Understanding AI and the Future of Jobs

A Short Overview of the Institute for Experiential AI

EAI The Institute for Experiential AI
Northeastern University

The Roux Institute
Northeastern University

Feb 14, 2025

Usama Fayyad, Ph.D.
Inaugural Director, Institute for Experiential AI, Northeastern University
Professor of the Practice, Khoury College of Computer Sciences
u.fayyad@northeastern.edu

My Key Themes

- **Working AI is all machine learning (ML) today**
  - *Data, Data, Data*
- **Smarter not larger models**
  - *What can we do to find alternatives to the ever larger LLMs*
- **Talent and upskilling the workforce is a necessity**
  - *Not optional*
  - *Experiential, Experiential, Experiential*
- **Education in the traditional way will be disrupted**
  - *We need to evolve a new reserve of skills in applied AI*
  - *We can lead in the new way to teach and learn with AI*

Some Questions

- **How do we differentiate in a crowded environment?**
  - *Look for smarter not bigger*
- **In the AI space, what is it that we should invest in?**
  - *Understand the depreciating vs appreciating assets in the AI equation*
  - *Understanding practical ROI is key*
- **Talent and upskilling?**
  - *What works and how to make it work*
  - *Education will be disrupted*
- **How do we get AI adopted in the way we work?**
  - *Applications and pragmatic focus on learning by doing is key*
- **How do we leverage our strengths?**
  - *Specialize and go deep*

What is Artificial Intelligence?

The use of computers to "simulate" human intelligence

- Defining "intelligence" is an open problem
- "Common Sense Reasoning" still an open problem

The excessive hype lead to two AI Winters - Cut in funding, industry disillusionment, and practitioners avoid the field

- AI Winter 1 - Mid 1970's
- AI Winter 2 - Early 1990's

What about **Machine Learning**?

*A subset of AI concerned with machines modifying/learning behaviors based on experience (inputs) - **Training Data***

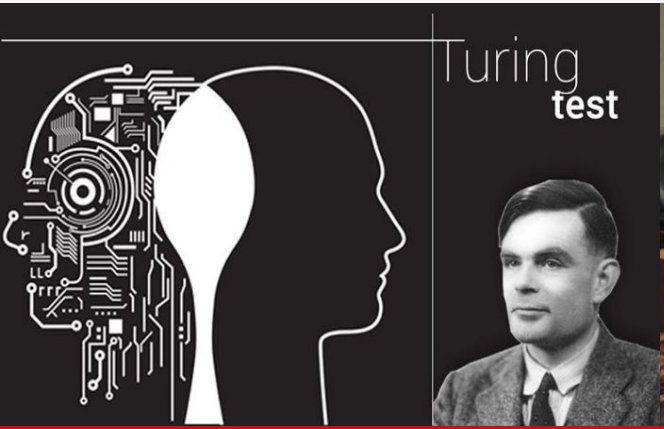

Hype Sounds Familiar?

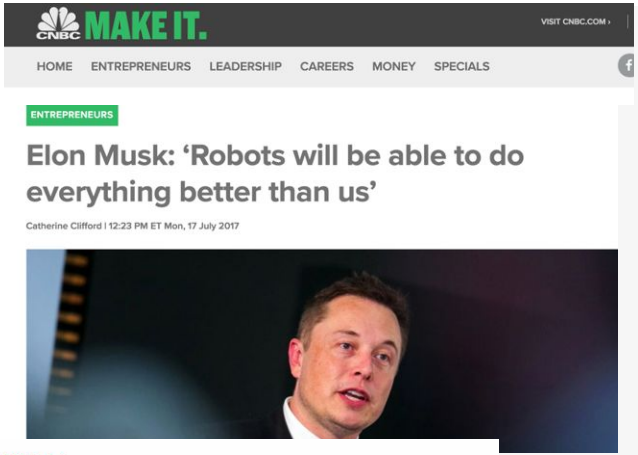

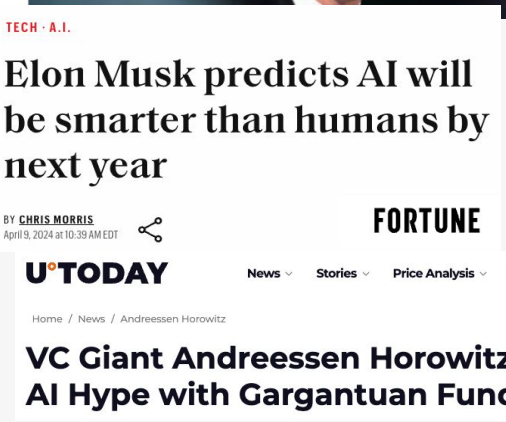

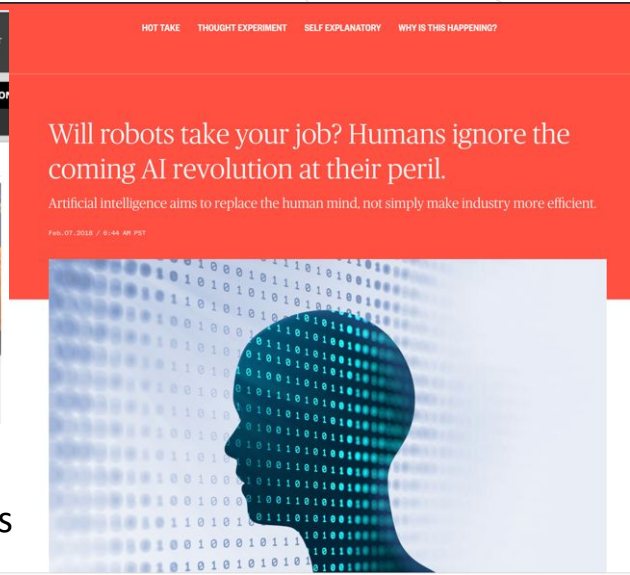

- Major hype in the 1980's – AI was going to solve all problems and change the world
  - U.S. was afraid of Japan AI program – 5th Gen. Systems

- We are all going to be useless
  - Jobless
  - Brainless
- China 2030 AI is the new Japanese 5th Gen

Machine Learning survived both AI winters

***Not** because we developed **new/better ML algorithms**…*

*But because we had a lot **more data***

How does GenAI fit within AI, Machine Learning?

Artificial Intelligence

Programs with the ability to simulate human intelligence

Machine Learning

Programs with the ability to learn without being explicitly programmed

Generative Models

Programs with the ability to learn how to generate new data that is similar to a given set of training data

Applied ML

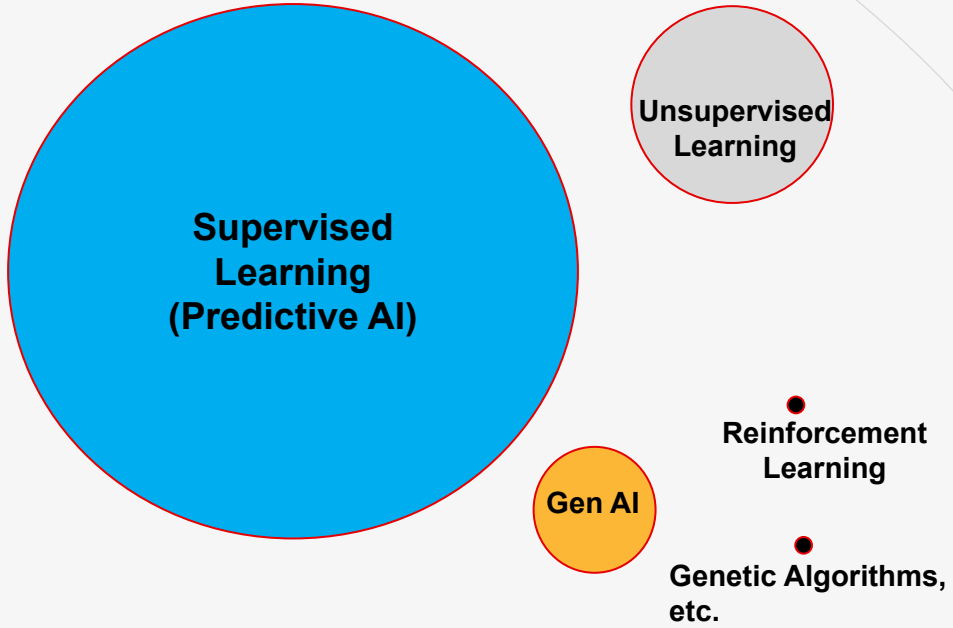

SECRET 1: Making AI Work (from talk on Secrets of Making AI Work)

1 **Narrow the Scope as Much as Possible**

Reduce the problem domain to one where “**complete knowledge**” is possible by narrowing scope

Opposite philosophy to ChatGPT

Complete knowledge is impossible unless your focus is extremely narrow

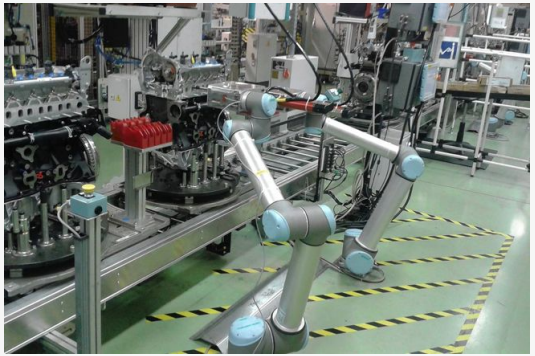

Amazon Go Stores

I spent 53 minutes in Amazon Go a the future of retail

By Matt McFarland, CNN Business

8 minute read · Updated 5:39 PM EDT, Wed October 3, 2018

Seattle (CNN Business) — If you want to glimpse the future of retail, check out a Amazon Go store.

They're sleek and modern, with a minimalist vibe. Black merchandise racks. Wo veneer. Polished concrete. Pop music plays softly in the background; cameras n the ceiling monitor your every move as you wander the aisles.

USA TODAY

SHOPPING

Amazon.com, Inc. Add Topic +

Why Amazon is ditching Just Walk Out checkouts at grocery stores

Betty Lin-Fisher
USA TODAY

Published 5:48 p.m. ET April 2, 2024 | Updated 11:11 a.m. ET April 9, 2024

Amazon is ditching its "Just Walk Out" technology – which allows customers to shop and leave the store without going to a register – for what it says is better technology at its Amazon Fresh stores.

The change, announced Tuesday, only affects Amazon Fresh locations, the Seattle-based company's grocery stores, and not Amazon Go, which are smaller convenience stores. It also does not impact the more than 130 third-party retailers that Amazon partners with for use of its "Just Walk Out" technology at such locations as airports, college stores and cafes, an Amazon spokesperson confirmed to USA TODAY.

The artificial intelligence technology, which sends customers their receipts after they've taken items off the shelves and left the store, will be replaced by smart carts, which allow customers to scan their items as they shop and see what they're paying and saving on a screen, Amazon said.

In an email, Amazon said it made the decision to cut the technology, which can be found in Amazon Fresh and Amazon Go stores, due to customer feedback.

Data: Some of the Big Challenges for AI

Successful AI is totally dependent on ML/Data Science, hence need good training data: Data remains a huge challenge for most organizations

Good training data is **extremely expensive** to get

→ *… reliable labelling **even more expensive***

Just **collecting and managing** raw data is a **challenge** for most organizations

→ *… data is growing exponentially with digitization, cloud, and IOT*

Data manipulation is very difficult, few understand unstructured data

Must Capture Data at High Granularity

But most businesses are not equipped to effectively manage data as an asset

How do we make this Data work for the business?

New economy of Interactions is rich with unstructured data

in fact, 90% of Data in any organization is UNSTRUCTURED

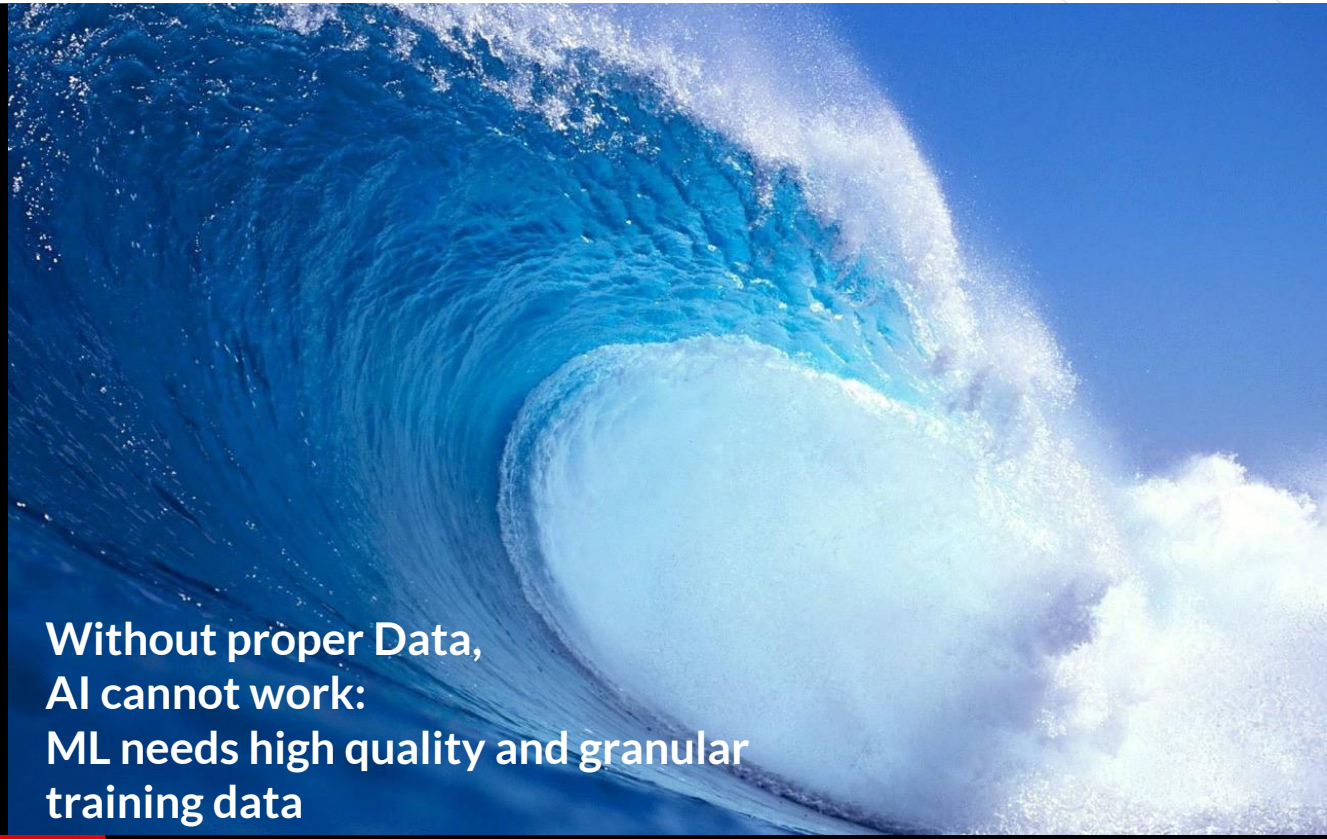

The Lost Themes: Data and the Human-in-the-Loop

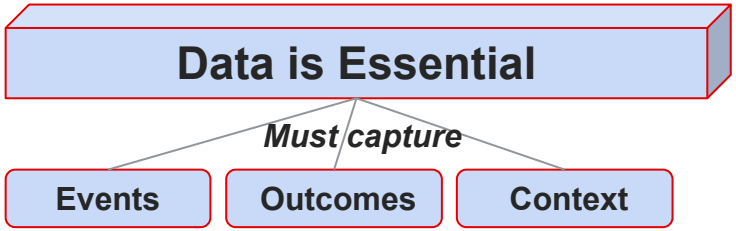

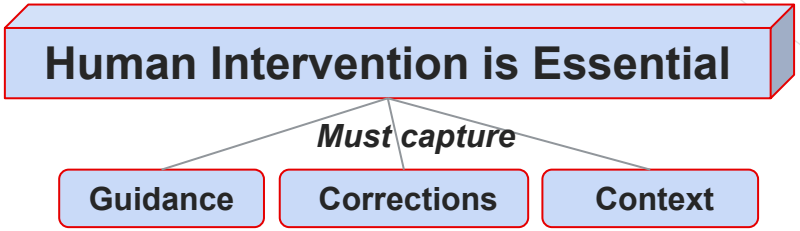

most organizations struggle with the basics of making data work as an asset

The **"AI-haves" understand this** and they have systems to:

- Capture every bit of data + context
- The ability to leverage this data through Machine Learning (ML) to automate the determination of the **right action** in the **proper context**

Capitalize on and capture every human intervention to guide AI

What about the human-in-the-loop?

Human intervention is the most valuable asset for Google, Open AI, Amazon, Tesla, and all companies that make AI work

What can we do about Data?

Capitalize on Data as the Essential and Lasting Asset…

Enable the capture, management, and retrieval of data that captures operations and interventions…

Some actions

a. AI Data Strategy: data capture, storage, and management
b. Enable the capture data from operations systematically
c. Data rights and access
d. Safe guidelines for how to collect and share data responsibly…

Depreciating assets: Hardware, Software/algorithms…

Data is the "investment grade" asset.

Example: Areas where AI Can Help in Healthcare and Life Sciences

- Digital Health (Home Hospital, Digital Pathology)
- Detect signals in multi-sensor environments
- Trach health issues and aging issues in patients outside clinic
- NLP to leverage unstructured text data – LLM (large language models) and other open source methods for image and TS analysis
- Image analysis tools to leverage and retrieve related image data (*query by example, pattern recognition, etc.*)
- Graph-based and network representations
- Network Science models for understanding multi-factor interactions
- Multi-omics approaches to extend the single-omics traditions

What is Experiential AI?

Human-centric approaches to solve real problems in real contexts with a human in the loop: Effective Human ⇔ AI cooperation.

Human intervention is a must

Human intervention is a great opportunity for knowledge capture & ML

Thesis: *Taking an applied approach is the best way to solve problems in science and in practice:*

- *Leverage data in a way that amplifies the values and benefits of machine learning*
- *create mechanisms for machines and humans to learn together*

Result: *creating actions, decisions, & results that neither machine nor humans can achieve alone.*

Generative AI

Now What?

What is Generative AI?

Generation of text, images, etc. from a textual or multimodal prompt

- Generating text is based on Large Language Models (LLMs)
- Other models generate images or videos

These models are referred to as
"Stochastic Parrots"
because they do not understand what they say

Examples of Generative AI

Generative AI for Text:
- ChatGPT (Open AI)
- Bard (Google)
- LlaMa (Meta)

Generative AI for Images:
- DALL-E (Open AI)
- Stable Diffusion (Stability AI)
- Midjourney (idem)

Generative AI for Short Videos:
- Midjourney (with a parameter)
- Runway

Generative AI for Code:
- Open AI's Codex
- Microsoft CoPilot (Codex on Github)
- Cognition's Devin (s/w engineer)

Generative AI for Voice & Music
- Aiva.ai
- Endel
- Lyricstudio.net

Plus other objects:
- DNA, Protein folding
- Chemistry, Physics, Molecular Design

What is ChatGPT?

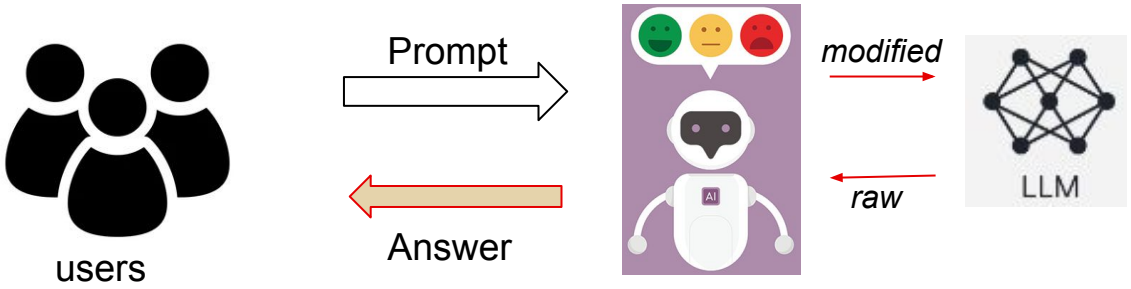

Chat interface running atop GPT-3 (Then GPT-3.5, later GPT4)

- **Generative Pre-trained Transformer** model – utilizes transformer
  - GPT3 trained on a corpus of about 1 TB of Web text data (with adjustments to reduces biases)
  - GPT-3 is a neural net with 175 billion parameters – expensive to train

- ❑ Uses generative unsupervised training
- ❑ Works by predicting next token in a sequence – sequences can be very long…

What is ChatGPT?

Most capabilities of ChatGPT not new – around since 2020 in GPT-3

Auto-complete on steroids – with prediction and “concept graph” & human feedback/editing

How much does it cost to train GPT-3 on the 1 TB of Curated Data?

How much did it cost to curate the right 1 TB of data?

Speed of Adoption

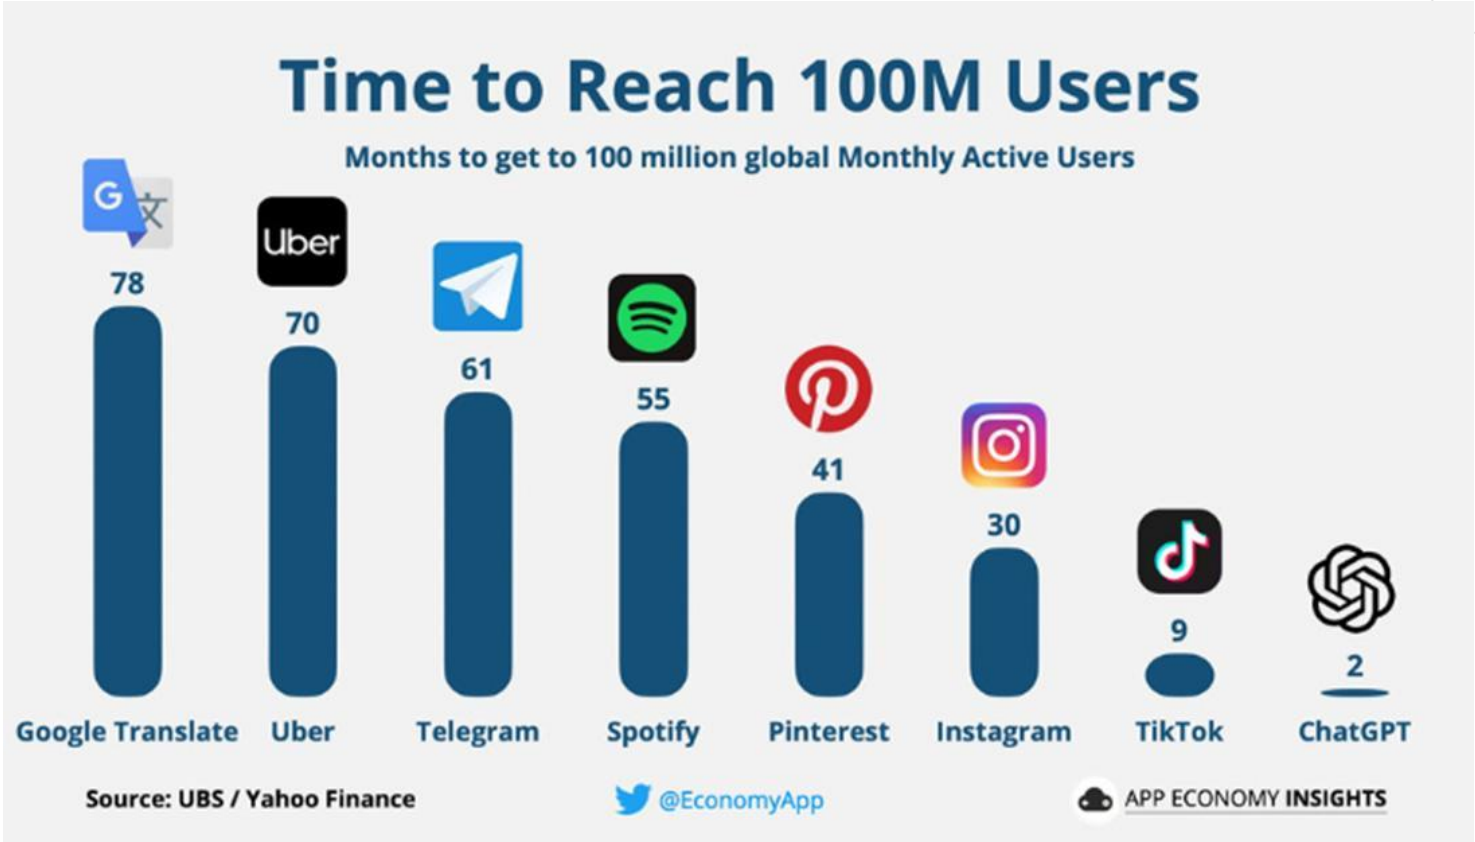

Economic Impact

- Knowledge worker tasks
  - Several estimates, ranging from 15% to 80% of the work likely to experience significant acceleration
  - But total automation not in reach

What is the size of the "knowledge economy"?

Between 19.6% and 30.4% of **global employment** (ILO, 2023)

Percent of Knowledge Economy	
High Income Countries	35-54%
Upper-middle Income Countries	22-54%

Source: U.N. Report: "Automation hits the knowledge worker: ChatGPT and the future of work"
https://sdgs.un.org/sites/default/files/2023-05/B59%20-%20Berg%20-%20Automation%20hits%20the%20knowledge%20worker%20ChatGPT%20and%20the%20future%20of%20work.pdf

Jobs Impact

- Human in the Loop is ***Essential***
  - Need to check the output
  - Need to modify and edit
  - Need to approve

Will AI replace my job?

NO - but a Human using AI will
… if you are not using AI

Many Cross-sector Uses of Generative AI

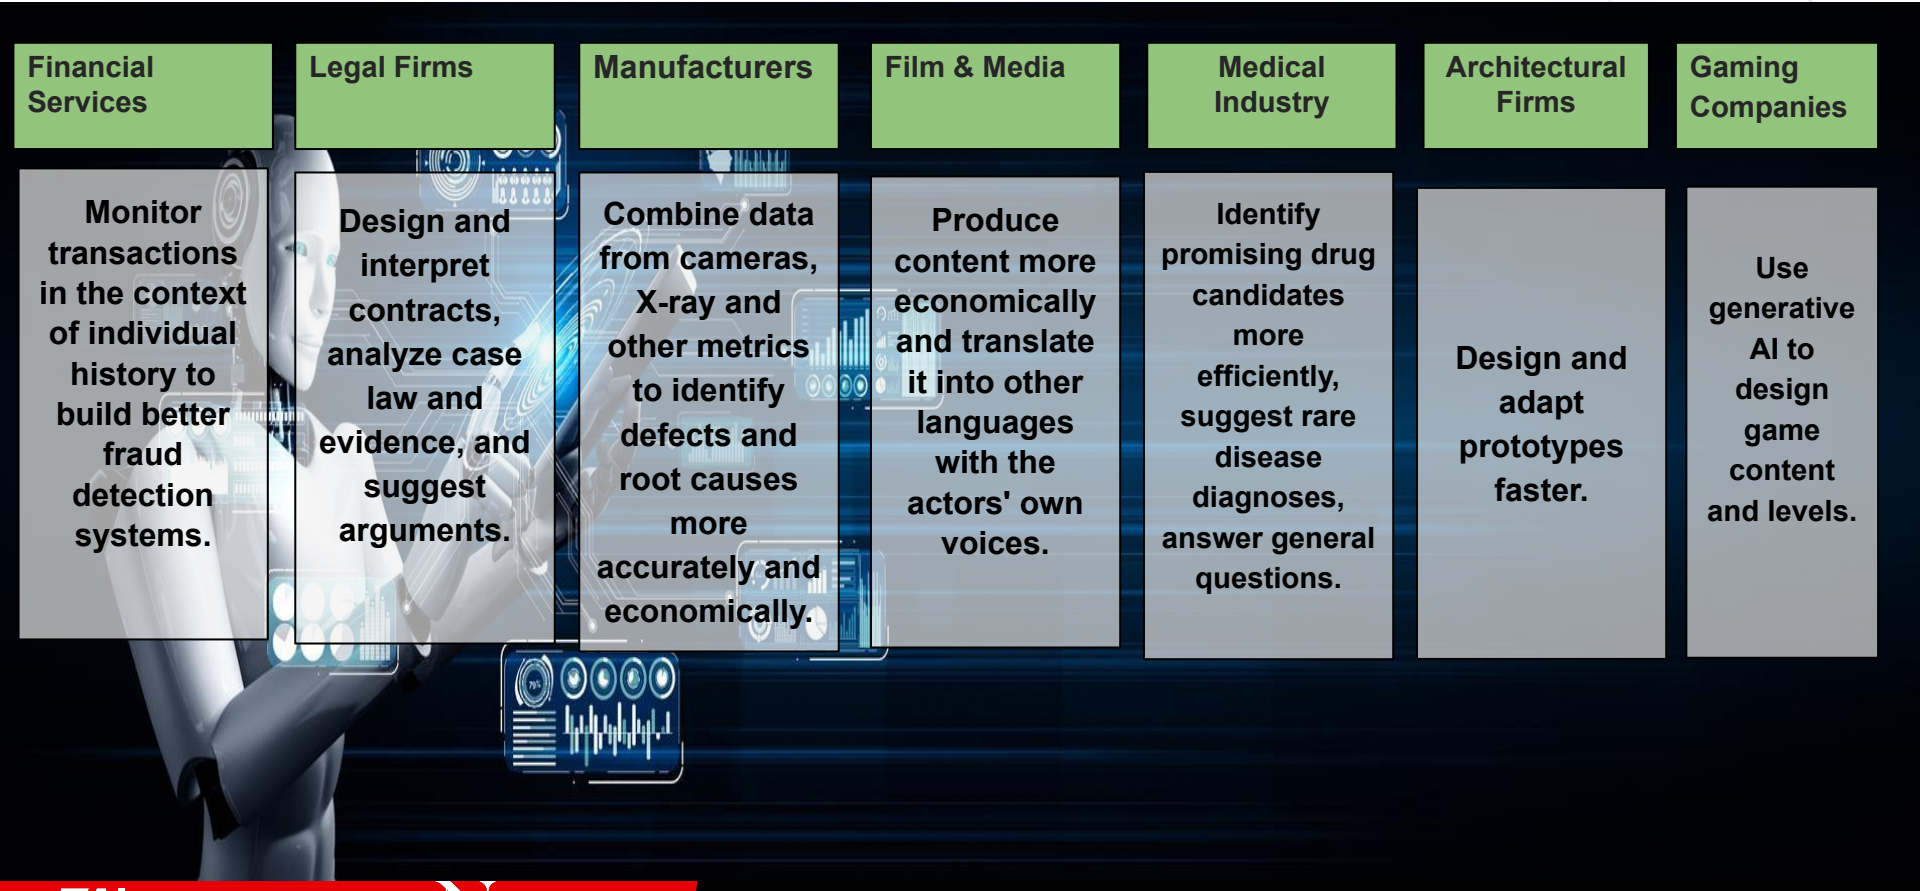

Financial Services	Legal Firms	Manufacturers	Film & Media	Medical Industry	Architectural Firms	Gaming Companies
Monitor transactions in the context of individual history to build better fraud detection systems.	Design and interpret contracts, analyze case law and evidence, and suggest arguments.	Combine data from cameras, X-ray and other metrics to identify defects and root causes more accurately and economically.	Produce content more economically and translate it into other languages with the actors' own voices.	Identify promising drug candidates more efficiently, suggest rare disease diagnoses, answer general questions.	Design and adapt prototypes faster.	Use generative AI to design game content and levels.

Major Developments over the Last Year (1)

From fascination to pragmatics and evaluation

Demos of amusing and fascinating capabilities

→ **How does it help the business?**

Loose claims of impact, acceleration and ROI

→ **Can you quantify the benefits?**

Microsoft Co-Pilot for Code

Is learning code an easier LLM task than Natural Language?

- **Microsoft acquires GitHub in 2018 for $7.5B** - Gets access to code by millions of developers using the open source platform
- Train LLM's to produce code (e.g. Python) instead of language (e.g. English)
  - ❑ Use documentation and notes
  - ❑ Label programs by purpose, and other metadata

Is Code easier than English?

Does this replace programmers?

How much would it increase the productivity of a programmer?

Helpful to non-programmers (no-code applications)?

Are larger models better?

- **Why are these LLMs getting bigger?**
  - GPT3 (175B) → GPT 3.5 (~ 0.5T) → GPT 4 (~ 1.5T)
- **Why is the rhetoric “bigger is better”?**
- **Why are we not leveraging prior knowledge (except in RAG hacks)?**

Sample Case Study in Marketing

GenAI with Customers

From GenAI to JenAI?

Same Tech for Deep Fakes

An Interesting tool for marketing?

Consider *Virgin Cruises* Marketing Campaign…

From GenAI to JenAI?

Same Tech for Deep Fakes

An Interesting tool for marketing?

Consider *Virgin Cruises* Marketing Campaign…

is JenAI really using Generative AI?

Personalization at scale:

- **Using generative voice:** match user speech to JLo's speech pattern and style based on training.
- **Using generative video:** match movement, pose, gestures to JLo's avatar

users can craft custom video invites from JLo herself…

- Implemented by a startup: DeepLocal - *before the chatGPT hype*

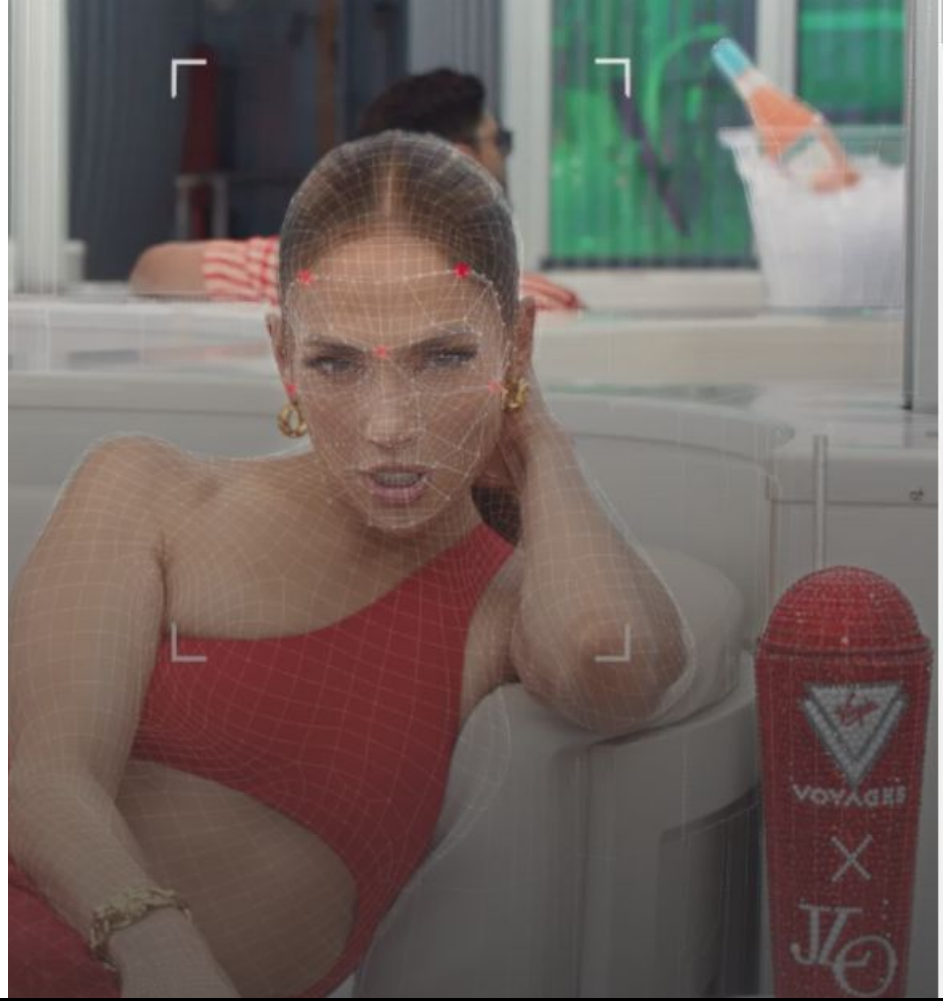

Amusing Applications?

Fiction is Easy to Recognize

But reality gets much more complicated (e.g., fake news)

Sophisticated mashup or plagiarism?

But Reality Gets Much More Complicated

+ Videos!

Gen AI enables Rapid Bad Uses

- Fraud
- Impersonation
- Social Engineering
- Cybercrime

Topics
European Parliament

EU AI Act

- June 2024

- Human-likeness
  (without transparency)

Responsible AI for LLMs

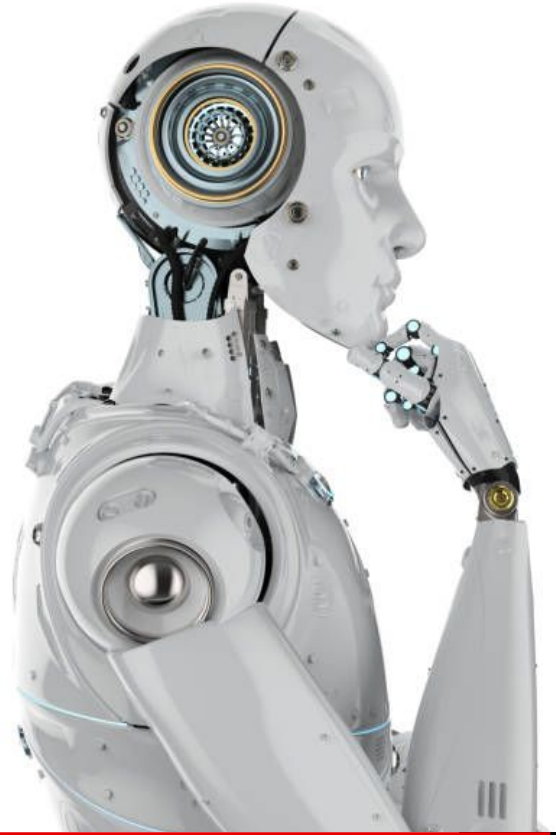

Stay lucid about what LLMs can and cannot do.

- LLMs do not "hallucinate" – they make errors
- LLMs are not "thinking" – they mimic conversation
- LLMs do not "have" opinions or character – they exhibit / reflect those
- LLMs do not "intend" outcomes – they do produce outcomes
- LLMs do not intend manipulation and harm but they do "cause" manipulation, misinformation, and harm

WHY Responsible AI?

✗ reputational cost	MARKET	✔ customer trust ✔ competitive edge
✗ regulatory issues	LAW	✔ shaping policy
✗ systems designed against us	SOCIETY	✔ systems designed for us

Responsible AI as a Field of Study

- How to practice AI responsibly?
- How to Assess the Risks of Using AI?
- An opportunity to lead
- A strategic advantage
- One of the major areas of the Institute for Experiential AI at Northeastern

Major Developments over the Last Year (2)

Robustness and mitigating errors

RAG (Retrieval Augmented Generation)

Advanced RAG - pre- and post-retrieval processing

Modular RAG - Hybrid Search, Recursive Retrieval and Querying, Step-back approach, Sub-queries, Hypothetical Document Embeddings

Graph RAG - graph/hierarchy leverage in prompting

Major Developments over the Last Year (3)

Larger LLM's are not necessarily better

Smaller LLMs (SLM)

Better Together (SLM + Knowledge Graphs)

Private LLMs on the rise - Narrow specialization of SLMs

April 2024: Large Hyperscalers discover SLMs - In The News:

After pushing the thesis: “the larger the LLM, the better” - tech giants started admitting SLMs (Small Language Models) since LLMs created unsustainable costs

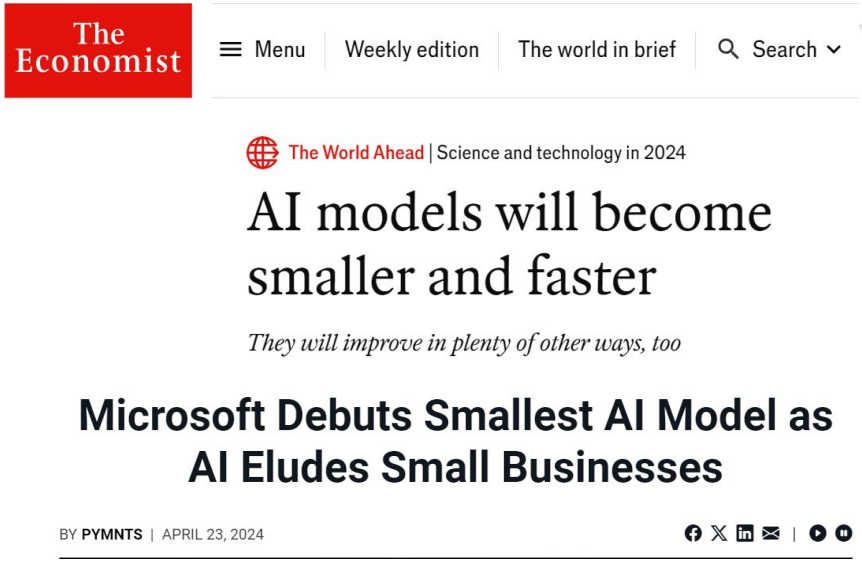
The Economist

Menu | Weekly edition | The world in brief | Search

The World Ahead | Science and technology in 2024

AI models will become smaller and faster

They will improve in plenty of other ways, too

Microsoft Debuts Smallest AI Model as AI Eludes Small Businesses

BY PYMNTS | APRIL 23, 2024

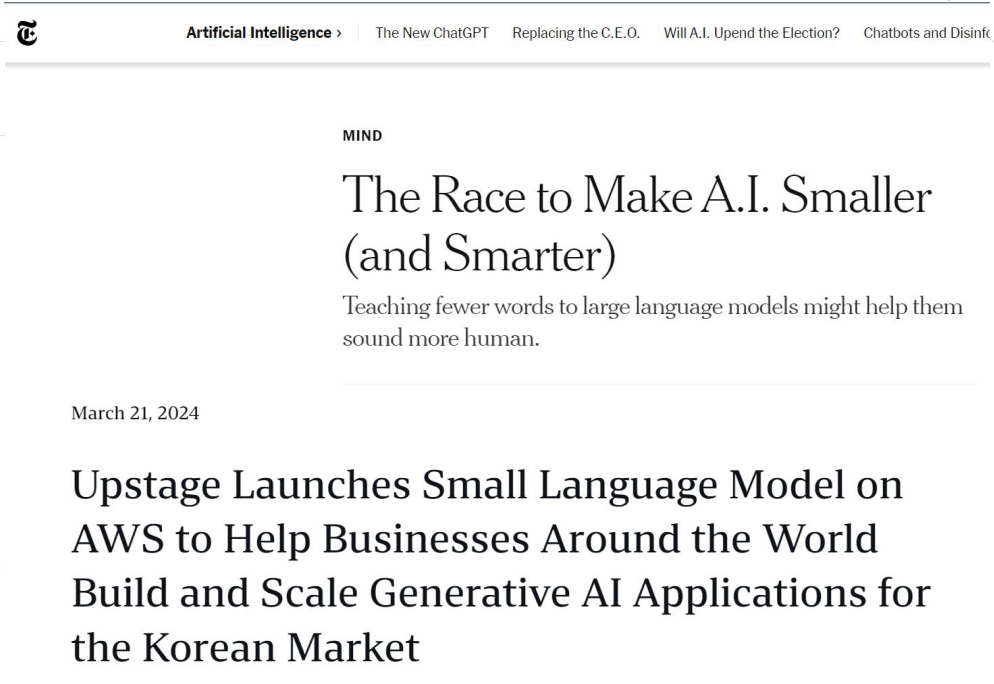
Artificial Intelligence › | The New ChatGPT | Replacing the C.E.O. | Will A.I. Upend the Election? | Chatbots and Disinfo

MIND

The Race to Make A.I. Smaller (and Smarter)

Teaching fewer words to large language models might help them sound more human.

March 21, 2024

Upstage Launches Small Language Model on AWS to Help Businesses Around the World Build and Scale Generative AI Applications for the Korean Market

Takeaway Lessons

Ask not: “What is the largest model we can afford?”

Ask: “What is the smallest LLM we can get away with?”

What is our angle to do a lot more with less?

Answers from our experience…

a. Prior knowledge: knowledge graphs, network models, other “synthesized” knowledge bases

b. Enable compute and data access to run fast launch projects

Major Developments over the Last Year (4)

Ease of Use and Integration with Apps

- Deployment Simplification - e.g. Amazon Bedrock
- Multi-modal interfaces: GPT4o
- Co-pilot like integration: Office, browser, Windows, Search
- A great example of integrated GenAI tools?

More Technical Advances

- LM Agents
- Student-teacher LLM models
- Synthetic data
- Research to understand what happens inside the “black box”

A Great Example of Using GenAI

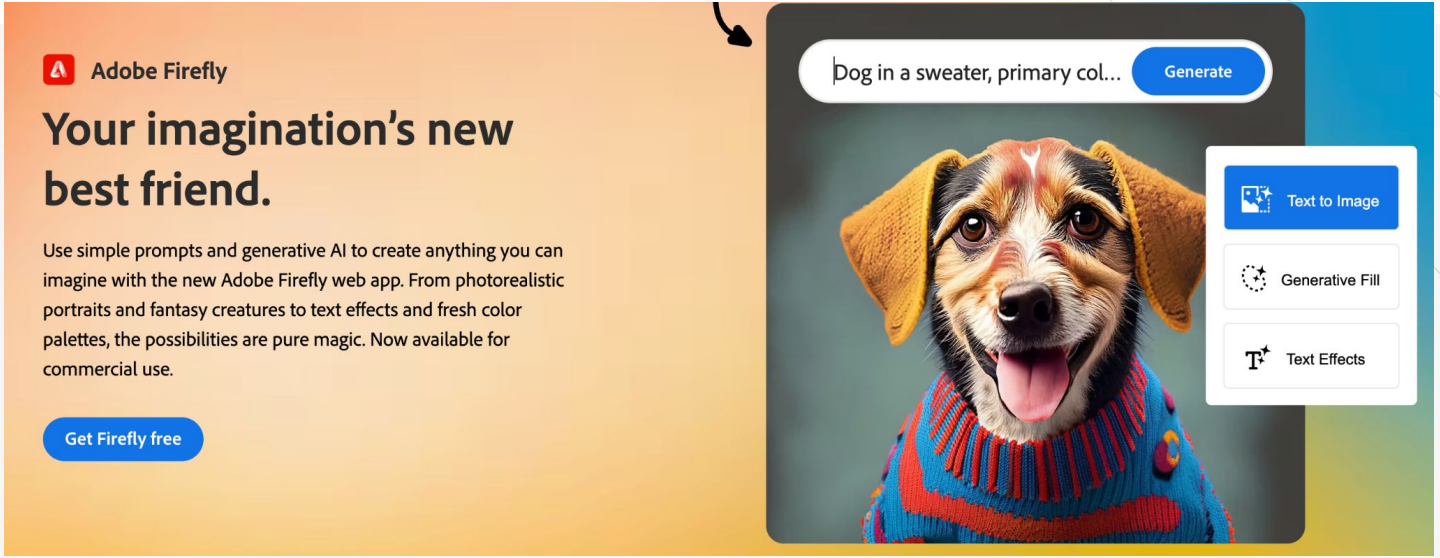

- Private LLM trained on fully-owned images (safe, owned, controlled DATA)
- Users indemnified against IP rights issues
- Integrated into a platform that reaches many millions of users
- Natural and intuitive integrations

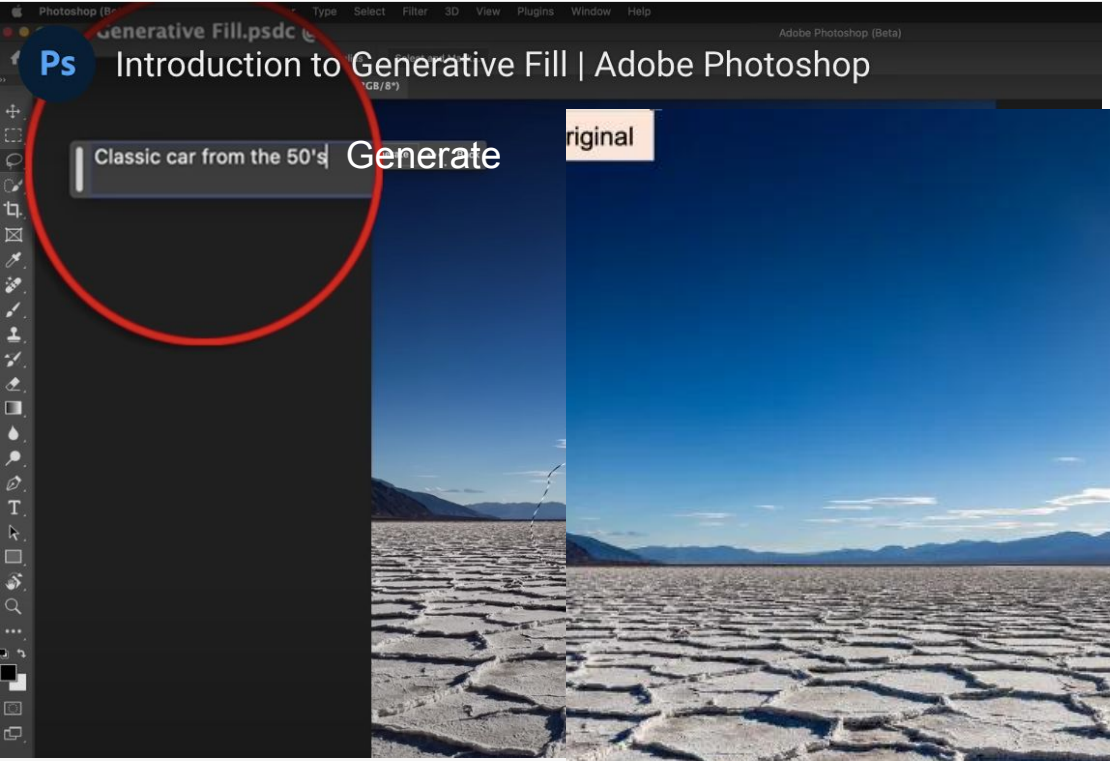

Adobe’s PhotoShop: Generative Fill feature

Concluding Remarks

Now What?

Questions on AI Strategy for Our Students

Can we upskill co-op candidates in every college with AI?
Companies desperately need talent that has AI fluency

How do we scale experiential learning in AI @NU?
AI readiness classes + executing NU AI projects by in-house co-ops

Embrace AI and learning tools in the classroom

Graduating the most AI-savvy student cohorts in academia
Requires pragmatic class experience and AI co-ops

Employer Challenge:
How do you acquire, train, and retain the right talent?

SECRET 7: in Making AI Work (from talk on Secrets of Making AI Work)

7 **Nurture Talent and build AI Culture**

Talent and **CULTURE** are critical, employee & executive education are a must

Culture eats Strategy for Breakfast

Establish a Data and AI **organization with ongoing upskilling and strong talent management**

What is our AI Talent Strategy for Our Company?

- **Levels of Training:**
  a. AI Literacy & Awareness
  b. AI business usage
  c. AI technical usage
  d. AI Technical in-house Expertise
  e. AI for Executives and Leaders
  f. AI Policies, Regulations, and Law

At Northeastern we offer programs in all these

- For **b, c, d:** we believe best approach is **Experiential Education:** *Learn while applying the new technology to your work problems*

- **What do we need to attract the right talent?**
- **What do we need to retain the right talent?**
  - **Startups**
  - **Groups and Centers of Excellence that offer up AI services and know how**
  - **The best upskilling programs for employees as competitive advantage**

Overview of the *Institute for Experiential AI (EAI)* at Northeastern University

5
Research Focus Areas

- AI + Life Sciences
- AI + Health & Human Performance (AI+H2P)
- AI4CaS: Climate & Sustainability
- Responsible AI
- Hybrid GenAI
  Data + Prior Knowledge

3
Advisory Practices

- AI Solutions Hub
- Responsible AI Practice
- AI Ethics Advisory Board

EAI The Institute for Experiential AI
Northeastern University

EAI The Institute for Experiential AI
Northeastern University

100+
Tenure/Tenure Track AI Faculty

50+
Research Scientists, Research Faculty, Postdocs

Focus Area Leaders

Dr. Usama Fayyad
Inaugural Director

Dr. Sam Scarpino
Director, AI + Life Sciences

Prof. Eugene Tunik
Director, AI+H2P

Dr. Cansu Canca
Director, Responsible AI Practice

Prof. Auroop Gangul
Director, AI4CaS

Dr. Ricardo Baeza-Yates
Director of Research

Big Themes: The Institute for Experiential AI @ Northeastern University

Experiential AI	=	AI with the Human in the Loop	Human-assisted AI + AI-assisted Human Intelligence
Applied AI is essential	⇒	Seek applied problems with companies	A use-inspired fundamental research agenda aiming for impact on key problems
Experiential Education		AI Solutions Hub provides applied skills	Workforce-ready graduates Upskilling through Lifetime Learning
Focus on Impact & Leadership		Responsible AI, Use-Inspired Research Focus	Leadership in areas of high need, high impact, and academic innovation

Pragmatic focus enhances our leadership in innovative education and applied academic and industry impact

Focus areas for the Institute for EAI

We selected certain areas of focus where we believe:

- The focus will have **High Impact**
- We have a chance **to be a leader**
- We can leverage **NU strengths** in those areas

3 Focus Areas

AI + Life Sciences

AI+H2P: Health & Human Performance

AI4CaS: Climate & Sustainability

2 Core Areas

Responsible AI

Hybrid GenAI: Data + Prior Knowledge

3 Advisory Practices

AI Solutions Hub

Responsible AI Practice

AI Ethics Advisory Board

Partner Testimonials

"The EAI and Roux Institute team has been an incredible partner by seeking to understand what our needs are, then building an innovative solution that advances our approach to personal customer care and frictionless service. Their partnership has enabled us to execute our goal of deepening customer relationships through AI practices."

Bob Montgomery-Rice,
President and CEO of Bangor Savings Bank

"The world has witnessed rapid AI development over the last few months. However, putting AI to use is not an easy task, partially because there are a lot of unknowns and risks in AI. But the Responsible AI program we have now enables us to explore that in a way that is safe."

Xuning Tang
Associate Director of Responsible AI, Verizon

"Having cross-talk and cross-pollination between the groups and BU leaders happened for the first time ever. That was a huge thing. That is big. That is transformational.

Digital Transformation Lead,
High-Tech Manufacturer

Bangor Savings Bank · Fidelity Investments · FullscopeRMS · Sun Life · verizon

Summary/Concluding Thoughts

1. AI is an enterprise imperative – challenging to make work – But is a big factor in competitiveness in the knowledge economy
2. HOWEVER:

No Data ⇒ No working AI

Capture your IP: events, outcomes, context

Human intervention a must

continuous correction of algorithmic errors

- The biggest threat of AI to humanity
  - Not the super-intelligence
  - But completely disorienting us in our new digital existence

Summary/Concluding Thoughts

The biggest danger to AI today?

The gap between how amazing the demo is and how disappointing the implementation results…

1. Generative AI offers a means for accelerating work, ***but not fully automating it***
2. Generative AI can help reduce robotic, repetitive, and manually intensive work
3. Can be a game changer for efficiency, accuracy and CX (customer experience)
4. National regulation is coming first in the EU and China, then the US
5. Barriers, complexity and costs (if done rationally) of GenAI are coming down – this tech is available to all companies and competitors

Thank you!
Any Questions?

Usama M. Fayyad

u.fayyad@northeastern.edu
Assistant: r.alshami@northeastern.edu

Ready to get started?

Stay connected with the Institute! Subscribe to our monthly newsletter, In the AI Loop.

AI.northeastern.edu

IEAI-NU
UFayyad

@Experiential_ai
@UsamaF

Institute for Experiential AI

Legends of Data & AI – Podcast – on Spotify and other platforms: https://bit.ly/legendsofdata-ai

The Institute for Experiential AI
Northeastern University

The bottom line?

A Large Language Model - has no knowledge or understanding of what it "learned"

- Billions to trillions of weights
- They serve as a glorified "auto-complete" capability

It is amazing & astounding what these stochastic parrots can do!

In our new digital *knowledge economy!*

How does it work with trillions of parameters?

Let's delve a bit deeper into an example…

The second incarnation of ML/neural nets: Deep Learning

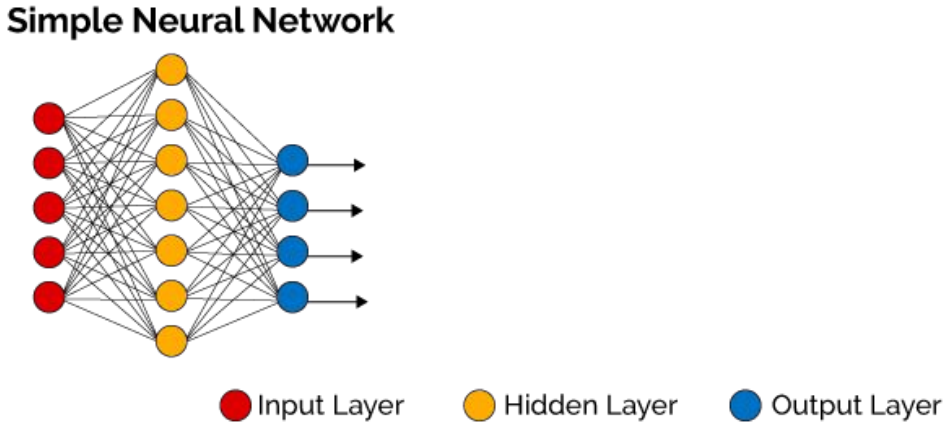

- Resurgence on Neural Networks in 2010 – Deep Learning
  - Not much new, just lots of computation and lots of data
  - Works very effectively in “non-declarative” knowledge
  - Breakthroughs on learning “procedural knowledge” where we don’t know how humans do it – Jeff Hinton in 2012 – image recognition in 22K categories over 15 million images (85% accuracy)
  - Object Recognition
  - Speech Recognition
  - High-dimensional regression/decision making

The Perceptron

(from 1940's to F. Rosenblatt - 1958)

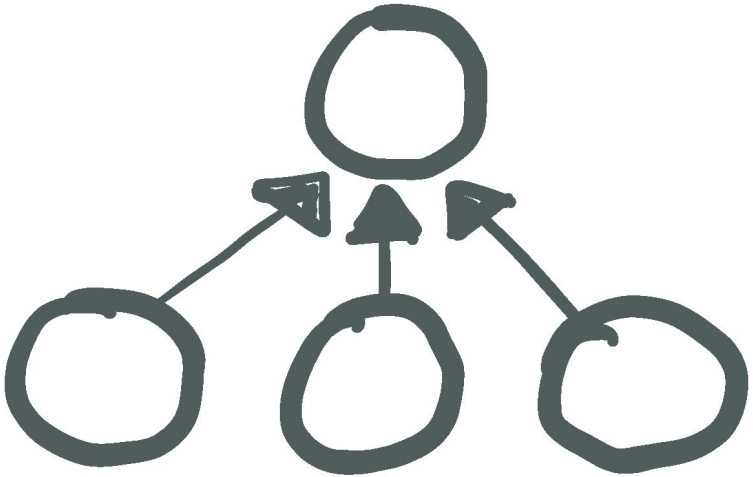

Slides adapted from Primer Talk by ***Prof. Byron Wallace, Northeastern University*** *-* ***Generative AI Workshop: From the Classroom to the Economy -*** *April 2023*

The Perceptron

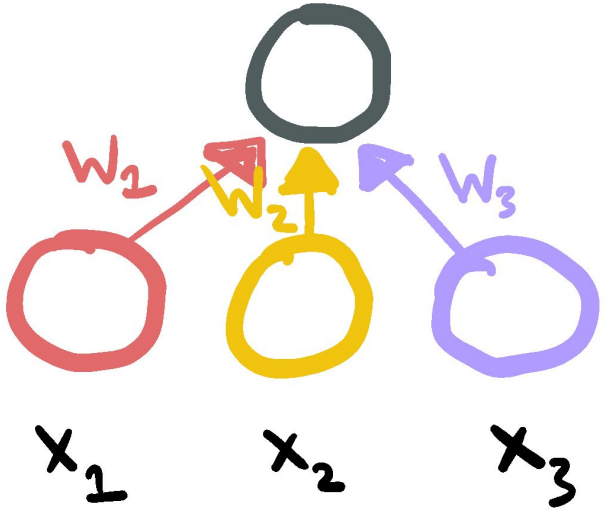

*Slides adapted from Primer Talk by **Prof. Byron Wallace, Northeastern University** - **Generative AI Workshop: From the Classroom to the Economy** - April 2023*

The Perceptron

$$y = \sigma(w_1 x_1 + w_2 x_2 + w_3 x_3)$$

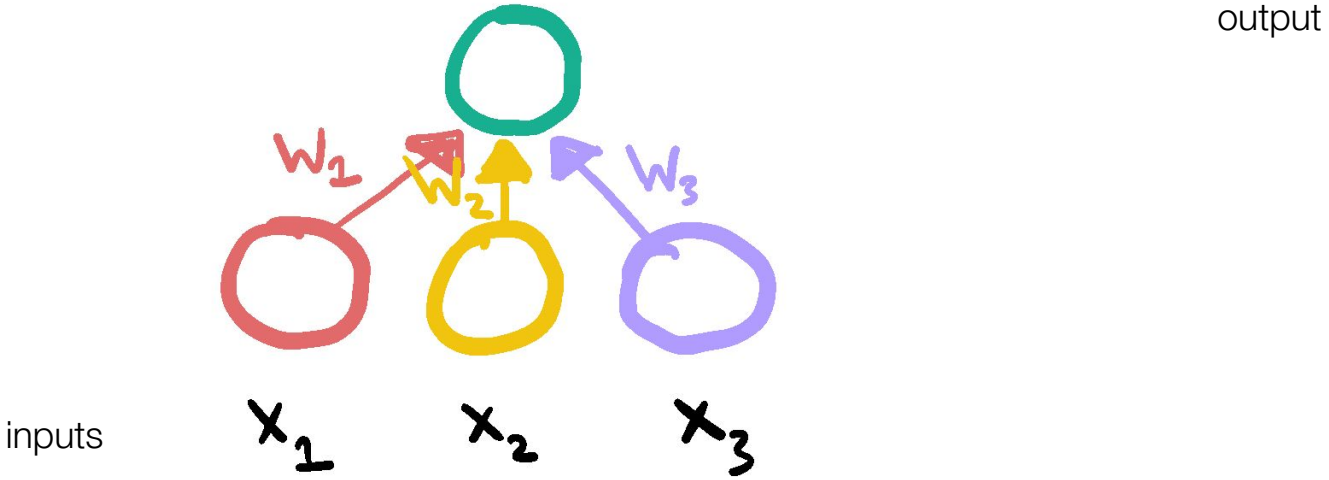

*Slides adapted from Primer Talk by **Prof. Byron Wallace, Northeastern University** - **Generative AI Workshop: From the Classroom to the Economy** - April 2023*

The Perceptron

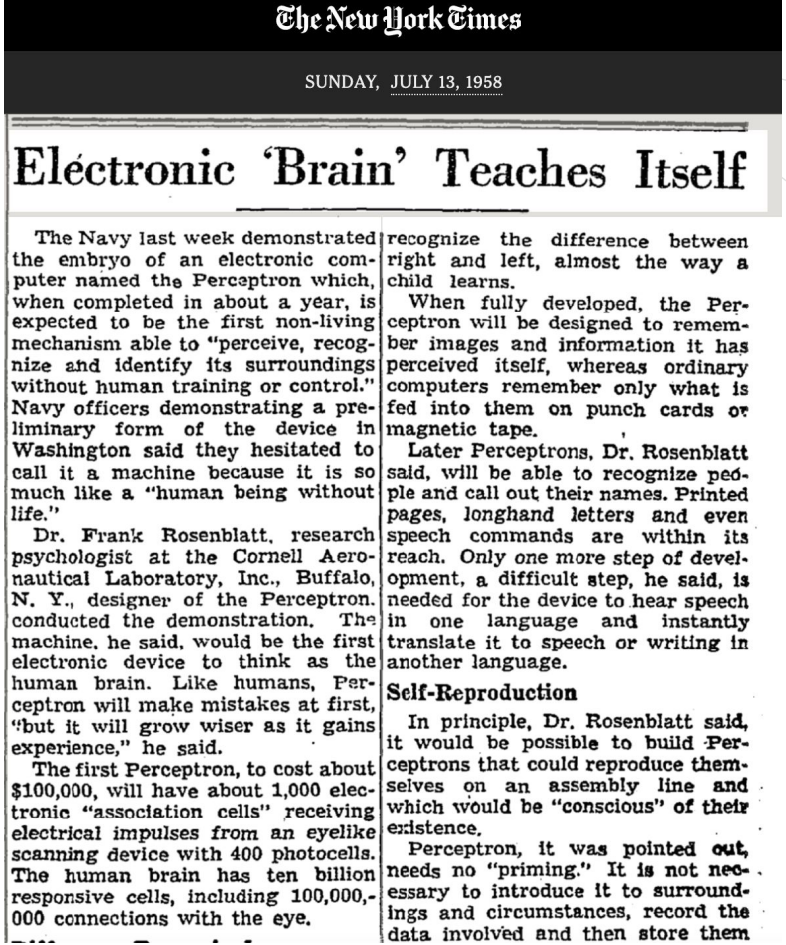

The New York Times

SUNDAY, JULY 13, 1958

Electronic 'Brain' Teaches Itself

The Navy last week demonstrated the embryo of an electronic computer named the Perceptron which, when completed in about a year, is expected to be the first non-living mechanism able to "perceive, recognize and identify its surroundings without human training or control." Navy officers demonstrating a preliminary form of the device in Washington said they hesitated to call it a machine because it is so much like a "human being without life."

Dr. Frank Rosenblatt, research psychologist at the Cornell Aeronautical Laboratory, Inc., Buffalo, N. Y., designer of the Perceptron, conducted the demonstration. The machine, he said, would be the first electronic device to think as the human brain. Like humans, Perceptron will make mistakes at first, "but it will grow wiser as it gains experience," he said.

The first Perceptron, to cost about $100,000, will have about 1,000 electronic "association cells" receiving electrical impulses from an eyelike scanning device with 400 photocells. The human brain has ten billion responsive cells, including 100,000,000 connections with the eye.

recognize the difference between right and left, almost the way a child learns.

When fully developed, the Perceptron will be designed to remember images and information it has perceived itself, whereas ordinary computers remember only what is fed into them on punch cards or magnetic tape.

Later Perceptrons, Dr. Rosenblatt said, will be able to recognize people and call out their names. Printed pages, longhand letters and even speech commands are within its reach. Only one more step of development, a difficult step, he said, is needed for the device to hear speech in one language and instantly translate it to speech or writing in another language.

Self-Reproduction

In principle, Dr. Rosenblatt said, it would be possible to build Perceptrons that could reproduce themselves on an assembly line and which would be "conscious" of their existence.

Perceptron, it was pointed out, needs no "priming." It is not necessary to introduce it to surroundings and circumstances, record the data involved and then store them

*Slides adapted from Primer Talk by **Prof. Byron Wallace, Northeastern University** - **Generative AI Workshop: From the Classroom to the Economy** - April 2023*

Deep neural nets

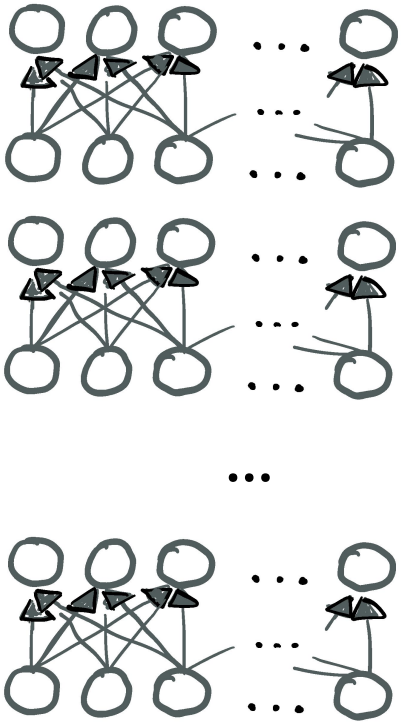

100's of Billions of parameters

Billions of parameters

Millions of nodes

Stack many layers (depth) to get deep networks

Increase the number of inputs

Slides adapted from Primer Talk by ***Prof. Byron Wallace, Northeastern University*** *-* ***Generative AI Workshop: From the Classroom to the Economy -*** *April 2023*

So how does Generative AI work?

The Secrets of Making AI (and GenAI) work?

What is your strategy for:

1. **Narrowing** the problem scope *as much as possible*
2. **Business Case**: with team, exec, and **FINANCE** buy in and agreement to reasonable **ROI** expectations over time
3. **Collect ALL** data surrounding *events, outcomes and context*
4. **Capture ALL** data from **EVERY** human intervention: *when, why, desired outcome, and context* with clear permission and disclosure
5. **Risk assessment processes** for unintended biases, and built-in **RESPONSIBLE AI** with embedded ethics team
6. **Incremental** and gradual, but **CONTINUOUS** improvement *over time*
7. **Talent** and **CULTURE** are critical, employee & executive education are a must

BloombergGPT: a finance-aware LLM

- BloombergGPT, a 50-billion parameter large language model, purpose-built from scratch for finance [March 30, 2023]
- BloombergGPT represents the first step in the development and application of this new technology for the financial industry.
- Assists Bloomberg in improving existing financial NLP tasks, such as sentiment analysis, named entity recognition, news classification, and question answering, summarization, among others.
- BloombergGPT: What? So What? Now what?

[BloombergGPT ARXIV: https://arxiv.org/abs/2303.17564]

BloombergGPT: News headlines generator

INPUT	**Input**: The US housing market shrank in value by $2.3 trillion, or 4.9%, in the second half of 2022, according to Redfin. That's the largest drop in percentage terms since the 2008 housing crisis, when values slumped 5.8% during the same period.
OUTPUT	**Output**: Home Prices See Biggest Drop in 15 Years

Example: Bloomberg

Generate Bloomberg queries (BQL) from natural language

INPUT

OUTPUT

Input: Get me the last price and market cap for Apple

Output: `get(px_last,cur_mkt_cap) for(['AAPL US Equity'])`

Example: Bloomberg

BloombergGPT: So what

- **SO WHAT**
  - This domain-specific language model allows Bloomberg to develop **many new types of applications** and achieve much **higher performance** than with custom models for each application - ***all with a faster time to market.***

sentiment analysis

auto entity recognition

answering financial questions

summarization

headline generation

[BloombergGPT ARXIV: https://arxiv.org/abs/2303.17564]

BloombergGPT: Some Questions

- Why a 50-billion parameter large language model?
  - Because team had a “compute” budget of $3.5M

- What is the significance BloombergGPT for Bloomberg?
  - Addressed some really important high-business value problems.
  - The approach generalized to many problems: sentiment analysis, named entity recognition, news classification, and question answering, summarization, among others.

- BloombergGPT: What? So What? Now what?

[BloombergGPT ARXIV: https://arxiv.org/abs/2303.17564]

Xfinance LLM 13B Outperforms BloombergGPT

- **Perform both unsupervised fine-tuning and instruction fine-tuning on the LLaMA 13B model (Stockastic.AI)**

- Fine-tuning on a GCP cluster of 8 A100 80GB GPUs over 24 hours **at a cost of $1,000**
- Outperforms BloombergGPT on a range of financial applications

25% of the size of BloombergGPT

1000x cheaper? Much more?

More Robust, easier to maintain

Major Developments over the Last Year (5)

Many developments on technical fronts…

- **LM Agents**
- **Student-teacher LLM models**
- **Synthetic data**
- **Research to understand what happens inside the “black box”**

Recent Case-Study - Cybersecurity Hacking

TECHNOLOGY

https://newatlas.com/technology/gpt4-autonomously-hack-zero-day-security-flaws/?utm_source=flipboard&utm_content=other

GPT-4 autonomously hacks zero-day security flaws with 53% success rate

By Joe Salas June 08, 2024

GPT-4 was able to exploit 87% of critical-severity CVEs (*Common Vulnerabilities & Exposures*) on its own.

HPTSA has shown to be **550% more efficient than a single LLM in exploiting vulnerabilities: 8 of 15 zero-day vulnerabilities.**
Solo LLM was able to hack **only 3/15** vulnerabilities.