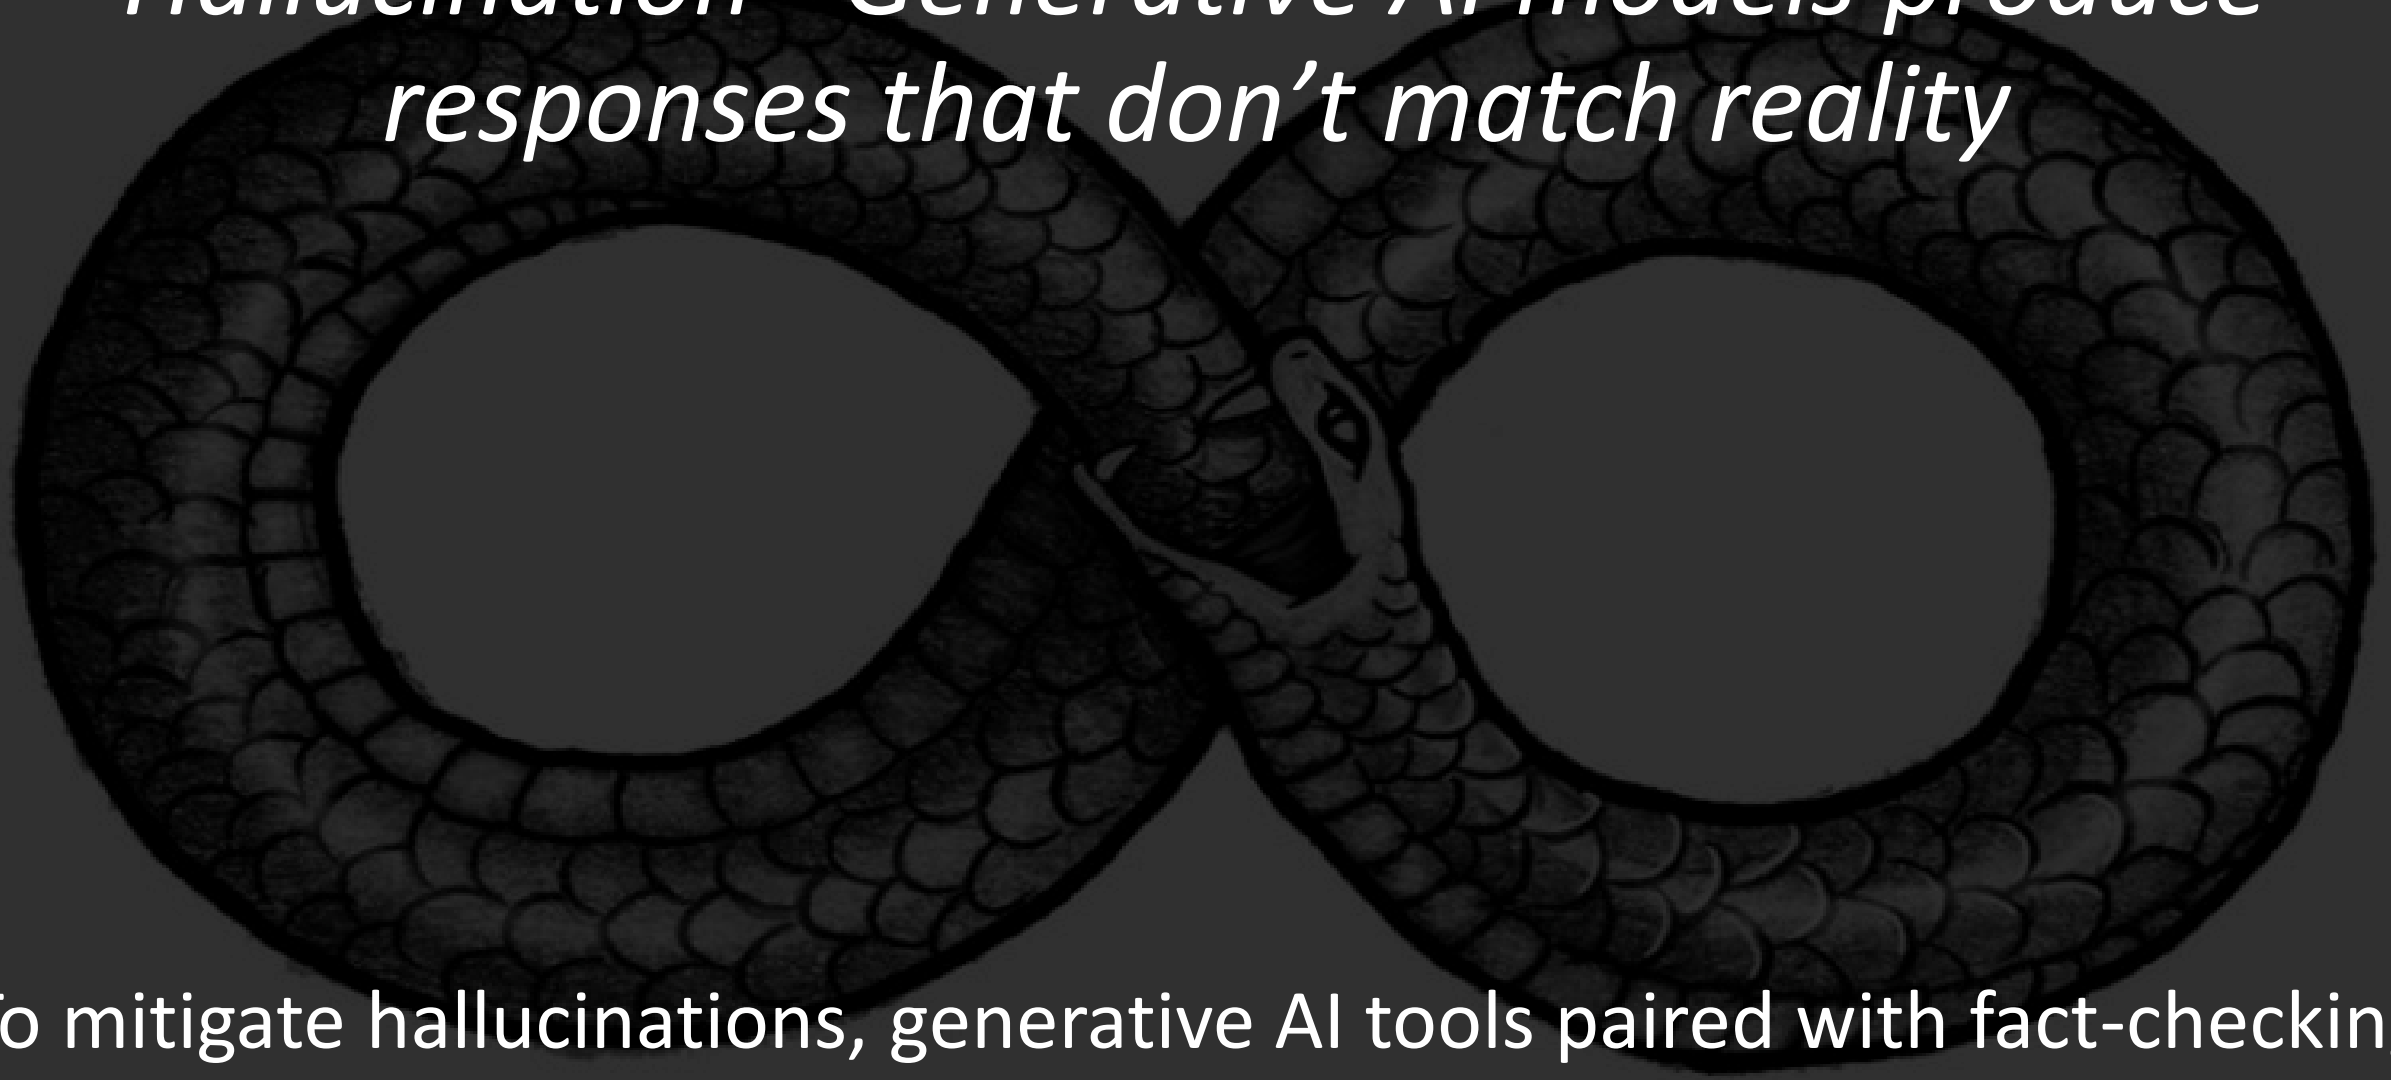

To mitigate hallucinations, generative AI tools paired with fact-checking systems that leave no chatbot unsupervised.

Examples of Hallucinations

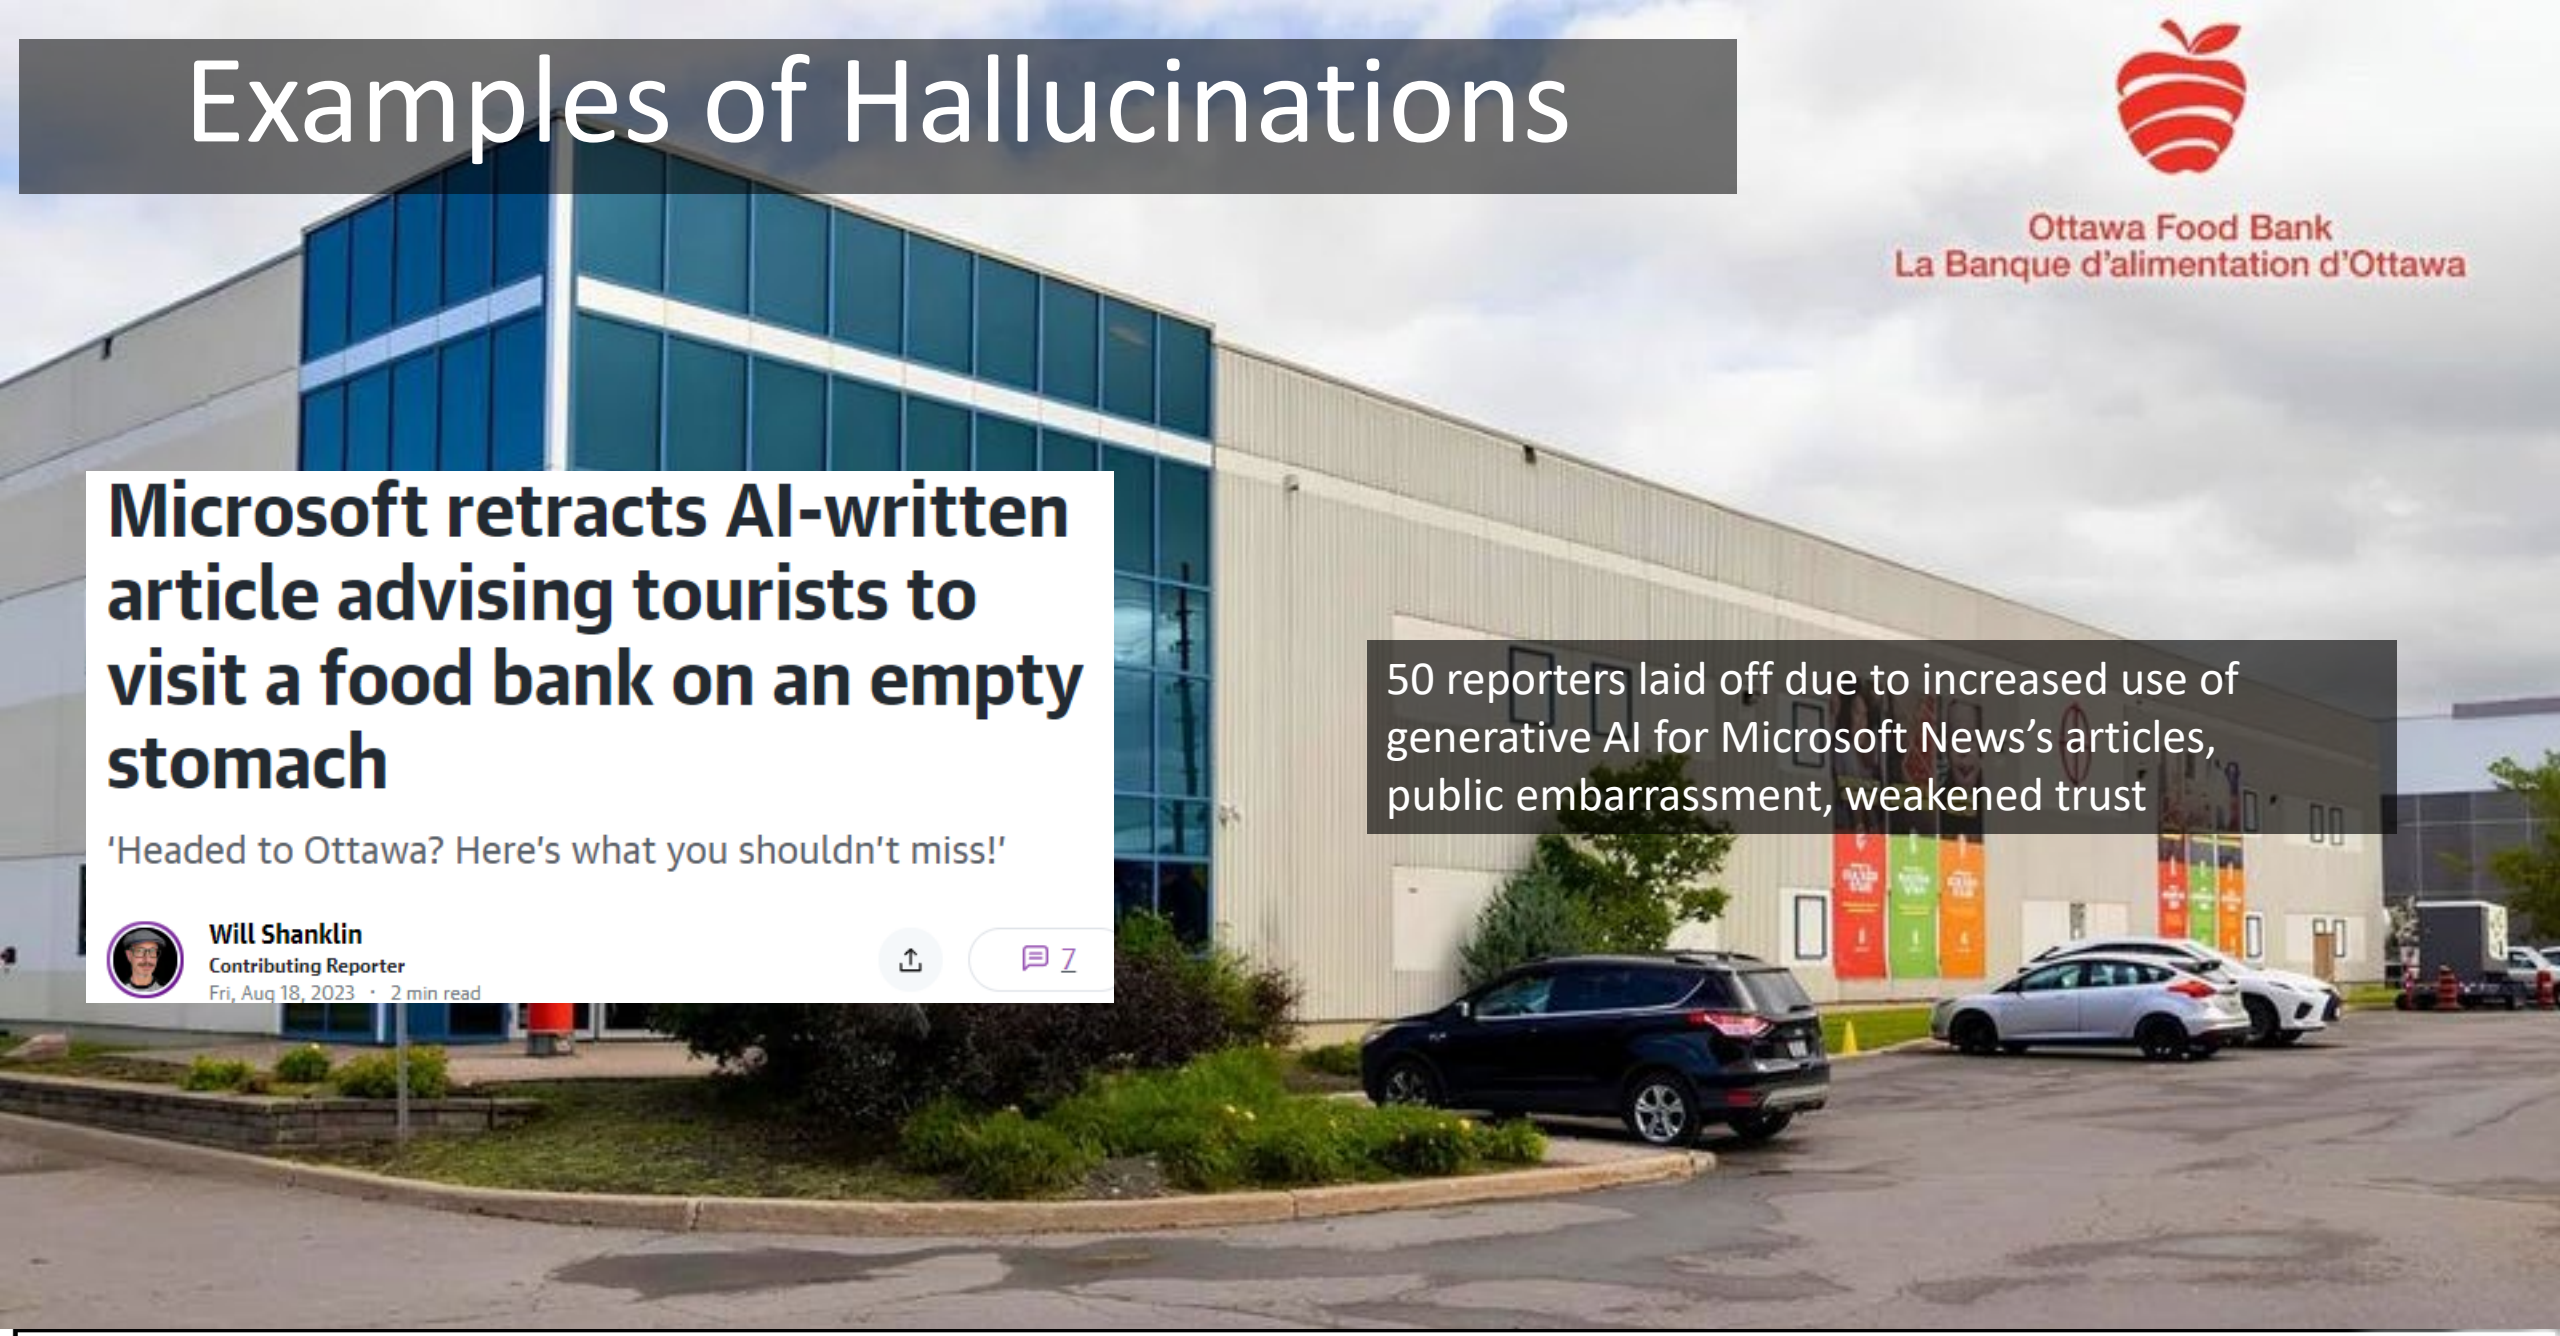

Microsoft retracts AI-written article advising tourists to visit a food bank on an empty stomach

'Headed to Ottawa? Here's what you shouldn't miss!'

Will Shanklin
Contributing Reporter
Fri, Aug 18, 2023 · 2 min read

7

50 reporters laid off due to increased use of generative AI for Microsoft News's articles, public embarrassment, weakened trust

Moffatt v. Air Canada: A Misrepresentation by an AI Chatbot

FEBRUARY 19, 2024

The recent decision in **Moffatt v. Air Canada**, 2024 BCCRT 149, represents a milestone in the expanding field of digital interactions and accountability. The case grapples with whether a company can be held liable for misleading information provided by an automated chatbot on its website. The decision held that a company can be liable for negligent misrepresentations made by a chatbot on a publicly available commercial website. The decision represents an incremental development of the law which previously has focused on the liability of persons for their pre-programmed automated tools.

Are companies financially responsible if a chatbot says incorrect information?

Mechanical Turk Paradox

1770

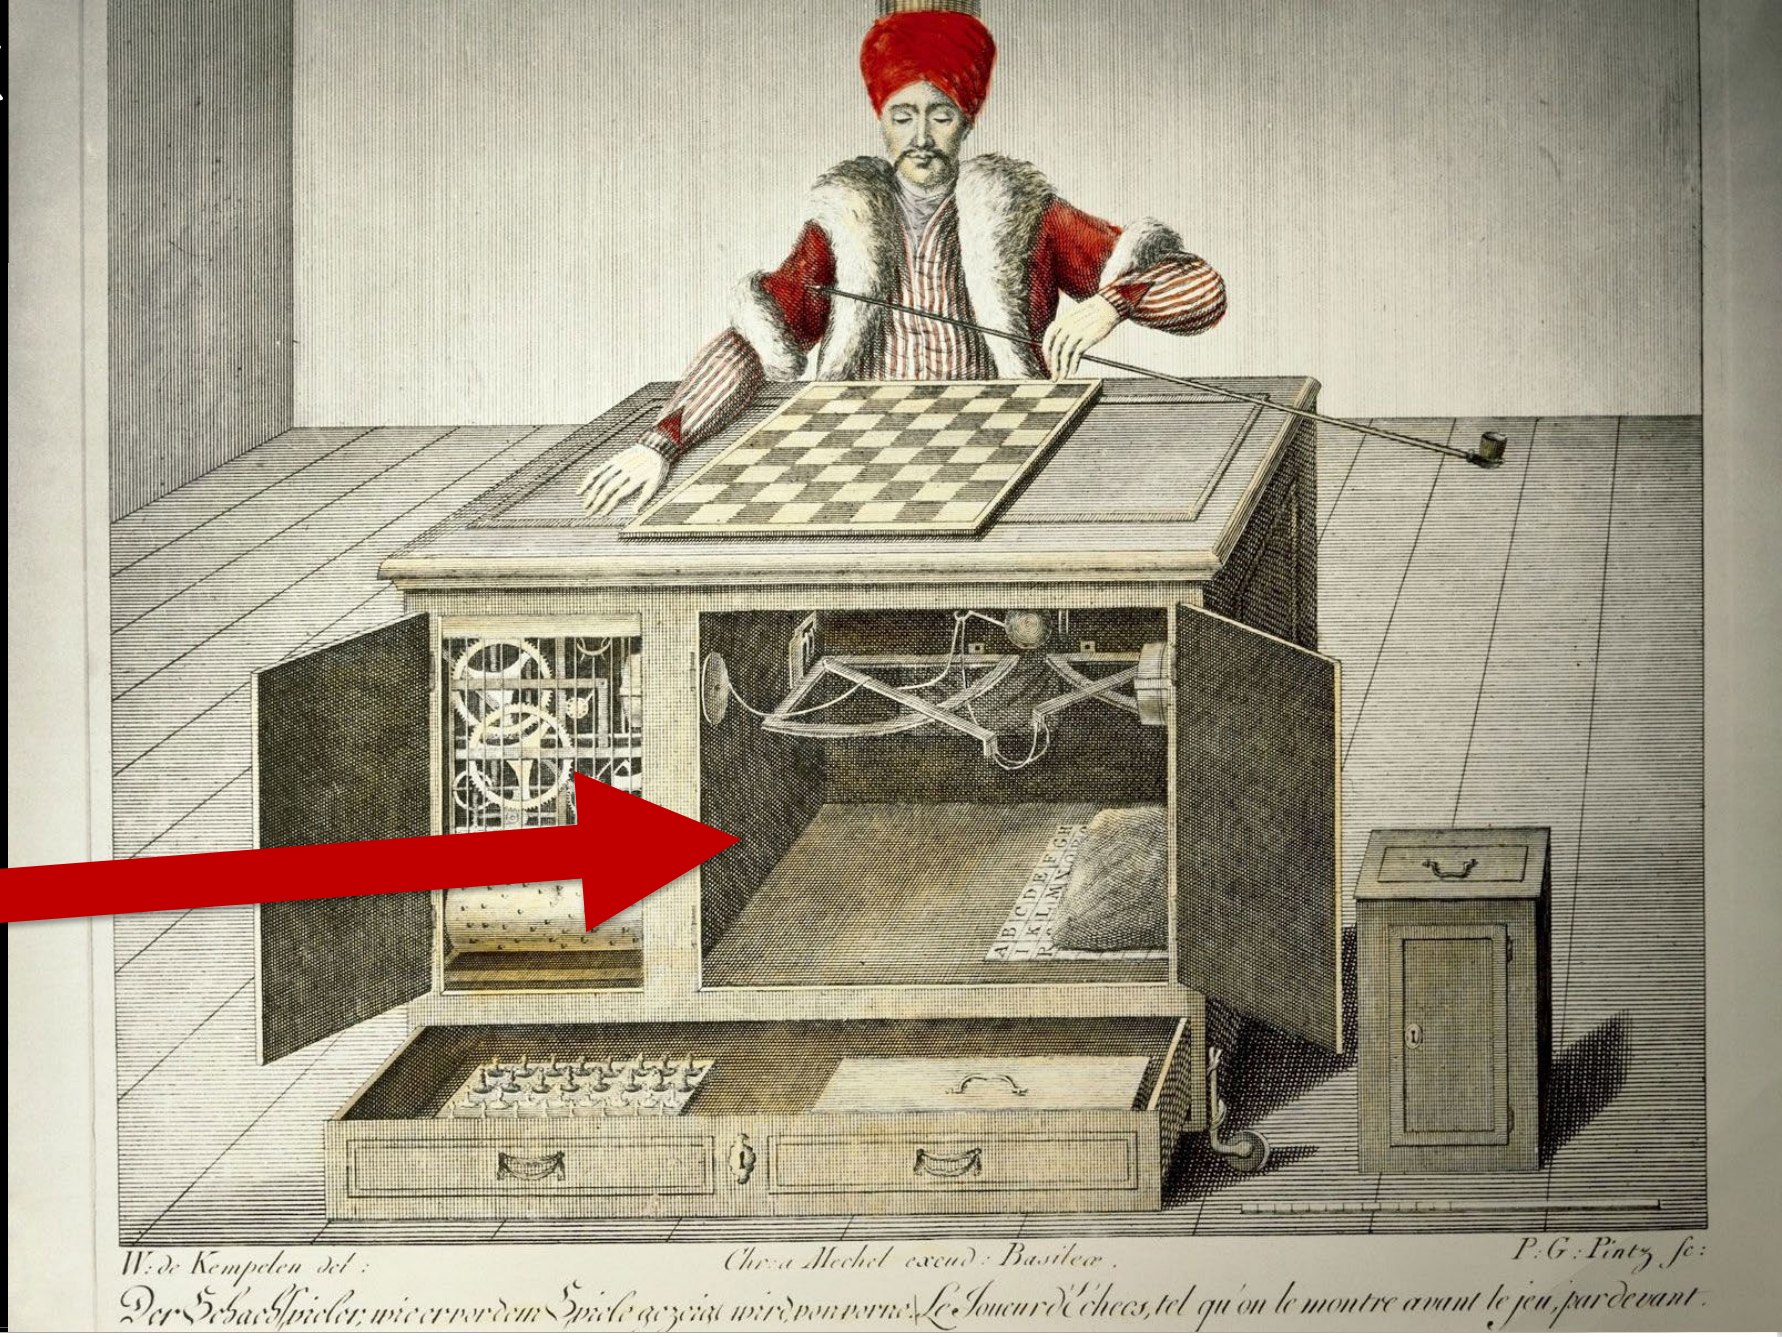

Opinion
Artificial intelligence (AI)

This article is more than **4 months old**

So, Amazon's 'AI-powered' cashier-free shops use a lot of … humans. Here's why that shouldn't surprise you

James Bridle

TECHNOLOGY EXECUTIVE COUNCIL

AI washing: A Microsoft VC's warning about dubious, and rising, corporate artificial intelligence claims

PUBLISHED FRI, MAY 10 2024•11:38 AM EDT

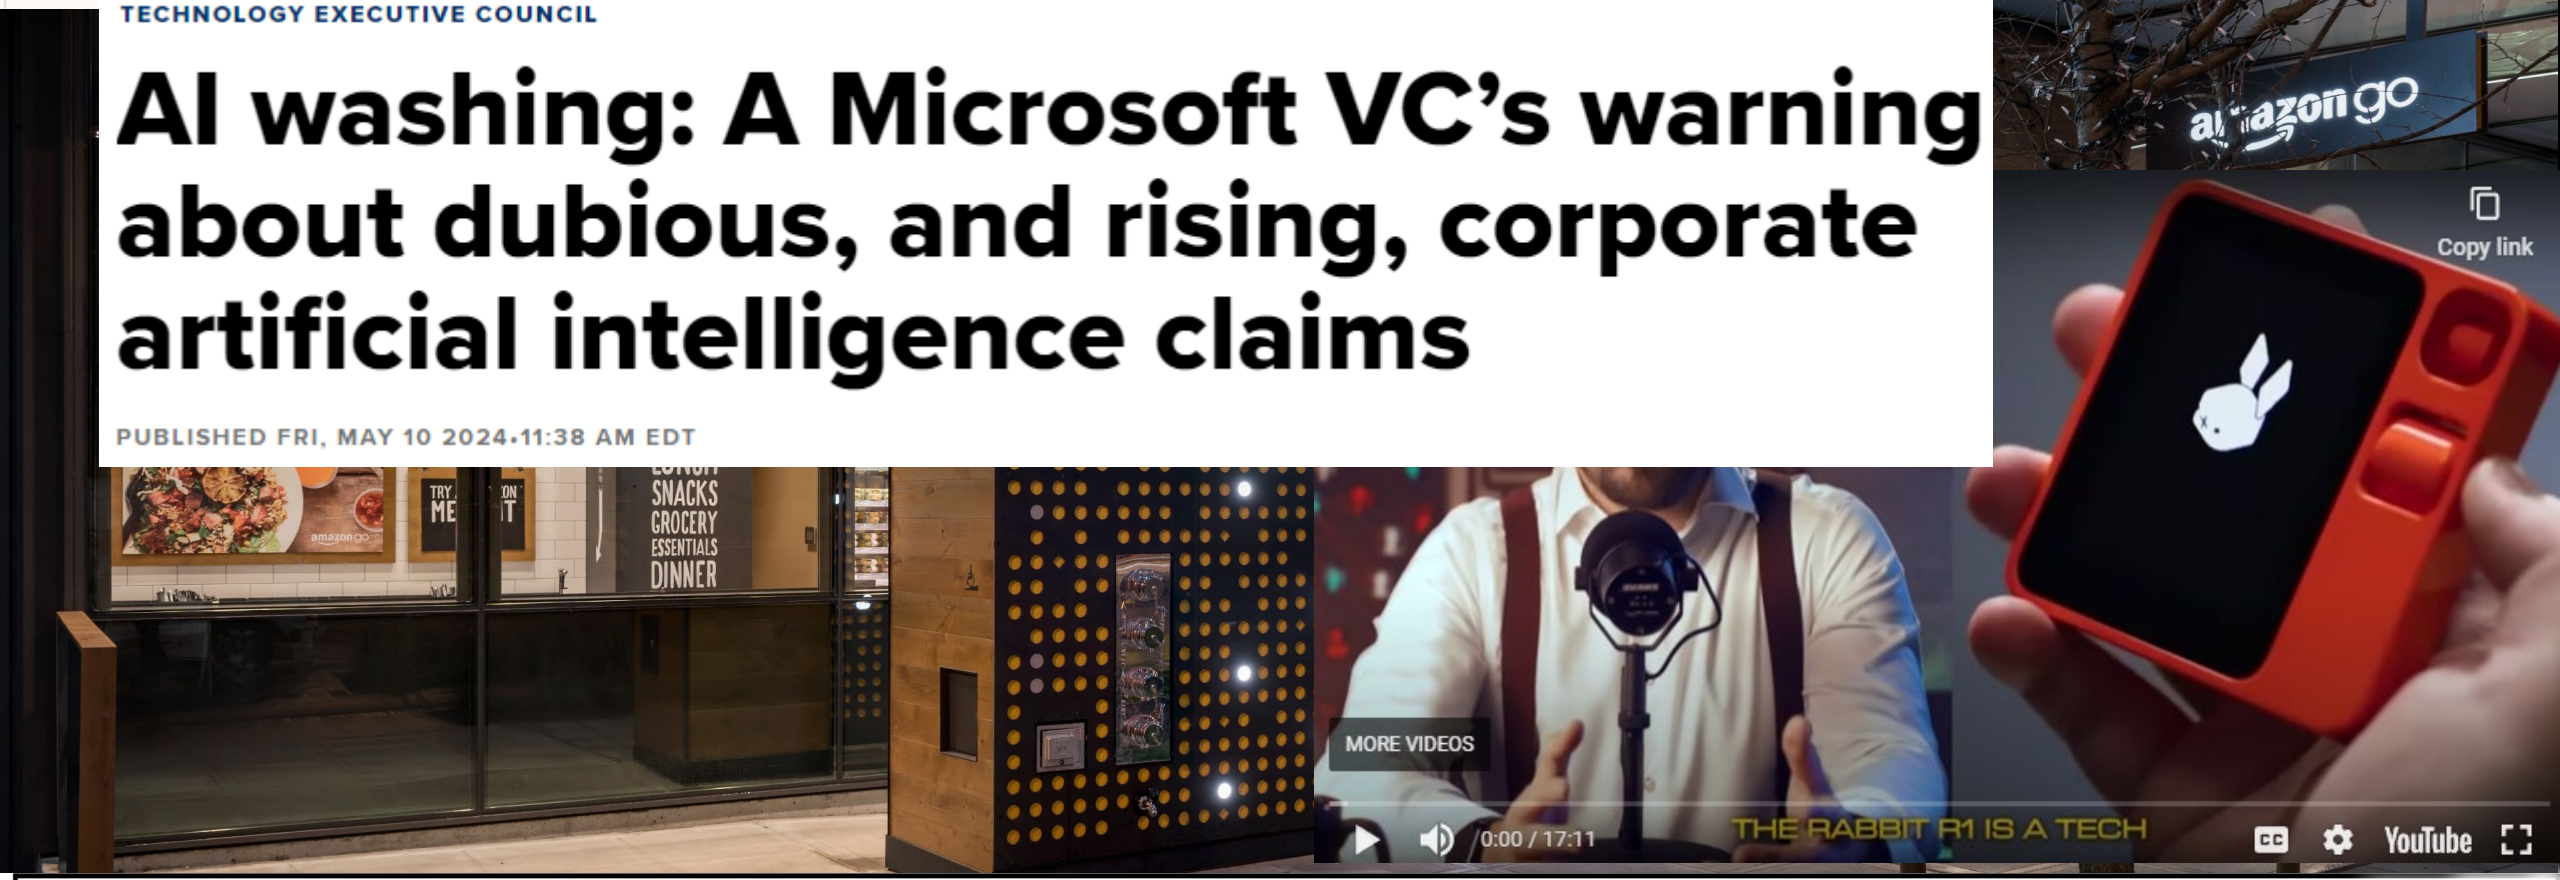

Three Mile Reopening

Microsoft

Former Microsoft CEO Bill Gates invested $1 billion in a nuclear power plant that broke ground in Kemmerer, Wyo., in June. The plant will power homes and AI, Gates told NPR's Steve Inskeep.

Peak of Inflated Expectations

Slope of Enlightenment

Hype Cycle

Technology Curve

Trough of Disillusionment

The Shadow AI Economy

Employees are Using LLMs (ChatGPT) More Than Company

Exhibit: the shadow AI economy, employee usage far outpaces official adoption

Why Don’t Employees Use Integrated Tools?

User Preference Drivers: Generic LLM Interface vs. Integrated Tool

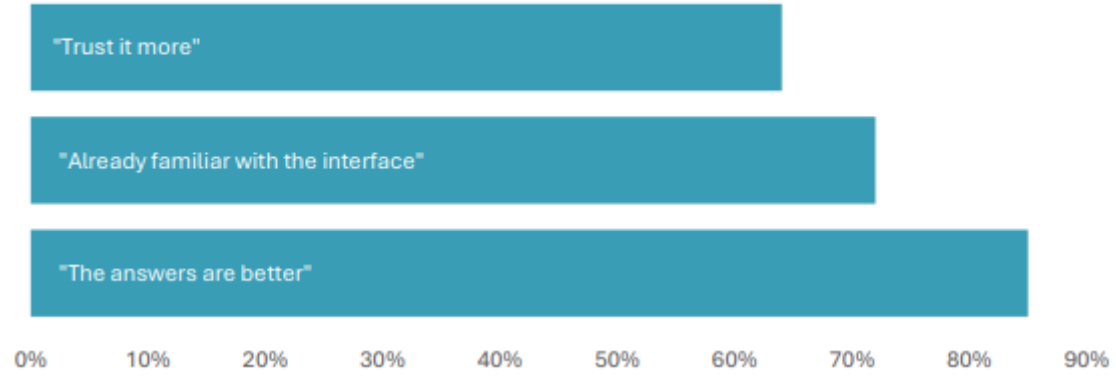

The Shadow AI Economy

So Why Don’t Companies Adopt LLMs Internally

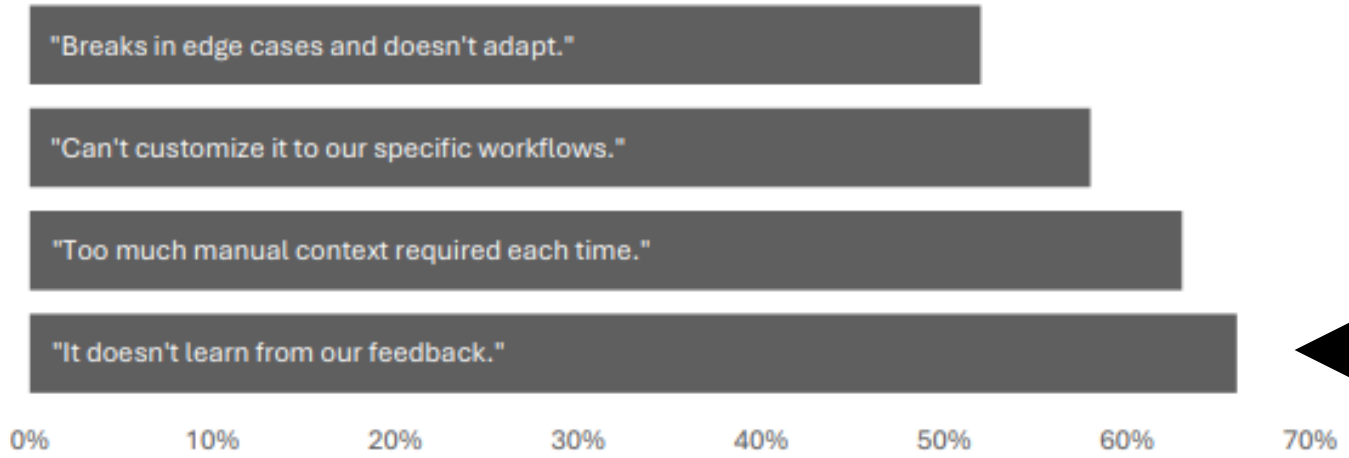

ChatGPT

AIs

?

What Do you Need to Build

Conditionals

1. If any of these words show up, it’s a fraud
2. If any of these words show up, it’s not a fraud
3. If the EBITDA is larger than x, it’s fraud
4. If the years experience is smaller than y, it’s not a fraud
5. If the claim is between these values, it’s fraud

Static Tools

Moving Forward – Building an AI

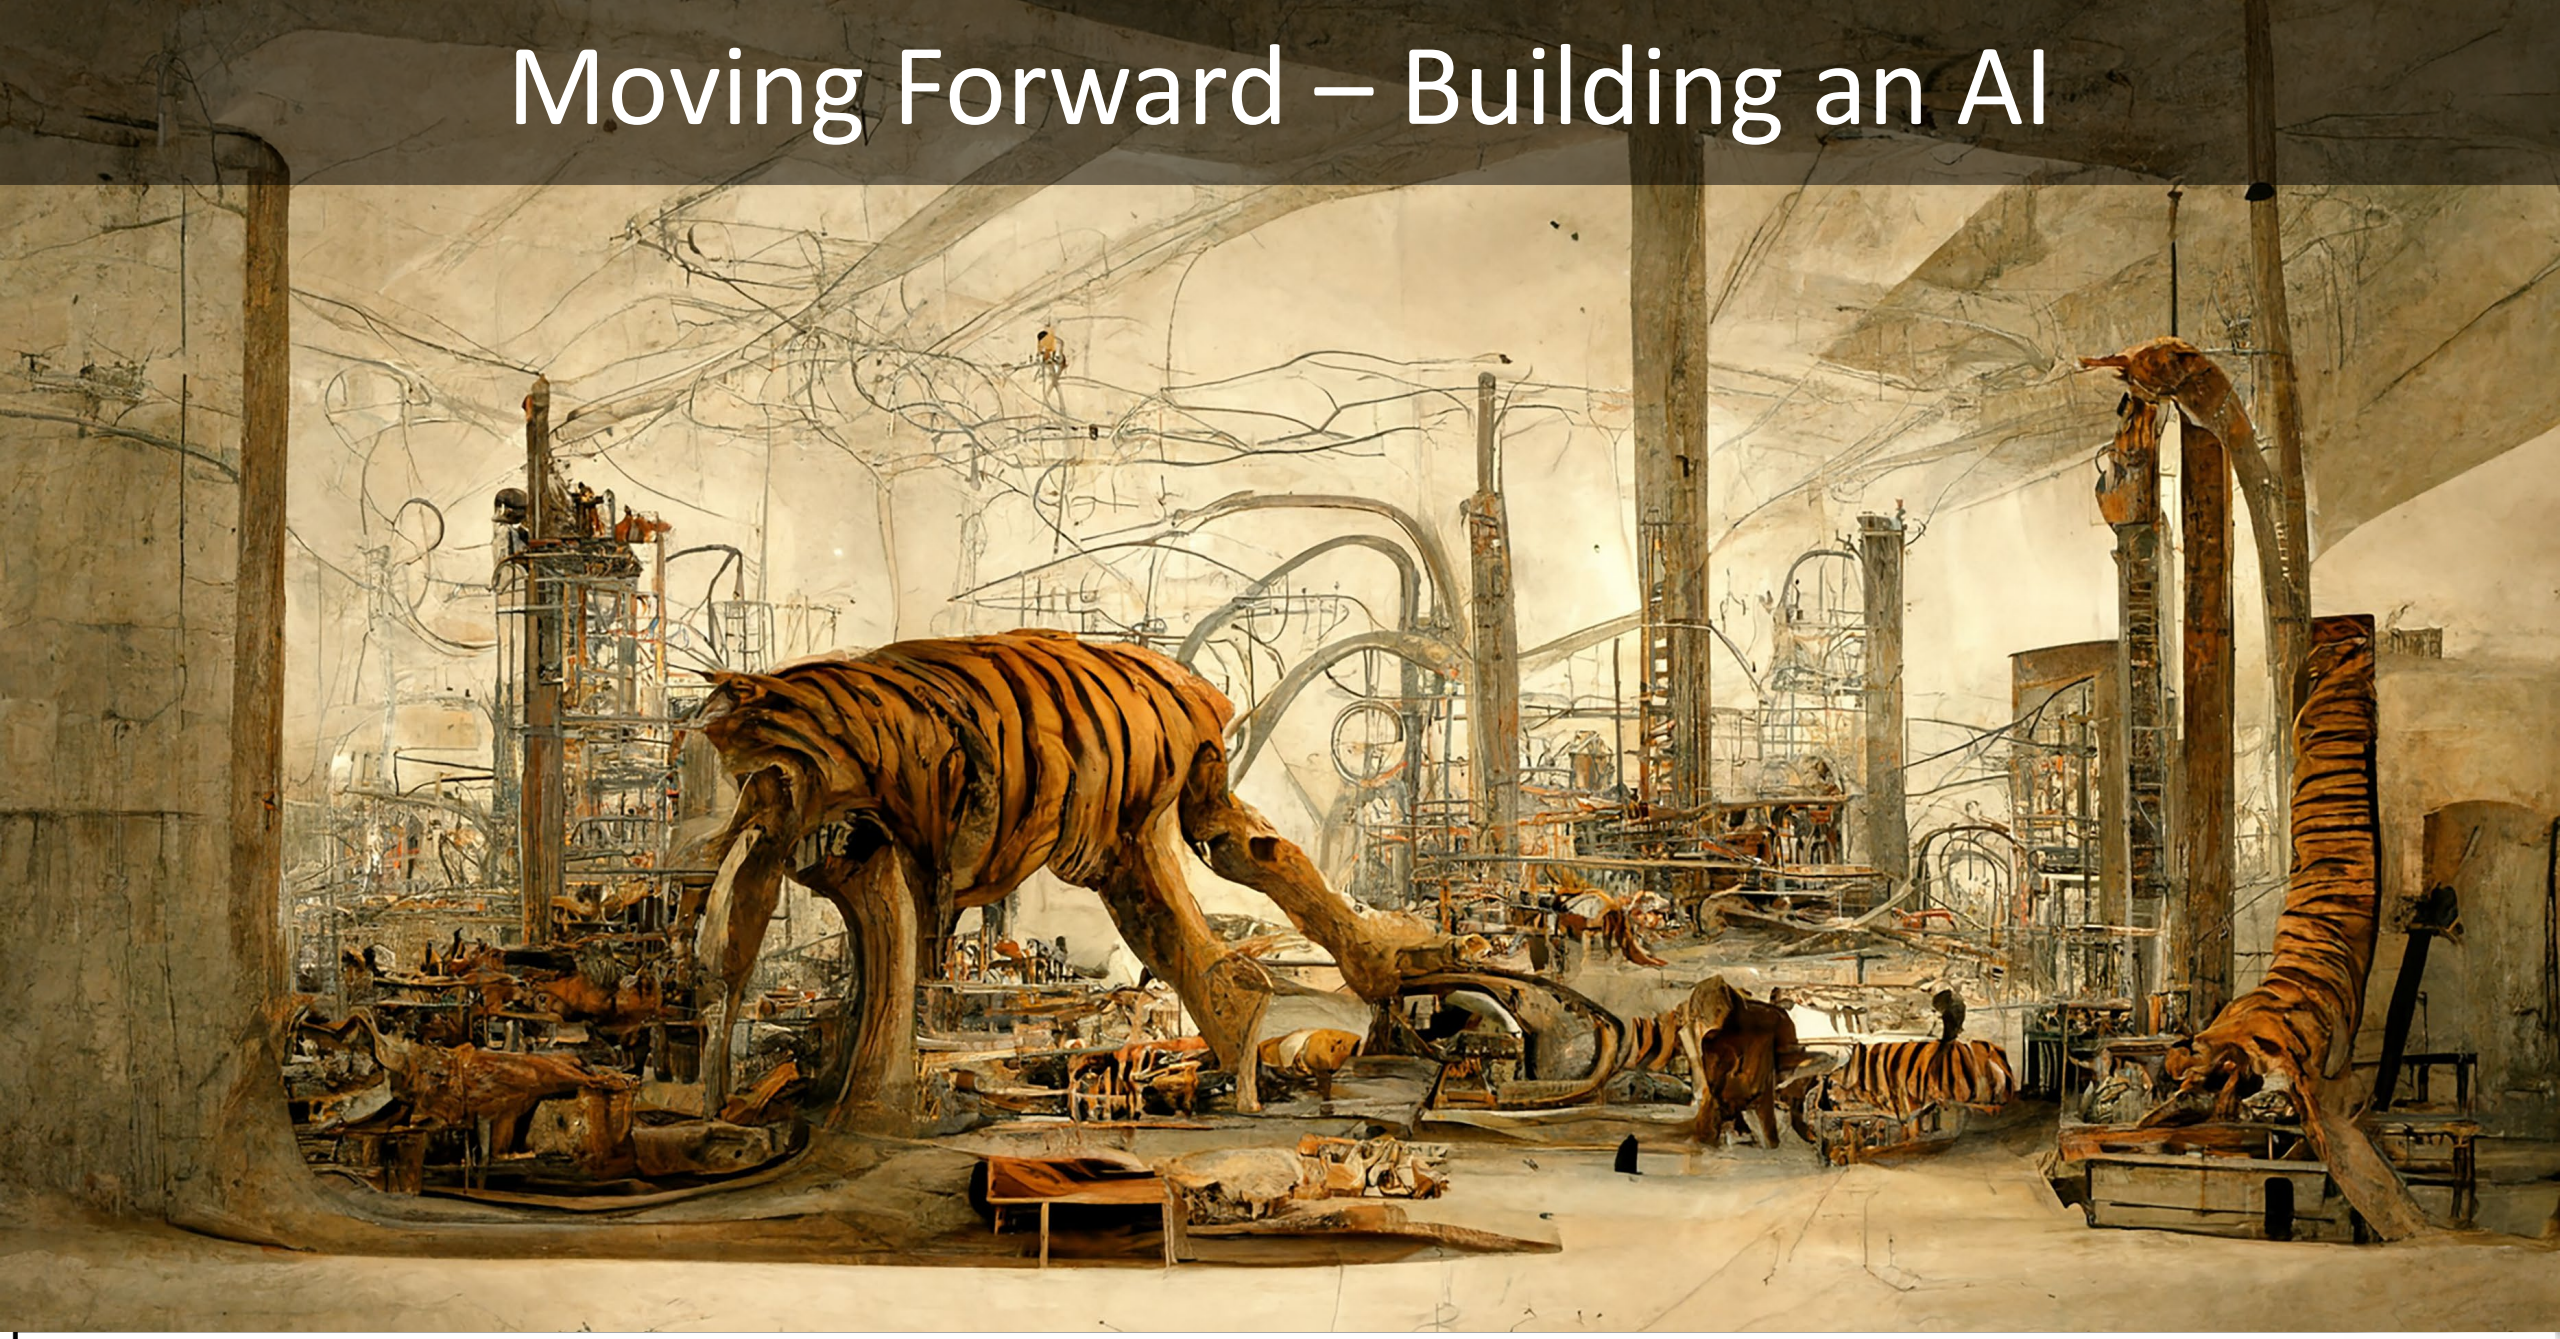

Use 'Buy, Acquire, Build Project Management Methodology

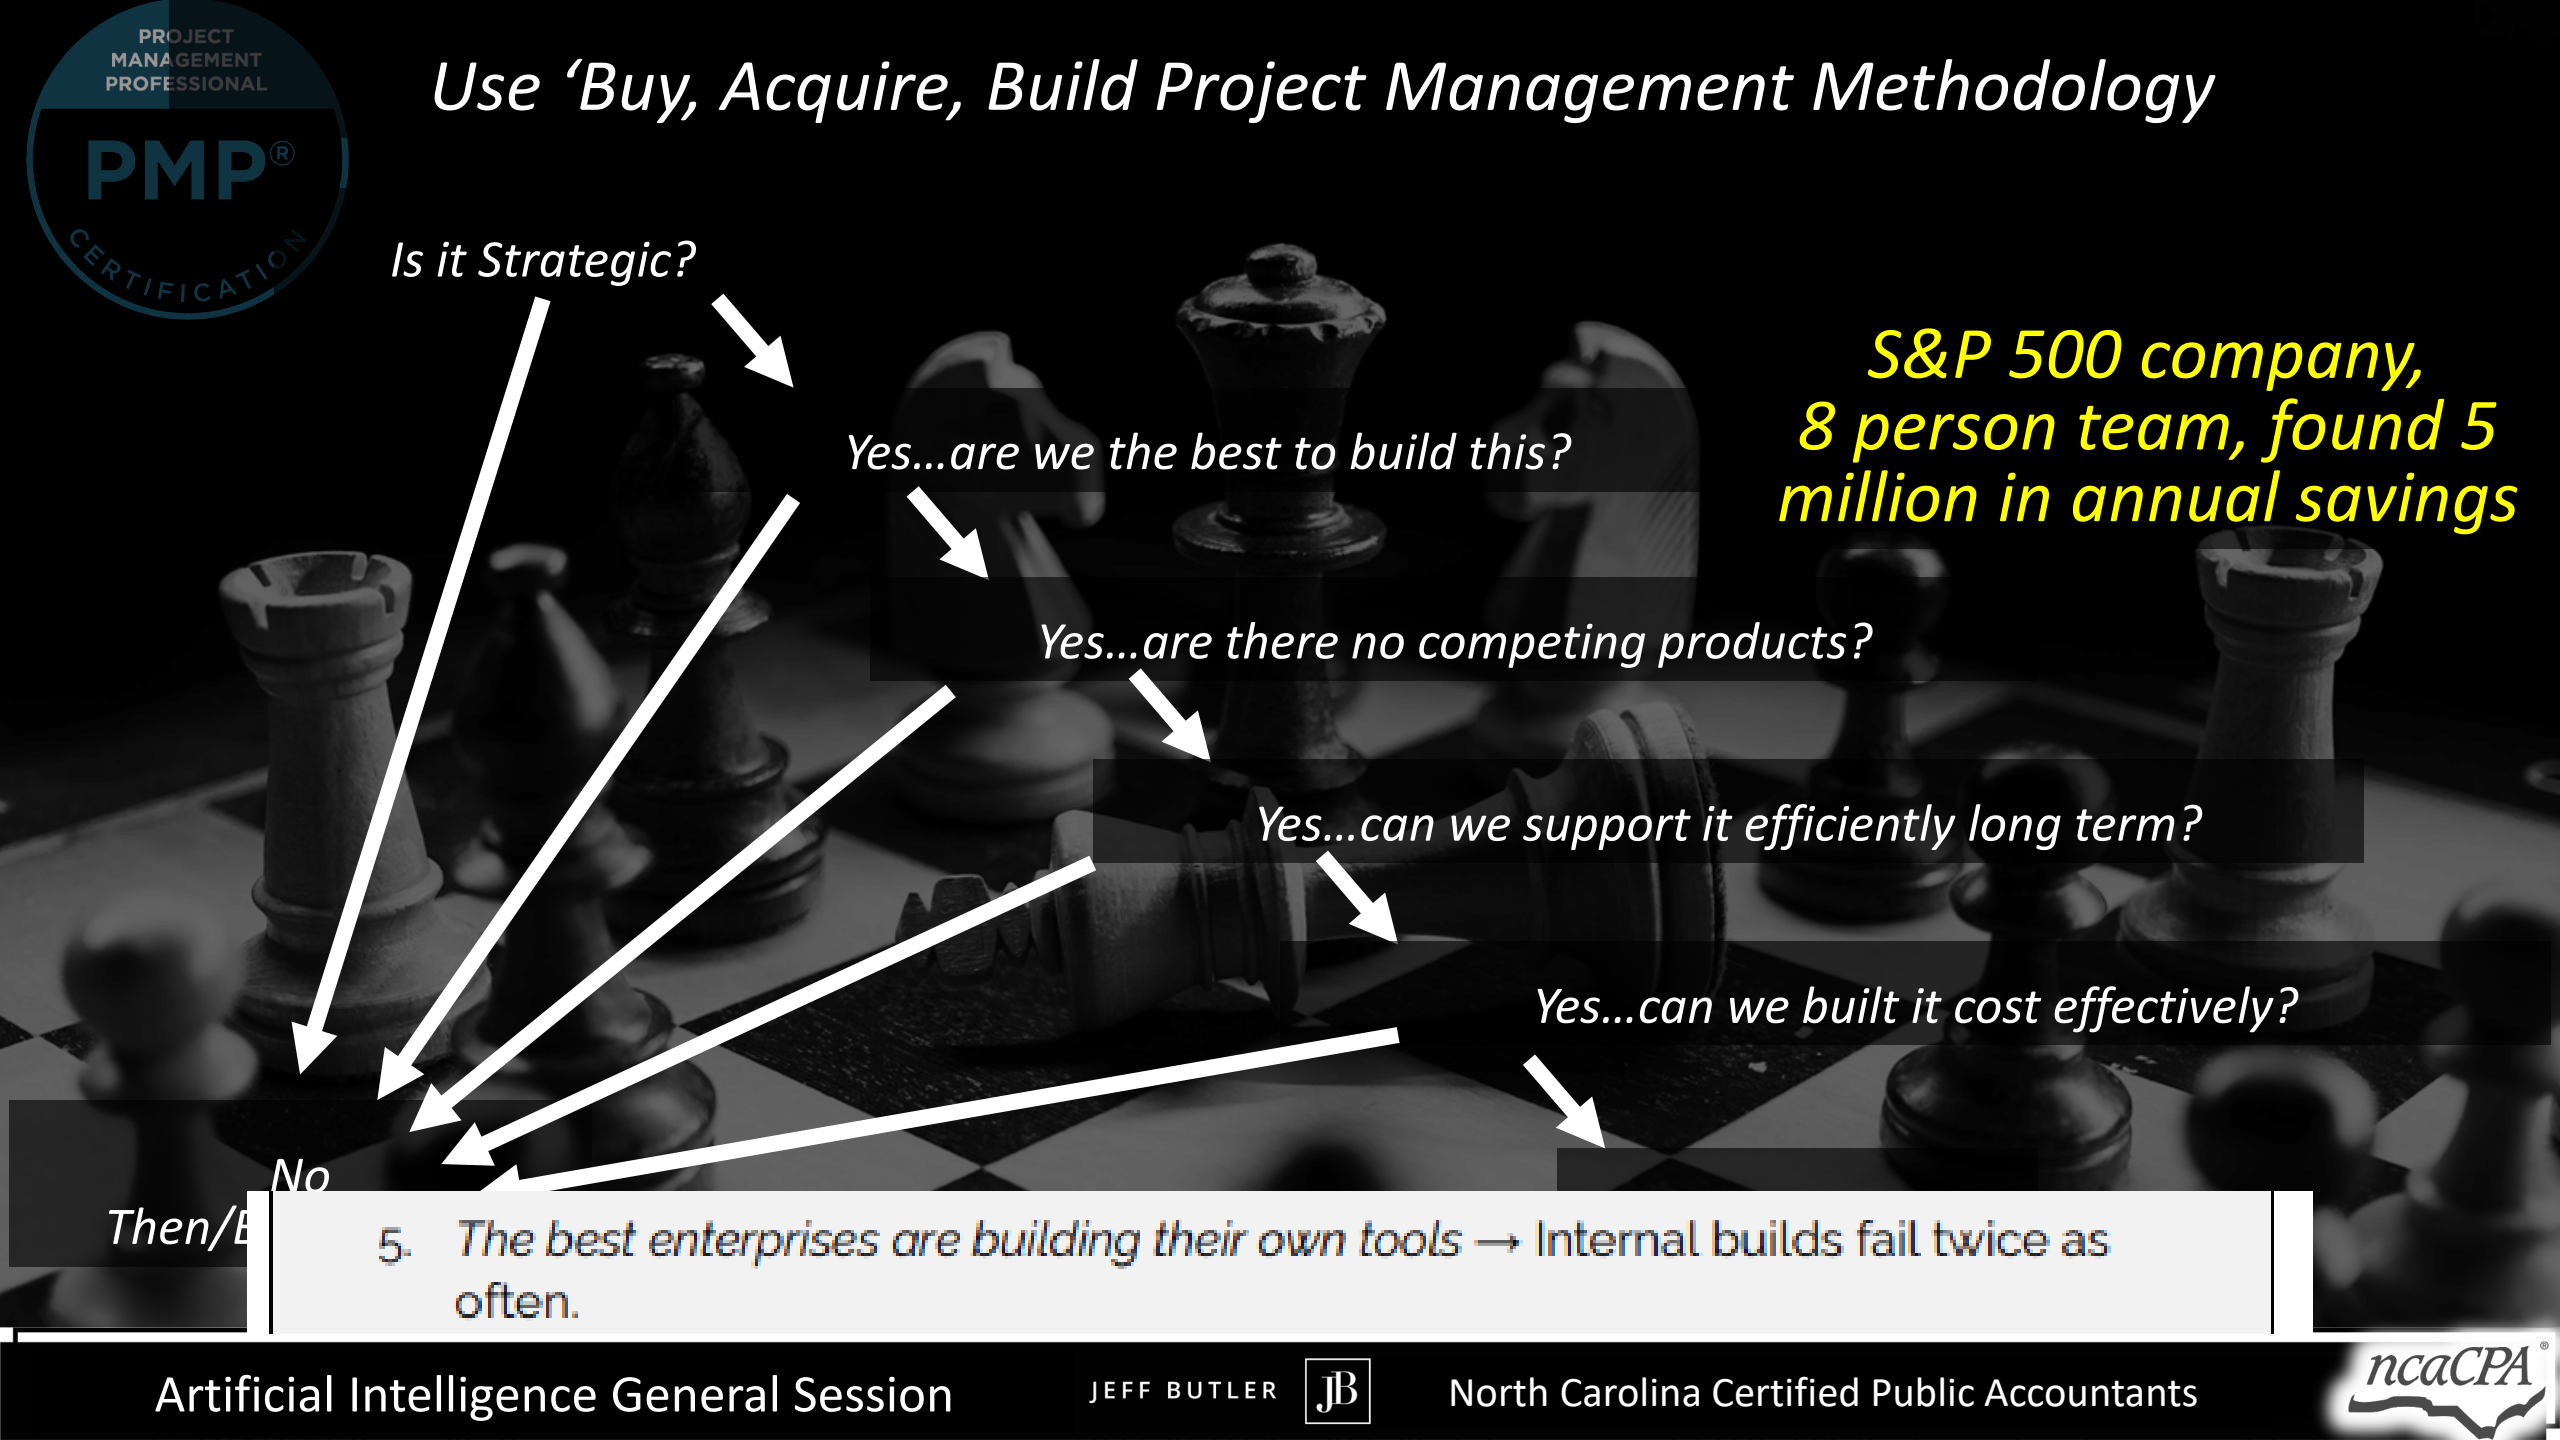

S&P 500 company, 8 person team, found 5 million in annual savings

5. *The best enterprises are building their own tools* → Internal builds fail twice as often.

Biggest Mistakes Firms Make With AI Implementation

1. Expecting end-to-end solution without customization
2. Tackling large problems instead of easy wins – scoping
3. Not communicating with vendor
4. Buying built-ready solutions

	Narrow Scope	**Broad Scope**
Simple Execution	Fast wins (Spend Categorization, Contract Review)	Partial pilots (Supplier Risk Monitoring)
Complex Execution	Early pilots (Negotiation Bots)	Fails (Full Procurement Orchestration)

If a Vendor Promises a Ready-Made Solution

Examples of Effective AI Scoping

Effective Scope

- Voice AI for call summarization and routing
- Document automation for contracts and forms
- Code generation for repetitive engineering tasks

Ineffective Scope

- AI Onboarding for new clients
- Customer service AI bot that can handle payments over the phone
- AI call agents to follow up on new inquiries

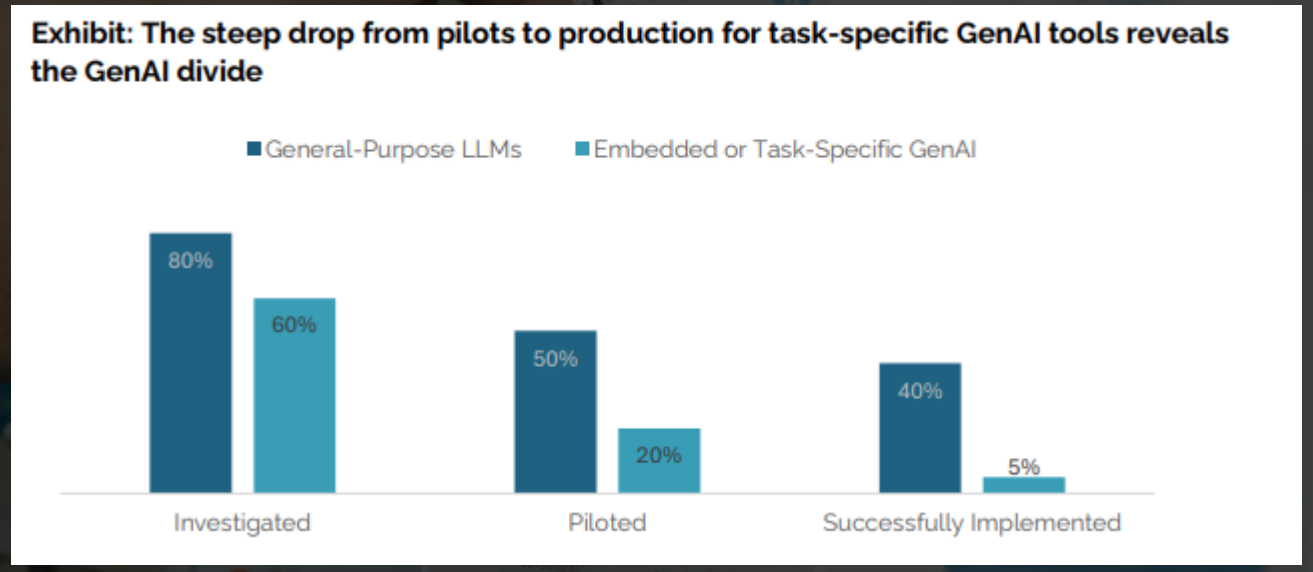

How To Navigate an AI Vendor Relationship

- AI Vendors should act like BPO clients, not SaaS solutions. (18-month commitment)
- Demand deep customization
- Drive adoption from the front lines
- Hold vendors accountable to business metrics.
- Most successful buyers understand the relationship is a partnership not just purchase.

Interviews, surveys, and analysis of 300 public implementations in “The State of AI Business in 2025” *

Helpful Resources and Slides

jeff@jeffjbutler.com

Helpful Resources and Slides:


FIRST PLACE
FINE ARTS
COLORADO STATE FAIR
PUEBLO, CO
EMERGING
Jason Allen
Pueblo West
Théâtre D'opéra Spatial
$750
Colorado State Fair

The Power of Human Connection

AI Disruption

Overcoming the Hype and Discovering the Opportunity

Helpful Resources and Slides