ReWork AI in Finance Summit New York - April 18 - 19, 2024

The Double-Edged Sword of AI in Finance: *Secrets of Making AI Work*

Keynote Talk

Usama M. Fayyad

Inaugural Director of the Institute for Experiential AI - Northeastern University

Professor of the Practice, Khoury College of Computer Sciences

EAI The Institute for Experiential AI
Northeastern University

Northeastern University

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***

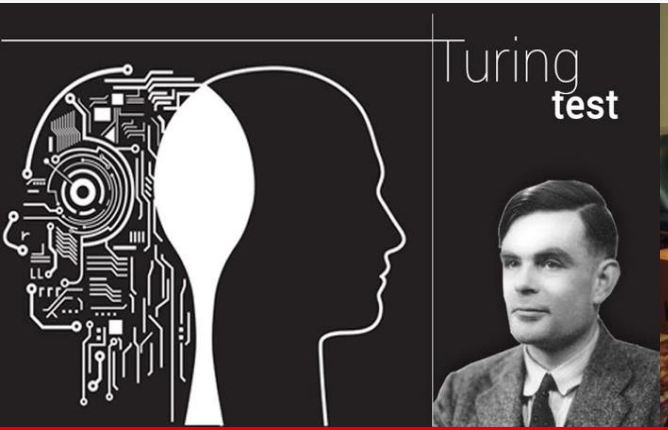

Sound like familiar hype?

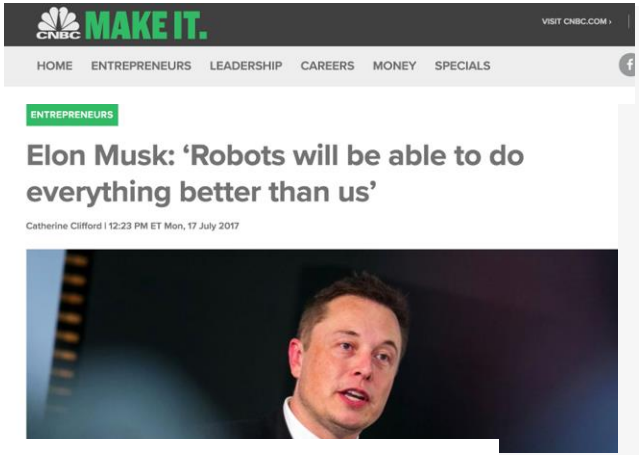

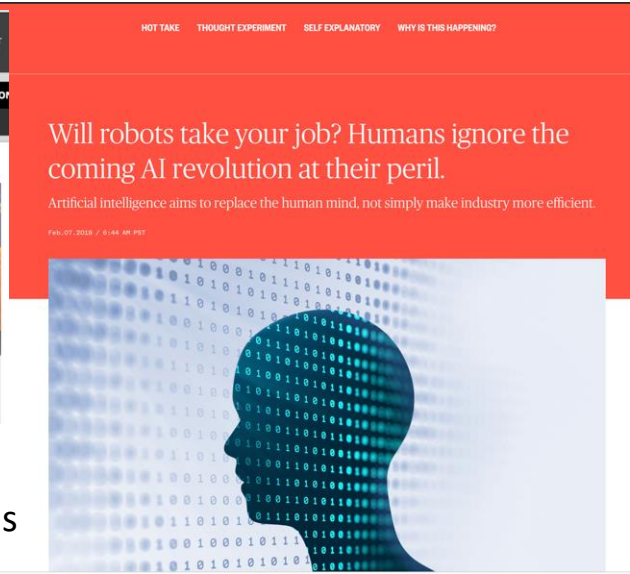

- Major hype in the 1980's – AI was going to solve all problems and change the world
  - U.S. was afraid of Japan AI program – 5th Gen. Systems

- 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

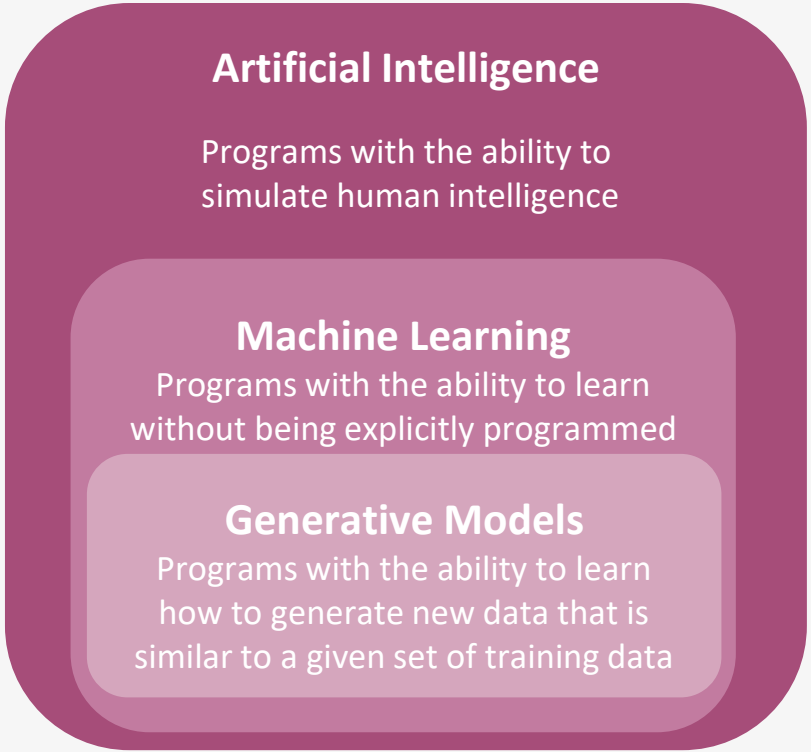

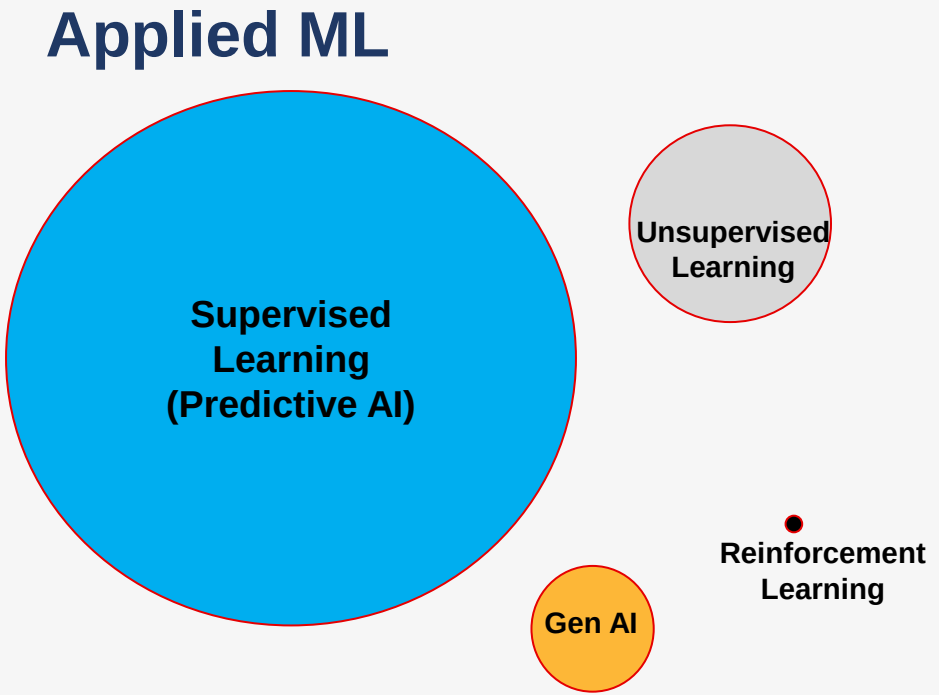

Human feedback

After “pre-training”, tune models to better align with *human feedback*

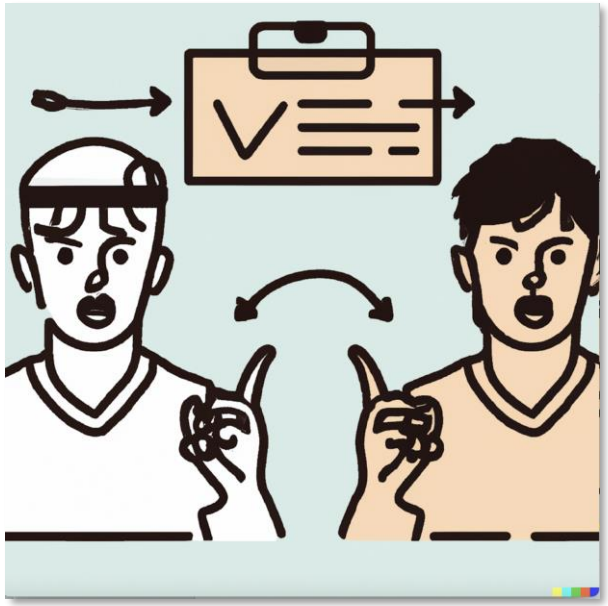

Slides adapted from Primer Talk by ***Prof. Byron Wallace, Northeastern University*** *-* ***Generative AI Workshop: From the Classroom to the Economy -*** *April 2023*

SECRET 1: in Making AI Work

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

Complete knowledge is impossible unless your focus is extremely narrow

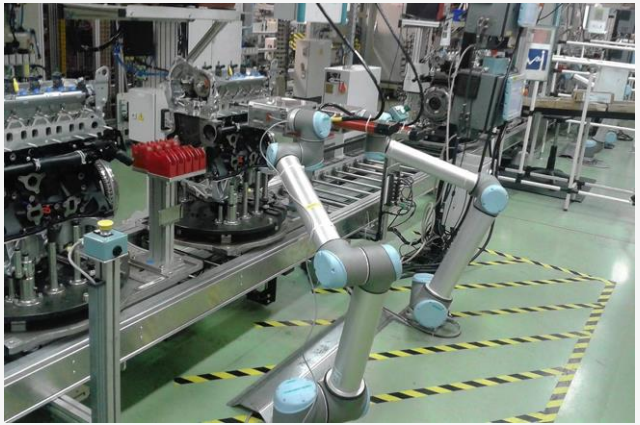

Reduce the problem domain to one where "complete knowledge" is possible

How about "Grand Challenge problems" like Machine Vision in Commerce?

- Look for simplifications - recognizing objects is too hard?
- Avoid image analysis completely

too hard to figure out items in a basket?

9 87654 32109 8

Any relation to how humans see?

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.

SECRET 2: in Making AI Work (reminder)

2 **Business Case**: with team, exec, and **FINANCE** buy in and agreement to reasonable **ROI** expectations over time

Latest in AI can be very expensive…

Depending on benefit expensive may be OK… Optimize later…

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

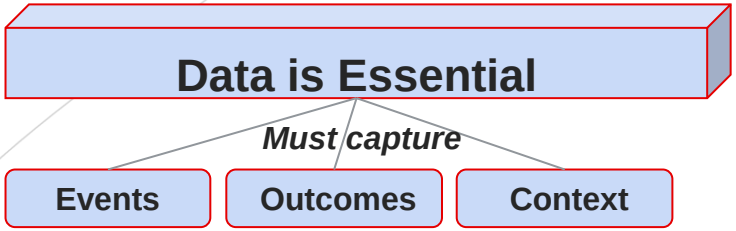

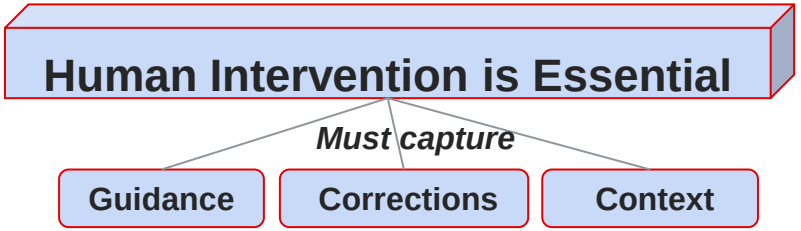

most organizations struggle with the basics of making data work as an asset

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

JPMorgan Chase in the Headlines (4/2023)

"Data" was the #1 most used word in the letter - 27 times. JPMorgan has identified 300 use cases for AI that are already in practice at Chase.

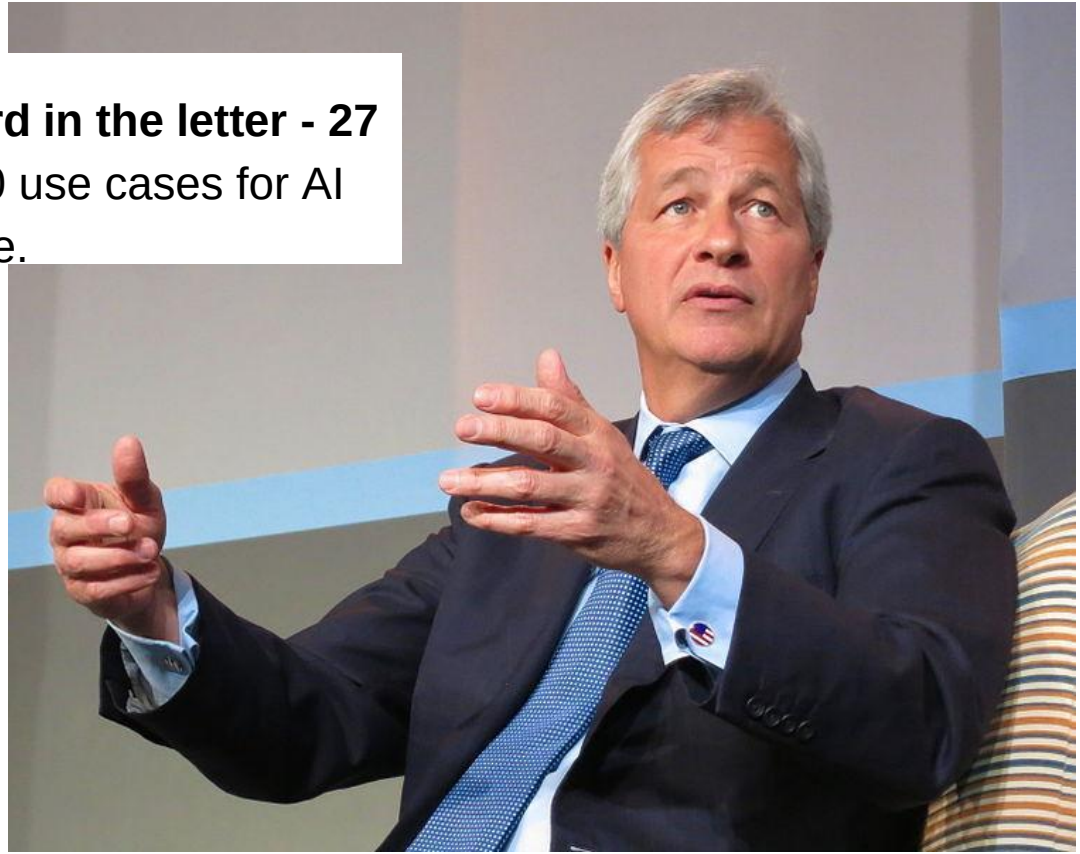

Jamie Dimon's annual Shareholder Letter for 2023

"Artificial intelligence (AI) is an extraordinary and groundbreaking technology. AI and the raw material that feeds it, data, will be critical to our company's future success — the importance of implementing new technologies simply cannot be overstated. We already have more than 300 AI use cases in production today for risk, prospecting, marketing, customer experience and fraud prevention, and AI runs throughout our payments processing and money movement systems across the globe. "

—Jamie Dimon, CEO of JPMorgan Chase

Example - From Traditional Banks to Digital Banks

Data is the key to restoring the lost customer intimacy in the digital interactions era

100 years ago

Now

Front office was intimate

- Direct interactions
- Personal knowledge - Staff knows all that is happening with client and family
- Personalized service automatic

Front office has no knowledge or intimacy

- More complex product line
- Data silos and high latencies
- No unified view and understanding of the customer
- Personalization a challenge

Back office was simple

- Easy to understand risk
- Easy to score and set limits by intuition
- KYC trivially easy and natural
- Controls straightforward

Back office is overcomplex/manual

- Difficult to scale because tech did not evolve
- Risk and Finance expensive hard to manage because the data is a mess and hard to access
- Controls a challenge

Examples of AI in Financial Services

Credit & Risk

- Underwriting & pricing model optimization
- Dynamic credit limits
- Use of 3rd-party and unstructured data
- Risk modeling and scenario analysis

Customer Service & Personalization

- Dynamic customer profiling, segmentation
- Personalized content, product/service recommendations
- Intelligent call routing and service intervention
- Agent support, Chatbot service augmentation

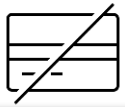

Fraud Detection & Prevention

- Credit Card Fraud
- Insurance claims fraud
- Money laundering

AI & Trading

- Data analysis and trading support
- Portfolio optimization
- Modeling and scenario analysis
- News headline generation

Back office & Process Automation

- Data extraction
- Document capture & processing
- Document & agreement review

Responsible AI

- AI Ethics Governance & Strategy
- Technical AI Audit & Ethics Assessment
- Responsible AI Training

FinTech: Why is AI necessary?

Scale is a must – human processing is not scalable or feasible

- **AI algorithms for "understanding" – of context and customer**
  - Leverage human judgement in delivery to build the right training data sets and KB
  - Need to make sure AI algos are subject to Responsible AI criteria (often overlooked)
- **Complexity of products requires "reasoning":** Still a big challenge in AI

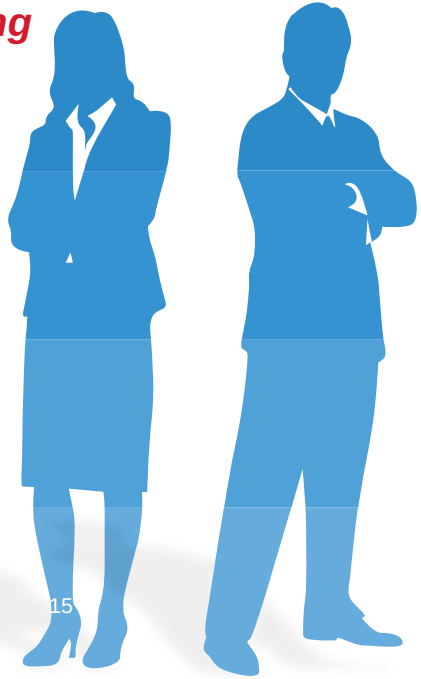

New capabilities (tech, network, new data sources require algortihms

- **Novel risk and credit scoring opportunities**
- **Integration into microservices**
- **Leverage networks (e.g. social, economic, commerce networks) and other viral services**
- **Customer service centers**
- **Process acceleration**

FinTech: Some of the Big Challenges for AI

Fairness and bias in algorithms can have big consequences

Financial decisions are much more **consequential** (than e.g. *targeting ads*)

Algorithms have little or no reasoning capability (or even common sense) – all they know is data (*w/ almost no context*)

Modeling complex decisions & consequences is a hard problem

Advanced tech and data enable **potentially deep intrusion** on **privacy & civil rights**

GenAI can be **Unstable** and in **unpredictable ways**

In new situations, random actions come out

FinTech: 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

SECRET 3: in Making AI Work

3 Capture ALL the data "**Events, Outcomes, and Context**" at as fine a granularity as possible

Remember: Data in Structured and Unstructured

Majority of Data in any organization is Unstructured (90% per Gartner)

Digitization Produces 100x the Data Flux

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

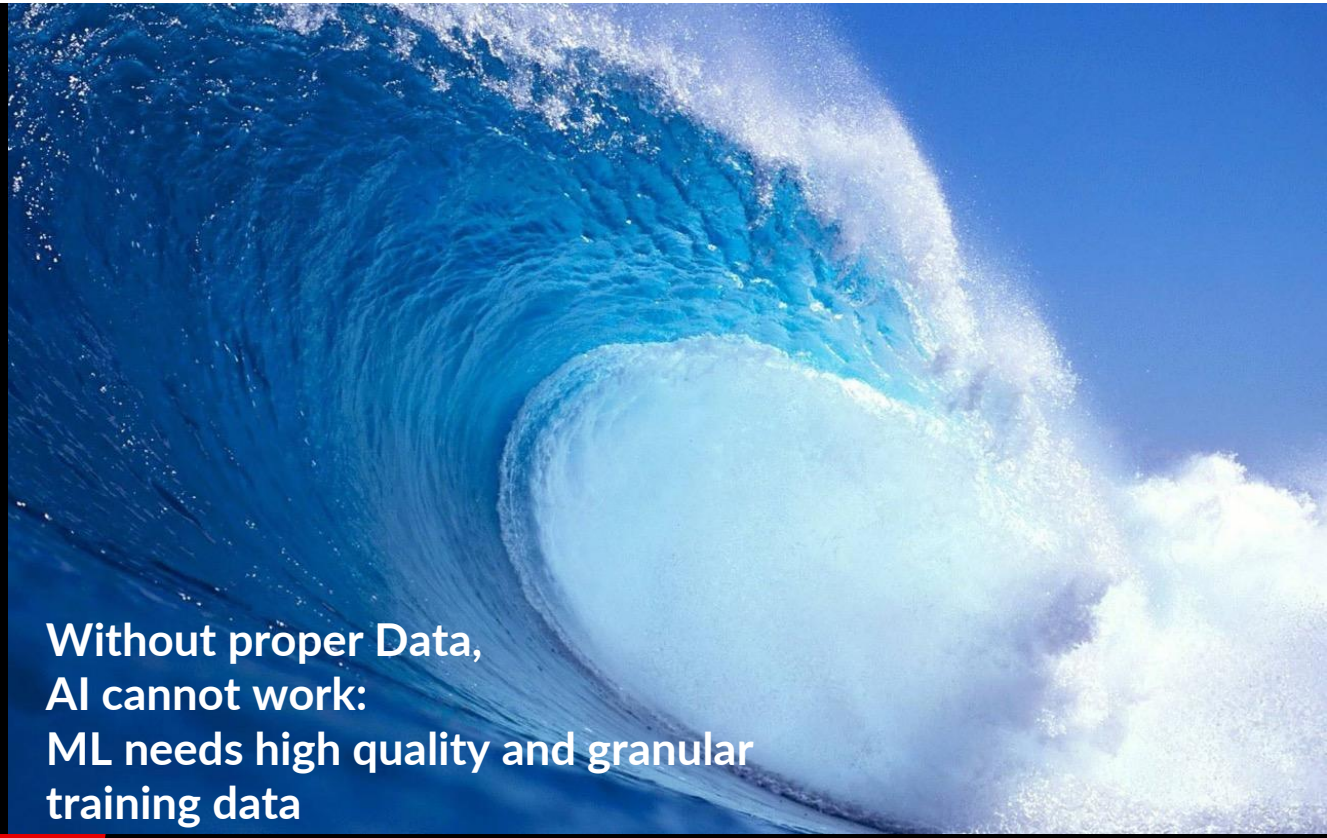

Without proper Data,
AI cannot work:
ML needs high quality and granular training data

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.*

Is There a Human in the Loop (HITL) in ChatGPT?

Much speculation about pure AI (AGI) or much human intervention?

- Strong evidence that human editorial review is applied
- Some questions are answered by humans
- Generally, this is a good sign in our opinion
- Does raise issues about “intelligence” and “reasoning”
- This is a best practice – we call it Experiential AI – many do it:
  - Google MLR
  - Amazon recommendations
  - Many intervention-based relevance feedback

https://mindmatters.ai/2023/01/found-chatgpts-humans-in-the-loop/

SECRET 4: in Making AI Work

4 **Capture ALL** data from **EVERY** human intervention: *when, why, desired outcome, and context* with clear permission and disclosure

Remember: in most organization this data is never Captured

Leakage of IP into Data Exhaust – a true waste of the most valuable resource…

Generative AI

Now What?

Speed of Adoption

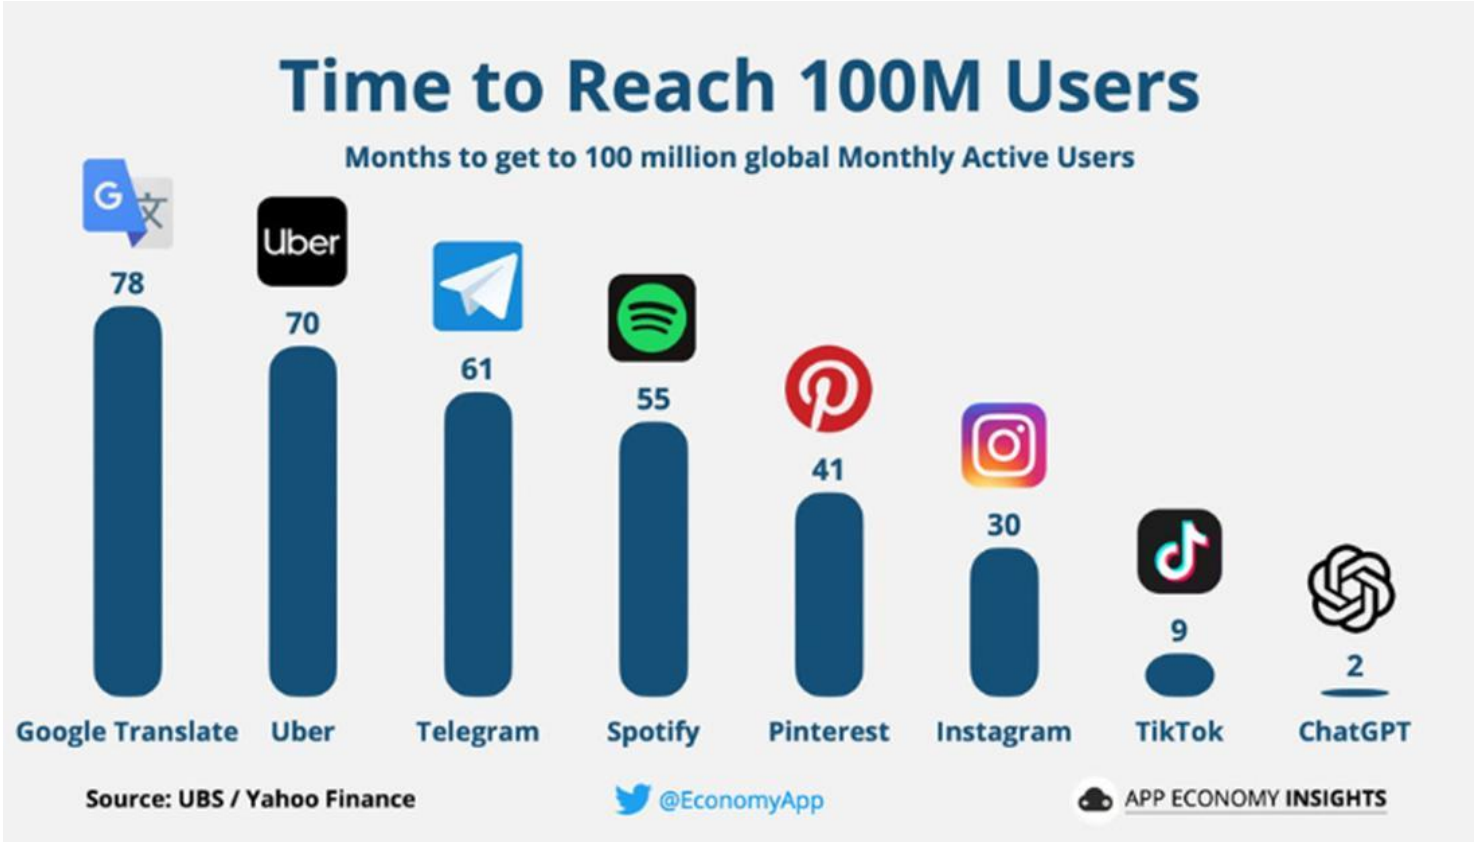

Economic Impact

- 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

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?

Ask me and we can explain it in 10 slides!

Human feedback

After “pre-training”, tune models to better align with *human feedback*

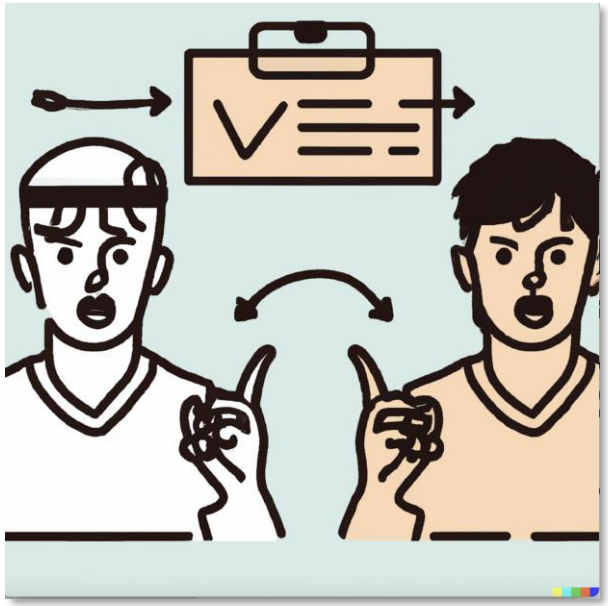

Slides adapted from Primer Talk by ***Prof. Byron Wallace, Northeastern University*** *-* ***Generative AI Workshop: From the Classroom to the Economy -*** *April 2023*

Generative AI: The good, the Bad, & the Ugly

The Good

- ✓ New Way to Search
- ✓ Productivity increases for knowledge workers
- ✓ Personalized Experiences
- ✓ Variety of Applications
- ✓ Less effort to produce results

The Bad

- ✗ Limited Context Window
- ✗ Loss of Factual Evidence
- ✗ Misinformation Heaven
- ✗ Heavy price in energy & compute infrastructure

The Ugly

- ✗ Errors - “Hallucinations”
- ✗ Data Confidentiality
- ✗ IP Rights and Issues
- ✗ Easy to hack

WHY Responsible AI?

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

ACM's Statement on Responsible Algorithmic Systems

1. Legitimacy and competence
2. Minimizing harm
3. Security and privacy
4. Transparency
5. Interpretability and explanation
6. Maintainability
7. Contestability and auditability
8. Accountability and responsibility
9. Limiting environmental impacts

RAI Governance

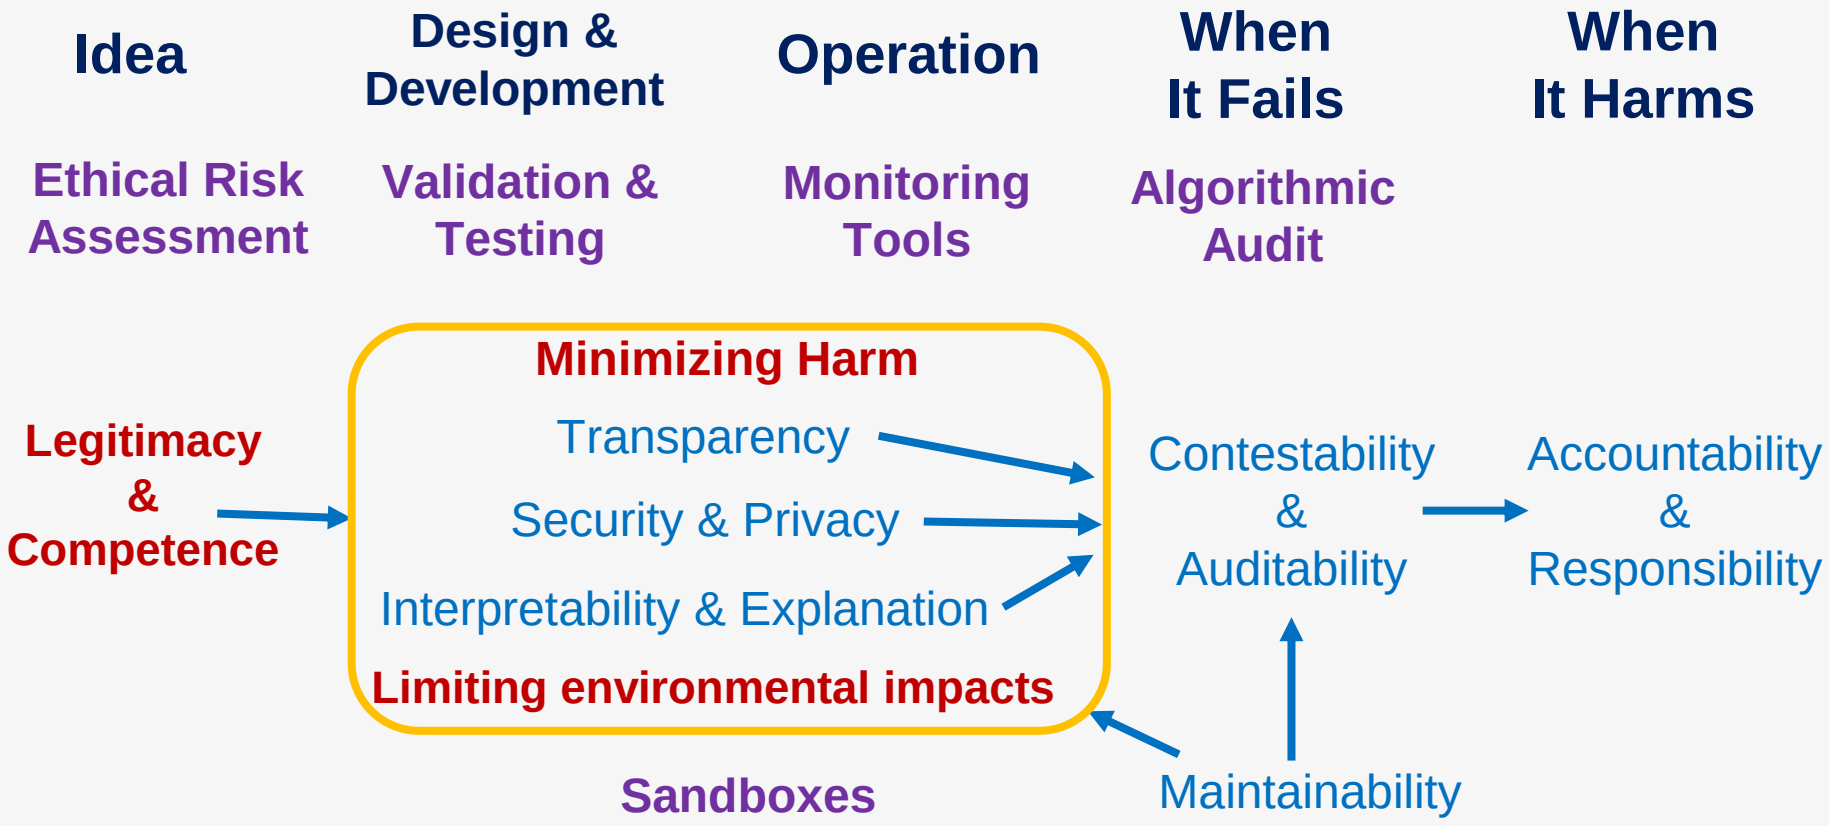

News@Northeastern

NORTHEASTERN LAUNCHES AI ETHICS ADVISORY BOARD TO HELP CHART A RESPONSIBLE FUTURE IN ARTIFICIAL INTELLIGENCE

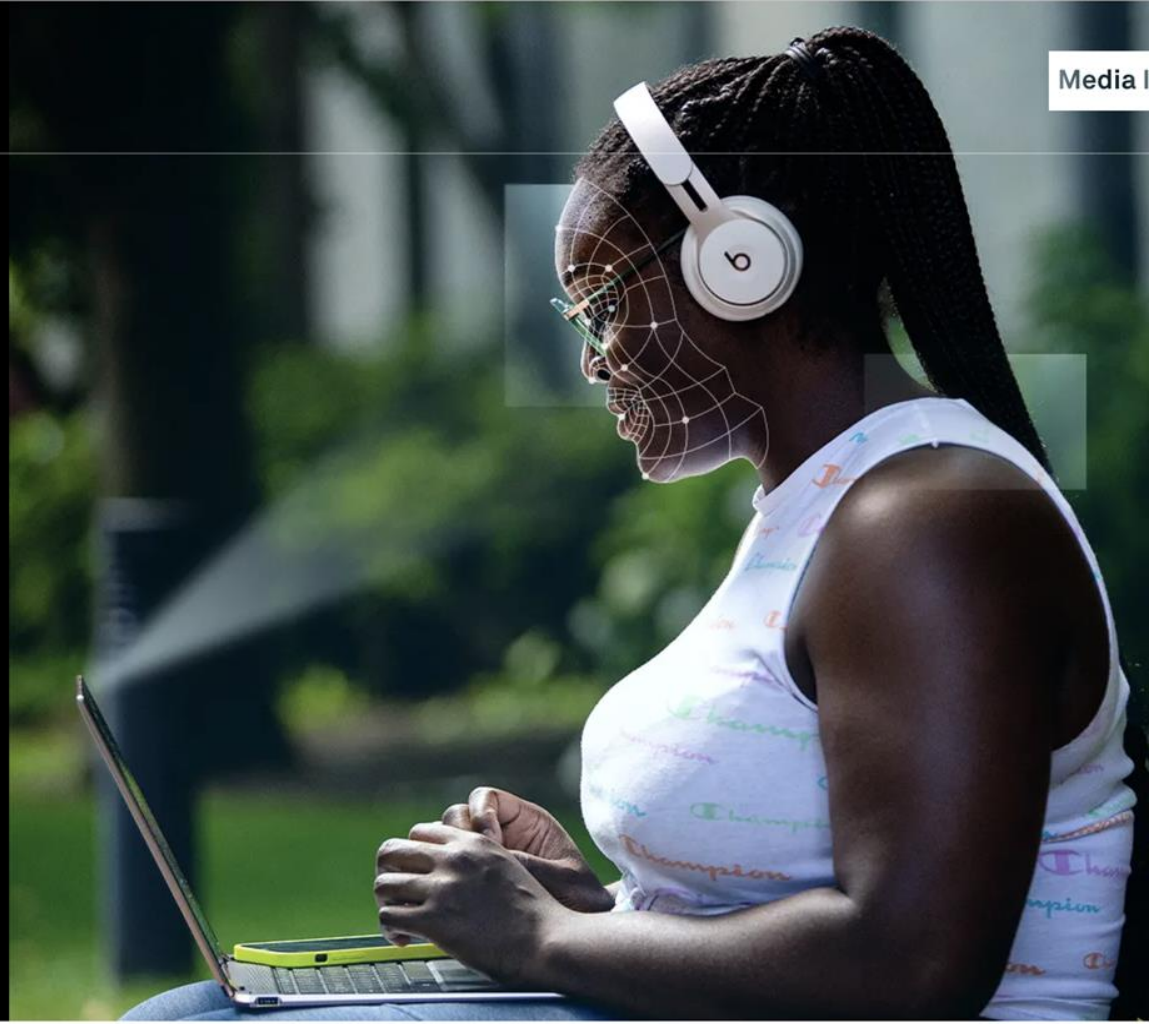

Illustration by Zach Christensen/Northeastern University

by **Cody Mello-Klein** July 28, 2022

CASE STUDY

verizon

The Institute for Experiential AI
Northeastern University

Experiential AI for the Largest Wireless Carrier in the United States

How Verizon Wireless partnered with experts at the Institute for Experiential AI to integrate Responsible AI into its innovation roadmap

OVERVIEW OF PROJECT

CLIENT:
Verizon Wireless

PROJECT:
Responsible
AI Governance

- AI Ethics Roadmap
- Ethics Risk Assessment Framework
- Fairness Monitoring Methodology

KALYIANI SEKAR
Chief Data & Analytics Officer

MEGHNA SINHA
VP of Artificial Intelligence and Data

XUNING TANG
Associate Director of Responsible AI

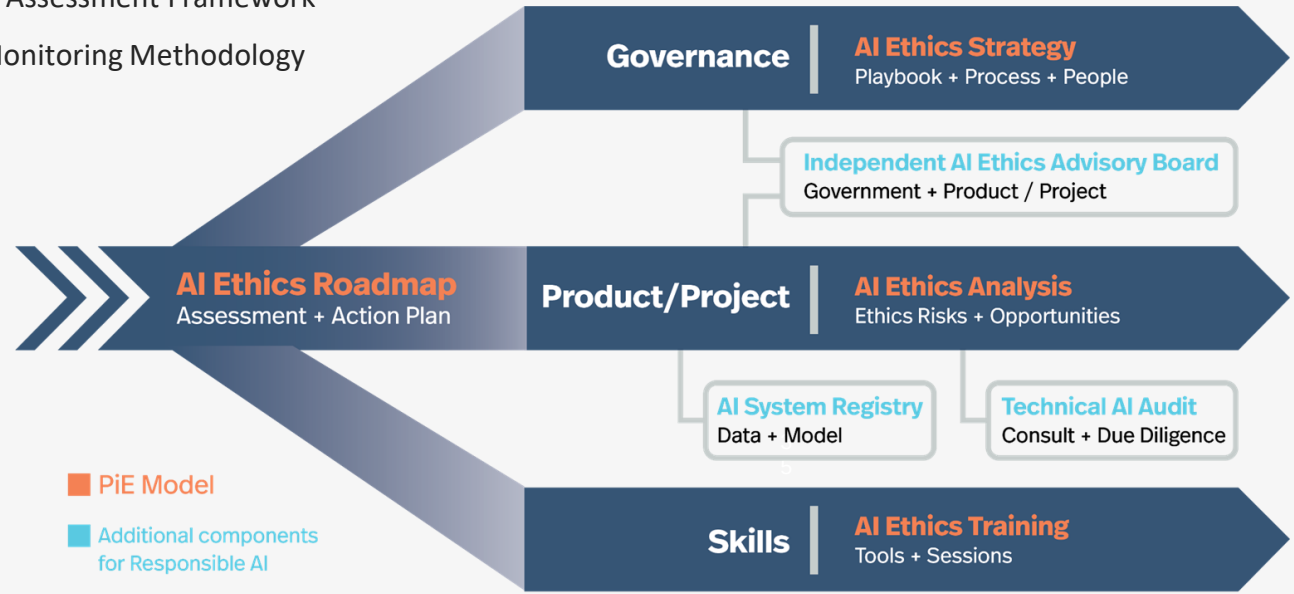

Example Project Structure | Comprehensive RAI Governance

RAI START

RAI Maturity Assessment
- Data collection
- Existing governance review
- Review of RAI workflow

Action-Plan Creation
- Stakeholder consultations
- Need and priority identification
- Action guiding roadmap

DEVELOP POLICY

Value Identification
- Define & operationalize organizational values

Tool+Guideline Dev.
- To mitigate risks and identify opportunities

RAI Training
- Leadership/workforce training to upskill for defined RAI roles

DEPLOY GOVERNANCE

Workflow Integration
- Integrating tools and guidelines from Phase 2 within all relevant teams (e.g. product design or data management)

Organizational Change
- Equipping teams to work together across the organization to solve RAI challenges as they appear

SECRET 5: in Making AI Work

5 **Risk assessment processes** for unintended biases, and built-in **RESPONSIBLE AI** with embedded ethics team

Remember: in most organization RAI is an after-thought

Risks can create debilitating Reputational Damage – and serious economic costs and liability…

Are bigger LLMs better?

Sample Case Study in Finance

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

- The language model quickly learned Bloomberg's BQL database language with just a few examples.

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

41

BloombergGPT: So what

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

- **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: 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

SECRET 2: in Making AI Work (reminder…)

2 **Business Case**: with team, exec, and **FINANCE** buy in and agreement to reasonable **ROI** expectations over time

Remember: in most cases Larger LLM is not necessarily better

Larger LLM is not only more expensive – but much more likely to be unstable and difficult to maintain…

Generative AI

The Issues

Generative AI: The good, the Bad, & the Ugly

The Good

- ✓ New Way to Search
- ✓ Productivity increases for knowledge workers
- ✓ Personalized Experiences
- ✓ Variety of Applications
- ✓ Less effort to produce results

The Bad

- ✗ Limited Context Window
- ✗ Loss of Factual Evidence
- ✗ Misinformation Heaven
- ✗ Heavy price in energy & compute infrastructure

The Ugly

- ✗ Errors - “Hallucinations”
- ✗ Data Confidentiality
- ✗ IP Rights and Issues
- ✗ Easy to hack

Bad Uses

- Fraud
- Impersonation
- Social Engineering
- Cybercrime

- Human-likeness
  (without transparency)

Responsible AI for LLMs

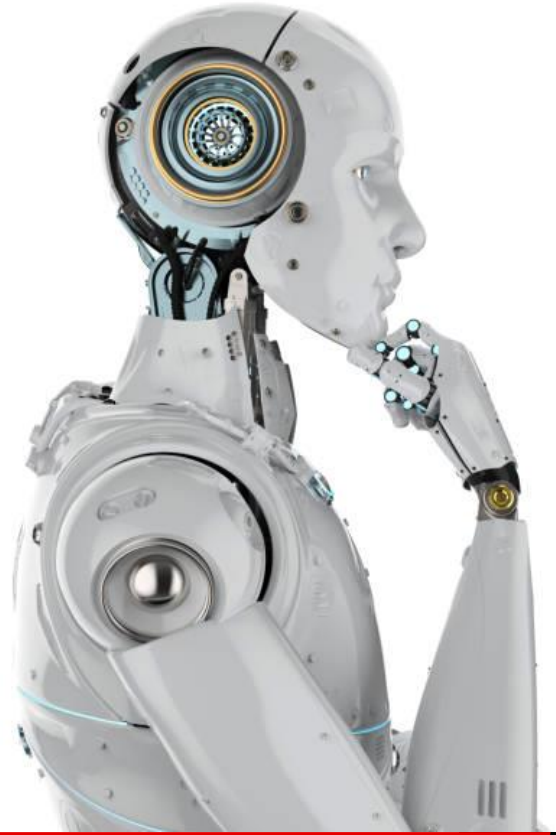

Stay lucid about what LLMs can and cannot do.

- 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

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!**

Concluding Remarks

Now What?

What is Your AI Strategy?

Are you capturing the right data?

Are you closing the loop with ML/Data Science?

Do you know where to use Predictive AI vs GenAI?

What applications of AI give you highest competitive advantage?

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

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

Summary/Concluding Thoughts

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. It 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 competitors and fintechs

Summary/Concluding Thoughts

1. AI is an enterprise imperative in the Digital Age, yet challenging to make work
2. It is a big factor in competitiveness in the knowledge economy
3. HOWEVER:

No Data ⇒ No working AI

Capture your IP: events, outcomes, context

Human intervention a must

continuous correction of algorithmic errors

- Digitization ⇒ 10x to 100x more data - good news for AI?
- Getting the data/context story right is the key enabler for business insights & AI
- There is a rational approach to getting to data assets - and capturing valuable human interactions ⇒ Experiential AI

Thank you!
Any Questions?

USAMA M. FAYYAD

U.Fayyad@northeastern.edu

Assistant: R.Alshami@northeastern.edu

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

AI.northeastern.edu
www.open-insights.com

IEAI-NU
UFayyad
OpenInsights-tech

@Experiential_ai
@UsamaF

Institute for Experiential AI
Open Insights

EAI The Institute for Experiential AI
Northeastern University

Usama Fayyad, Executive Director, Institute for Experiential AI (EAI) at Northeastern University
Professor of the Practice, Khoury College for Computer Sciences, Northeastern University

Goal: Make AI and Data usable, useful, manageable - democratize the responsible use of AI across fields

Education

Large Orgs

Startups

Education

- **Ph.D. Computer Science & Engineering (CSE) in AI/Machine Learning**
- **MSE (CSE), M.Sc. (Mathematics)**
- **BSE (EE), BSE (Computer Engin)**

Academic achievements

- Fellow: Association for the Advancement of Artificial Intelligence (AAAI) and Association for Computing Machinery (ACM)
- Over 100 technical articles on data mining, data science, AI/ML, and databases.
- Over 20 patents, 2 technical books.

- First in industry: Chief Data Officer at Yahoo!
- First Global Chief Data Officer & Group Managing Director at Barclays Bank, London
- Chaired/started major conferences in Data Science, Data Mining, AI
- Founding Editor-in-Chief on two key journals in Data Science