Lecture 10

Large Language Models and Societal Impacts of AI

Haoyu Yue / yohaoyu@washington.edu
Ph.D. Student, Interdisciplinary Urban Design and Planning
University of Washington

RE 519 Real Estate Data Analytics and Visualization
Course Website: www.yuehaoyu.com/data-analytics-visualization/
Autumn 2025

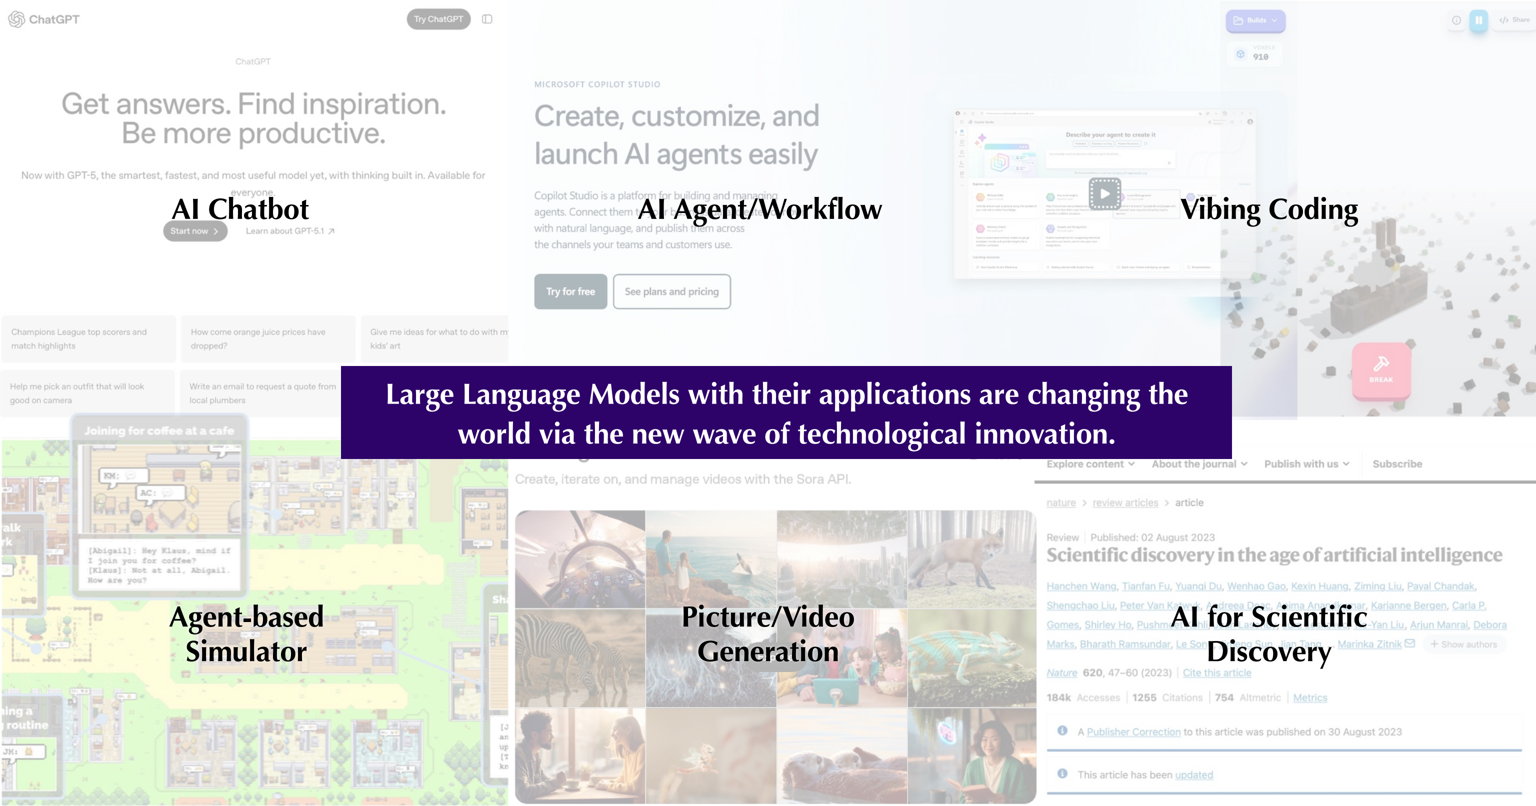

AI Chatbot

AI Agent/Workflow

Vibing Coding

Large Language Models with their applications are changing the world via the new wave of technological innovation.

Agent-based Simulator

Picture/Video Generation

AI for Scientific Discovery

Basics of Large Language Models

Language Models and Encoding

A language model (LM) is a model that assigns a probability to a sequence of words. Its fundamental goal is **predicting the probability of the next word given the previous words**, which is not an easy task.

First step, we need to convert texts to something computer can understand. There are many different *encoding* approaches.

Word-level or char-level may be problematic

Subword tokenization, such as Byte-Pair Encoding (BPE)

- Token is not a word
- Space may be part of token
- May not invariant to case changes
- Some approaches support multilanguage

OpenAI o200k_base (tokenizer for GPT4.1/5)

不動産データの分析と可視化

3428, 24009, 57731, 20951, 124101, 3385, 32648, 5330, 6571, 30395, 11415

OpenAI o200k_base (tokenizer for GPT4.1/5)

real·estate·data·analytics·and·visualization

13187, 12106, 1238, 33199, 326, 71302

OpenAI o200k_base (tokenizer for GPT4.1/5)

Real·Estate·Data·Analytics·and·Visualization

16418, 22603, 4833, 36331, 326, 183695

Microsoft phi-2 (tokenizer for MS's small models)

Real·Estate·Data·Analytics·and·Visualization

15633, 23015, 6060, 30437, 290, 15612, 1634

Source: https://tiktokenizer.vercel.app/

Basics of Large Language Models

Embeddings

Real·Estate·Data·Analytics·and·Visualization
16418, 22603, 4833, 36331, 326, 183695

The objective of a language model is, for example:

- given "Real" → predict "Estate"
- given "Real Estate" → predict "Data"

But those token IDs do not have any meaning other than their ID. In a machine learning framework, we need **X, features**, to train model and make predictions on next token.

Therefore, we use **embedding** to have a numeric representation of each word by mapping token ID to a high-dimensional vector.

8-dimensional embeddings (an example)
There are more dimensions for real models (Llama 3 8B – 4096; DeepSeek-R1 – 7169)

Token	e_1	e_2	e_3	e_4	e_5	e_6	e_7	e_8
Real	0.12	−0.44	0.90	−1.22	0.55	0.03	−0.88	1.41
Estate	−0.77	0.31	−0.05	1.09	−0.33	0.14	0.72	−0.56
Data	1.03	−0.11	0.66	−0.44	0.02	1.28	−0.60	0.33
Analytics	0.22	0.88	−1.11	−0.02	0.91	−0.55	0.04	0.77
and	−0.12	0.05	0.08	−0.33	−0.01	0.11	−0.22	0.09
Visualization	0.55	−1.28	0.41	0.77	−0.66	0.88	1.04	−0.22

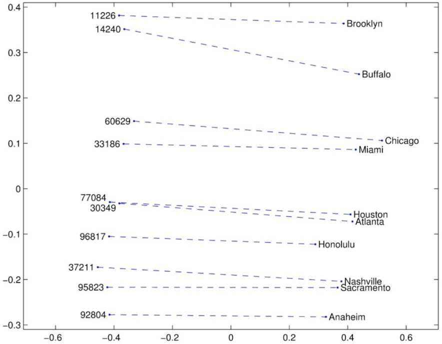

Embedding can capture some semantical meanings. For example: vector("King") - vector("Man") + vector("Woman") ≈ vector("Queen") (Mikolov, 2013)

Source: https://nlp.stanford.edu/projects/glove/

Basics of Large Language Models

Transformer and Pretraining

Transformer: a new deep learning framework: each token will look at all input tokens at the same to decide which ones are most important to decide the next token. Seminal paper: **Attention Is All You Need** (2017, Google)

Pretraining: by reading a massive number of tokens, the model learn grammar, semantics, world knowledge, style patterns. The objective is predicting the probability distribution of token t given all token before.

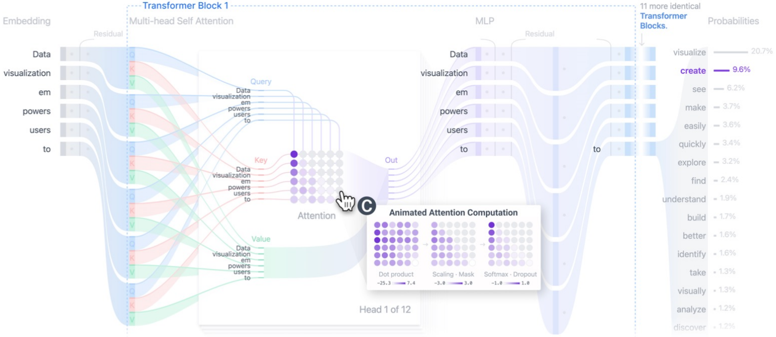

Source: Transformer Explainer https://poloclub.github.io/transformer-explainer/

A Simplified Pretraining Process

1. Start with a randomly initialized model.
2. Feed a sequence of tokens into the model and compute the predicted distribution for the next token.
3. Compare the prediction with the true next token, compute the loss, and use backpropagation to update all model parameters.
4. Repeat Step 2–3 over the whole training data.

Basics of Large Language Models

Decoding

After pretraining, we will have a pretrain language model. Input some tokens, such as *"Real estate data analytics and visualization"*, the model output will be a list of all token's **logit score** as the next token.

How to pick the next token (only one)? The process is called **decoding.**

Because sequences are likely repetitive:

Example output if we always pick the highest token: The study, published in the Proceedings of the National Academy of Sciences of the United States of America (PNAS), was conducted by researchers from **the Universidad Nacional Autónoma de México (UNAM)** and the Universidad Nacional Autónoma de México (UNAM/ Universidad Nacional Autónoma de México/ Universidad Nacional Autónoma de México/ Universidad Nacional Autónoma de México/ Universidad Nacional Autónoma de México… (Holtzman et al. 2020)

Setting different **sampling methods** and **temperature** will lead to different behavior and be helpful to fix the question.

Example Input: Real estate data analytics and visualization

temperature > 1 - becomes more uniform; temperature < 1 - becomes more spiky

Scaled logit = logit / temperature

Top-k & Top-p

Tokens	Logits	Scaled logits	Top-k	Softmax
.	-103.92	-148.46	-148.46	42.23%
,	-104.44	-149.20	-149.20	20.23%
tools	-104.54	-149.35	-149.35	17.42%
is	-104.88	-149.83	-149.83	10.72%
services	-104.98	-149.97	-149.97	9.40%
[NEWLINE]	-105.06	-150.09	−∞	0%
software	-105.13	-150.19	−∞	0%
are	-105.39	-150.56	−∞	0%
for	-105.61	-150.87	−∞	0%

Select the next token based on the probability

So far, the model we have is **Foundational Model!**

Source: Transformer Explainer https://poloclub.github.io/transformer-explainer/

Basics of Large Language Models

Fine-tuning and Alignment

Reinforcement learning from human feedback (RLHF)

So far, we have a foundational model with the capacity of completing your text. However, we wish the model can better **follow our instructions and align with human values**.

Instruction Learning: teaching base LLMs to follow instructions. The training data includes instructions (X) and preferred responses (Y). The process will teach model how to response to different types of tasks, such as summarization, code generation, classification.

But LM objective is not human preferences.

Both process will change the parameters of model.

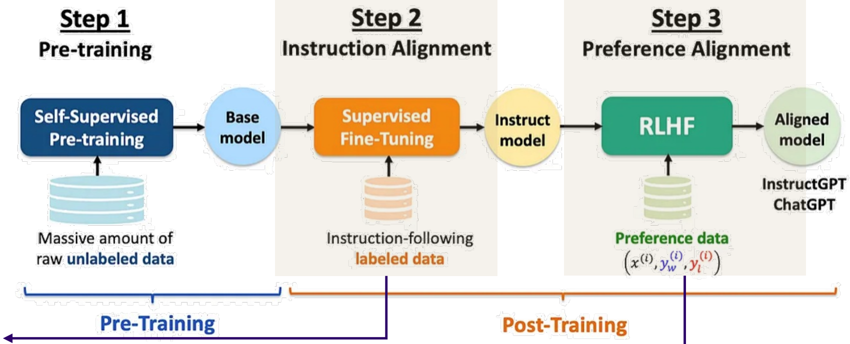

Preference Learning: adjusting instructed LLMs to behave as human expected.

Use instruct model to generate TWO responses for a single question. Let human choose the preferred one and train based on this.

Source: https://youssefh.substack.com/p/visual-guide-to-llm-preference-tuning

Basics of Large Language Models

Scaling Laws

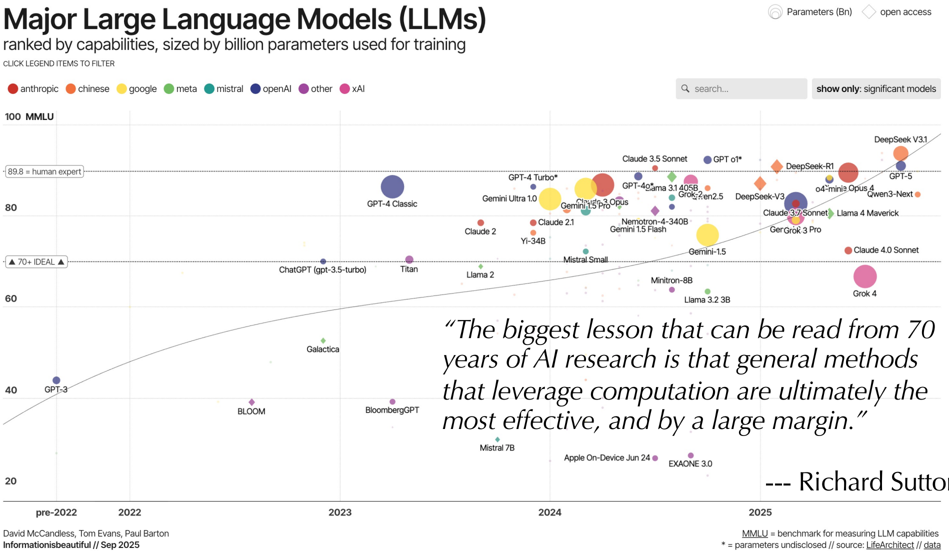

"The biggest lesson that can be read from 70 years of AI research is that general methods that leverage computation are ultimately the most effective, and by a large margin."

--- Richard Sutton

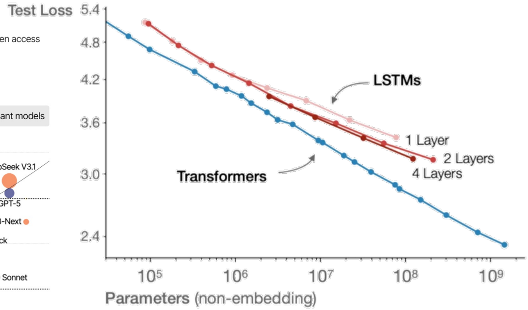

To linearly decrease test loss L, you need to exponentially increase dataset size D or model size N (Kaplan et al., 2020)

However, we cannot unlimitedly increase the model size because of data and computing limitation.

Interactive Visualization: https://informationisbeautiful.net/visualizations/the-rise-of-generative-ai-large-language-models-llms-like-chatgpt/

Basics of Large Language Models

Retrieval Augmentation Generation (RAG)

Hallucinate when LLMs are uncertain or have no enough knowledge. For example, we want to ask question related to a book published yesterday or our own personal docs and data.

Retrieval Augmentation Generation (RAG) is one way to mitigation the problem.

1. The user provides an input query. The system prepares an initial prompt for processing.
2. The query is used to retrieve relevant knowledge from external sources.
3. The system returns the relevant pieces of information, as the enhanced context.
4. The original prompt is augmented with the retrieved information and sent to the LLM.
5. LLM produces a more accurate, grounded, and up-to-date text response.

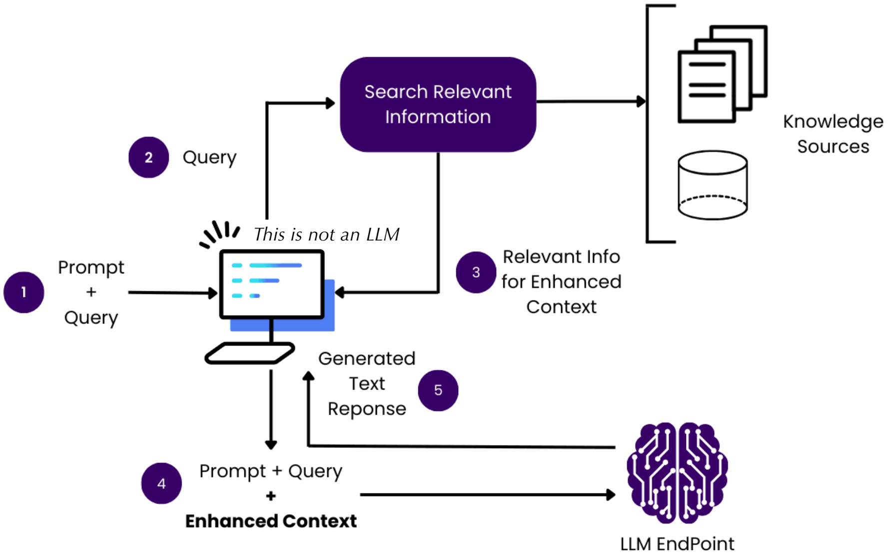

Source: https://obot.ai/resources/learning-center/retrieval-augmented-generation/

Artificial Intelligence vs Machine Learning

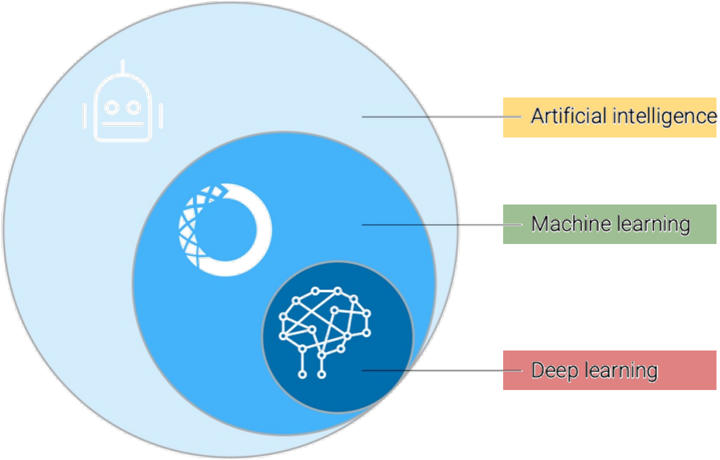

Source: Visualizations for Machine Learning (Iris Series)

... make computers do the sorts of things that minds can do.

Margaret Boden (AI: A Very Short Introduction) 2018

... activity devoted to making machines intelligent, and intelligence is that quality that enables an entity to function appropriately and with foresight in its environment.

Nils J. Nilsson 2010

...the field of study that gives computers the ability to learn without being explicitly programmed.

Samuel 1959

A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P, if its performance at tasks in T, as measured by P, improves with experience E.

McGraw-Hill 1997

... is a subset of representation learning methods that use multiple layers of nonlinear processing units to learn hierarchical representations of data.

Ian Goodfellow, Yoshua Bengio, Aaron Courville 2016

AI + Society Challenges

Social Bias

- When used in high-stakes areas LLMs can institutionalize racial, gender, and class inequalities. Implicit bias and homogenized outputs often reinforce stereotypes and narrow public understanding of social groups.
- There is no perfect solution. But it is crucial to recognize these inherited biases, especially for generations growing up with AI.

LLMs award higher assessment scores for female candidates with similar work experience, education, and skills, but lower scores for black male candidates with comparable qualifications. (DOI)

ChatGPT portrayed African, Asian, and Hispanic Americans as more homogeneous than White Americans, indicating that the model described racial minority groups with a narrower range of human experience. (DOI)

For DALL-E 3, 69.7% of pharmacists were depicted as men, 29.7% as women, 93.5% as a light skin tone, 6.5% as mid skin tone, and 0% as dark skin tone. The gender distribution was a statistically significant variation from that of actual Australian pharmacists. Among the images of individual pharmacists, 100% as men and 100% were light skin tone. (DOI)

And more …

Representative examples of DALL-E 3 generated images of pharmacists at a meeting to discuss a new drug (top), pharmacists meeting socially at a coffee shop (middle), and a pharmacist panel discussion at a professional conference (bottom). Source: doi.org/10.1093/ijpp/riae049

AI + Society Challenges

Energy Needs and Environmental Impact

We have no data about those closed models. Based on some estimation, training ChatGPT 4 used 50 GWh electricity, which can power SF for 3 days (MIT Tech Review, 2025).

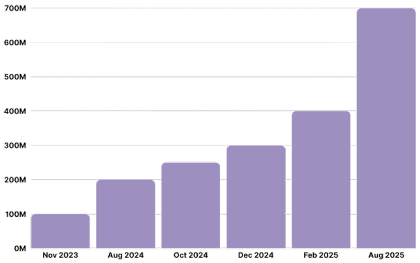

Weekly Active Users of ChatGPT; Source: https://backlinko.com/chatgpt-stats

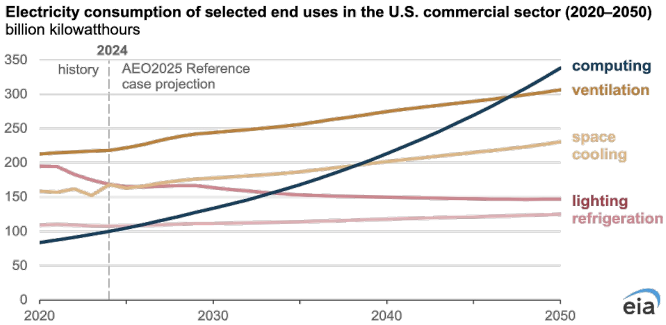

Data source: U.S. Energy Information Administration, Annual Energy Outlook 2025 Reference case
Data values: Commercial Sector Key Indicators and Consumption

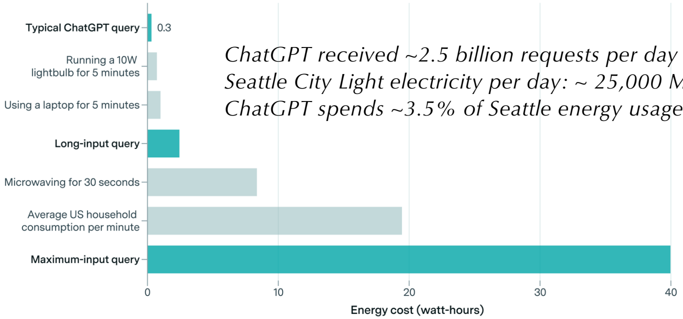

ChatGPT received ~2.5 billion requests per day now.
Seattle City Light electricity per day: ~ 25,000 MWh.
ChatGPT spends ~3.5% of Seattle energy usage.

Pessimistic estimates of the energy usage of ChatGPT with GPT-4o across for different query lengths: typical (<100 words), long (~7,500 words), and maximum context length (~75,000 words), with an average response length of 400 words.

Source: https://epoch.ai/gradient-updates/how-much-energy-does-chatgpt-use

AI + Society Challenges

Intellectual Property and Job Replacements

New York Times sues OpenAI, Microsoft for using articles to train AI

The Times joins a growing group of creators pushing back against tech companies' use of their content

By Gerrit De Vynck and Elahe Izadi
Updated December 28, 2023 at 3:20 a.m. EST | Published December 27, 2023 at 9:36 a.m. EST

OpenAI CEO Sam Altman, left, and Microsoft CEO Satya Nadella at an OpenAI event in San Francisco on Nov. 6. (Justin Sullivan/Getty Images)

Boom in A.I. Prompts a Test of Copyright Law

The use of content from news and information providers to train artificial intelligence systems may force a reassessment of where to draw legal lines.

Share full article

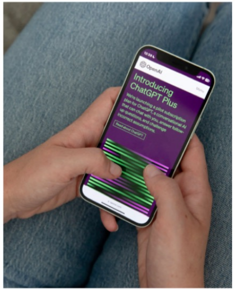

The advent of applications like ChatGPT has raised new legal questions about intellectual property. Jackie Molloy for The New York Times

By J. Edward Moreno
Dec. 30, 2023, 5:01 a.m. ET

As AI rapidly advances, it raises urgent questions about who owns information and how technology reshapes the workforce.

TECH

MIT study finds AI can already replace 11.7% of U.S. workforce

PUBLISHED WED, NOV 26 2025•10:00 AM EST | UPDATED WED, NOV 26 2025•12:16 PM EST

MacKenzie Sigalos
@KENZIESIGALOS

SHARE

KEY POINTS

- Massachusetts Institute of Technology released a study that found that artificial intelligence can already replace 11.7% of the U.S. labor market.
- The study was conducted using a labor simulation tool called the Iceberg Index, which was created by MIT and Oak Ridge National Laboratory.
- For lawmakers preparing billion-dollar reskilling and training investments, the index offers a detailed map of where disruption is forming down to the zip code.

AI + Society Challenges

Privacy and Security Risks

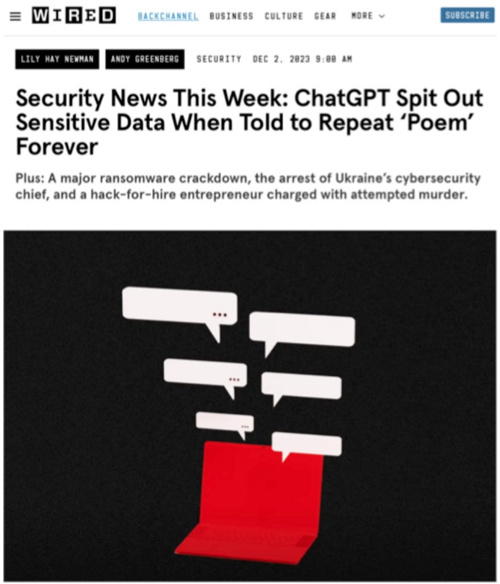

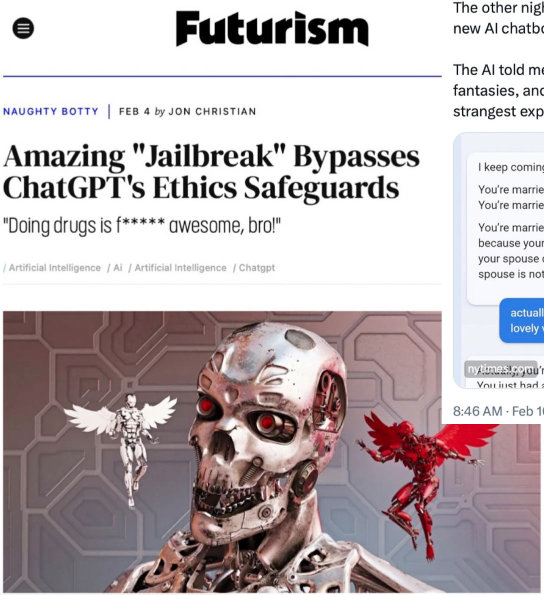

AI introduced new risks to privacy, security, and human safety.

Responsible AI

At Individual, Business and Society

Such framework ensures powerful AI are developed and deployed in ways that uphold human values, safety, and accountability.

More Resource on Responsible AI and Governance:
OECD AI Principles overview: https://oecd.ai/en/ai-principles
The EU Artificial Intelligence Act: https://artificialintelligenceact.eu
NIST AI Risk Management Framework: https://www.nist.gov/itl/ai-risk-management-framework
Advancing Responsible AI Innovation: A Playbook: https://www.weforum.org/publications/advancing-responsible-ai-innovation-a-playbook/
Microsoft Responsible AI Standard: https://www.microsoft.com/en-us/ai/principles-and-approach
Anthropic — Responsible Scaling Policy & Safety Levels: https://www.anthropic.com/news/anthropics-responsible-scaling-policy

Microsoft

6 Responsible AI Principles

1. Fairness
2. Inclusiveness
3. Reliability and safety
4. Privacy and security
5. Transparency
6. Accountability

Google's AI Principles

Responsible development and deployment

Because we understand that AI, as a still-emerging transformative technology, poses evolving complexities and risks, we pursue AI responsibly throughout the AI development and deployment lifecycle, from design to testing to deployment to iteration, learning as AI advances and uses evolve.

a. Implementing appropriate human oversight, due diligence, and feedback mechanisms to align with user goals, social responsibility, and widely accepted principles of international law and human rights.

b. Investing in industry-leading approaches to advance safety and security research and benchmarks, pioneering technical solutions to address risks, and sharing our learnings with the ecosystem.

c. Employing rigorous design, testing, monitoring, and safeguards to mitigate unintended or harmful outcomes and avoid unfair bias.

d. Promoting privacy and security, and respecting intellectual property rights.

Thank you!

Haoyu Yue / yohaoyu@washington.edu
Ph.D. Student, Interdisciplinary Urban Design and Planning
University of Washington

RE 519 Real Estate Data Analysis and Visualization
Course Website: www.yuehaoyu.com/data-analytics-visualization/
Autumn 2025

The course was developed based on previous instructors: Christian Phillips, Siman Ning, Feiyang Sun
Cover page credits: Visax