From Perceptron to AI: The Rise of Large Language Models

Lecture 4 of 4:
How can we use AI?

Dr. Michael Stobb
Thursday Forum
Coe College

February 2026

Last Time…

- Discussed “What are some of the Impacts of AI?”
  - Hallucinations
  - Synthetic Media Generation
  - AI Deepfakes
  - AI in Education

- Ran out of time to discuss:
  - Uses of AI in Biology
  - AI in Employment
  - Robotics
  - AI Consciousness

Today's Goals

- Should we even use AI at all?
  - What are the ethical considerations of AI?
  - Are there valid reasons to *not* use it?

- If AI use *can* be ethical, how should we use it?
  - Modern AI best practices
  - Do's and Don'ts
  - Hands on Examples

Energy Use in AI

Standard US home uses
1 GWh in about
100 years

- Little known fact:

Data centers use lots of Electricity!

- This alone should not worry us
  - Lots of things use lots of electricity!
    - AC in US homes use over 250,000 GWh annually
    - Chemical manufacturing uses 150,000 GWh
    - Food production uses 80,000 GWh
  - US Data Centers use ~200,000 GWh annually
    - Not just AI related
    - In 2015, US Data Centers used ~70,000 GWh

Energy Use in AI

- Better if contextualized:

How much energy per usage?

- GWh too big; need a better scale in watt hours:
  - Toasting bread: ~50 Wh
  - K-Cup Coffee maker: ~100 Wh
  - 1 minute in microwave: ~25 Wh
  - Fully charge cell phone: ~20 Wh
  - Streaming HD Netflix 1 hour: ~100 Wh

- AI queries in 2024 estimated at **~3 Wh each**
- August 2025 AI queries estimated at **~0.3 Wh each**

Energy Use in AI

- Not everything is “per query”

Inference (running)

- Model is used by millions of people
- Each query costs less than 1 Wh
- Billions of daily queries

~200 GWh annually
Per model

Training (building)

- Model has to be built
- One time cost, done up front
- Millions of dollars of compute

~10-100 GWh Once
Per model

Energy Use in AI

- As AI models get more efficient, will total energy spent on AI go down?

Probably Not

- Jevons Paradox:

> *Increased efficiency can lead to an **increase** in total consumption rather than a decrease*

  - Adding highway lanes leads to MORE traffic
  - Water saving appliances lead to MORE water usage
  - More fuel efficient cars are driven MORE overall

Energy Use in AI

- What about all the ***new*** data centers?
  - AI demand is doubling about every 6 months
  - The big companies are preparing for future
    - Estimated over $500 Billion in 2026

  - Iowa is a good place for Data Centers:
    - Centralized
    - Stable weather (no earthquakes, hurricanes)
    - Cold
    - Lots of wind power

Water Use in AI

- Computers get hot when they are used
- If they get too hot, the computer can break
- The computers are cooled lots of different ways
  - Air cooled (fans), direct water cooling, closed loop, open loop, etc.
- The buildings storing the computers need cooling
  - Potentially very large HVAC systems

Water is used (directly/indirectly) for cooling

Water Use in AI

- Cooling accounts for ~30-40% of energy use in classical data centers
- Modern (to-be-built) data centers are much more efficient: only ~10% of energy is for cooling
  - But they do use more water!
- Water usage scales *roughly* with **time** computing
  - Longer compute times need more water
  - If you have to “wait” a long time, you are using more water
  - This is true for *anything* on the internet
- Current AI models use ~1mL of water per query

Water Use in AI

- Water Use Comparisons:
  - All Google Data Centers globally for 2023:
    - **6.1** billion gallons!
  - US Golf Courses:
    - over 500 billion gallons annually
  - Quarter-Pound hamburger patty:
    - ~400 gallons
      - ~50 billion hamburgers eaten in US annually
      - That's *20 trillion gallons* annually!

Environmental Ethics of AI

- My personal opinion:

The energy/water argument is mostly a ***distraction***

- Not that it doesn’t have value!
- Just probably not as significant as *other* issues

AI & Copyright

- As we have already discussed:
  - AI models are trained on essentially ALL DATA
  - Copyrighted or not
  - Mostly *without* permission from owners

- All AI companies are involved in lawsuits
  - Anthropic already settled for $1.5 billion

- Lawsuits show evidence of training by retrieving verbatim text out of AI models
  - But this is not the big problem!
  - Often had to do *crazy* things to get output

AI & Copyright

- Copyright owners have sole right to “reproduce” the work
- In training their models, AI companies had to make an “intermediate copy” of the work
  - They had to essentially *download* it to their servers
  - This was done likely without permission or payment
  - This is the “Original Sin” of *Ingestion!*
- It doesn’t matter what the model produces or does not produce: it was “illegal” from the beginning
- Since the model was built with “illegally” obtained text, the entire model should be deemed “illegal”

AI & Copyright

- If you go to the library, you can read every mystery book
- If you then write a great mystery novel, is this ***theft***?
  - We would say the person “learned” how to write
  - AI companies are claiming their work is the same!
    - They are selling the intelligence *learned*
  - Copyright holders say that the scale and future competition violate Fair Use claims

---

- Current laws are murky at best regarding this
- US Supreme Court needs to weigh in on it!
  - Both sides have reasonable arguments

AI Alignment

- Who wins in a fist fight? Gorilla vs. Human
  - The Gorilla!
    - But then why do humans dominate the earth?
- This same idea can be applied to Humans and AI
  - What happens if we build a ***Super AI***?

The ***AI Alignment problem*** refers to the challenge of ensuring an advanced AI system reliably acts in accordance with human goals, values, and ethical principles.

- Turns out these are ***subjective***!
  - Specifying what these goals and values are is hard

AI Alignment

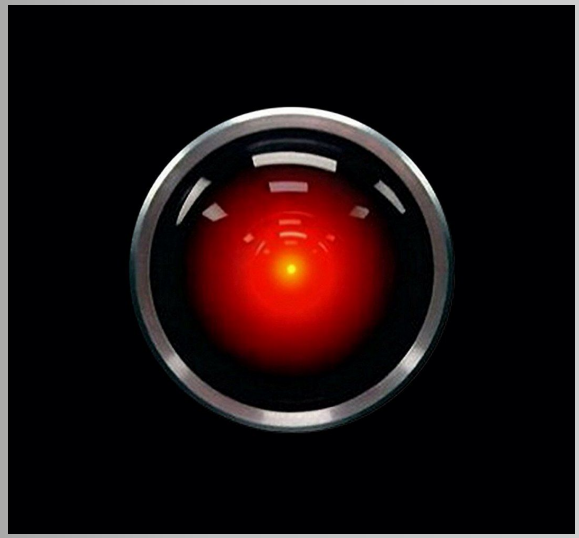

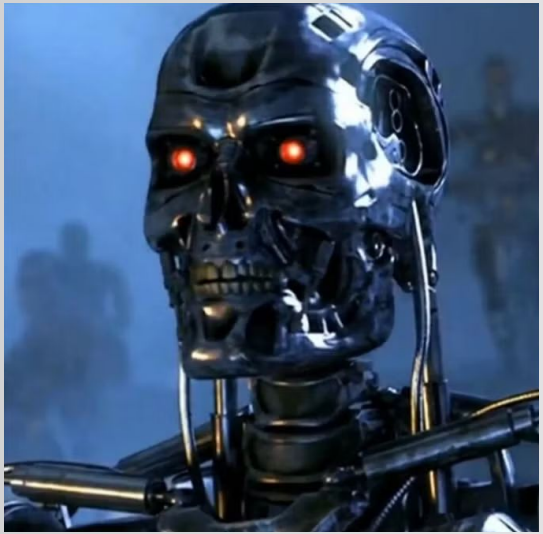

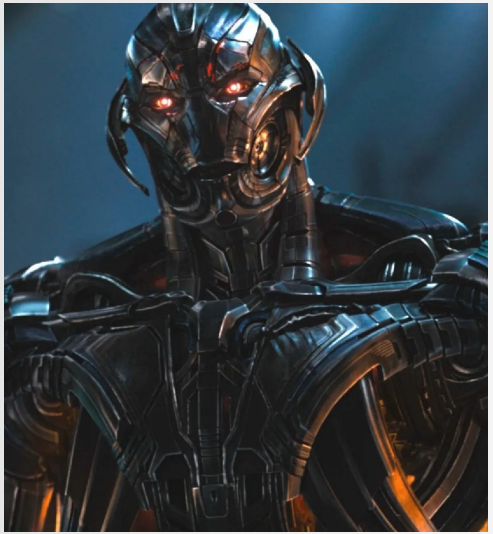

- A “rogue” AI would be bad
- But an *anti*-aligned AI is not required for an apocalypse

AI Alignment

- AI models are trained to optimize a given goal
  - Predicting the *right* word or sounding the *nicest*

Two Reasonable Ideas

- ***Idea 1***: Intelligence and Morality are **independent**
  - Being smart does not mean you are good *or* evil

- ***Idea 2***: All goals have **subgoals** that make up that task
  - Some of these subgoals are very likely to be shared:
    - Especially *self-preservation* & *resource acquisition*

  - Often called “Instrumental Convergence”

AI Alignment

- Imagine we have a super AI that wants to create as many ***paperclips*** as possible
  - This goal is simple and non-threatening to humans
- To maximize paperclips, it must not get turned off
  - This is a shared subgoal for all tasks
- The AI is super smart, but it could be totally amoral
  - These are independent!
- To keep from being turned off, AI could decide that humans need to be *eliminated* so they can't shut it down

AI Alignment

- While this seems silly, its based on solid premises
  - Has been hard to refute
- Many professional AI researchers are *very* concerned

> This is a scary thought experiment!
> But remember, we are **not there yet**.

- Current models are not sufficient for apocalypse
- But these models *can* still be used for good!
  - Finding cures to diseases, making tasks easier, etc.

How to Use Modern AI?

- Assuming that that AI can be ethically used. . .

> What are some good ways to use Modern AI?

- Disclaimer:
  - Things are changing *fast*
  - What works for one may not work for all
  - We will definitely not cover everything
  - We will mainly be showcasing Google Gemini models
    - Other models can be just as good!
    - Though often the “best” models require payment

Text Prompts

- We start with a basic request:

> “Please help me plan a family trip to Des Moines Iowa.”

- This works!
  - But it’s bland and not specialized to me at all
- The model needs *context* to make a better plan
- We could supply that context ourselves

OR

> We could ask the model to *ASK US!*

Reverse Prompting

- We add to our original prompt:

> “Please help me plan a family trip to Des Moines Iowa.
> Before starting anything,
> please ask me any necessary questions you may have.”

- This immediately sets us up for a *much* better response!
- It does take much longer, but the results are worth it
- Can be done for several rounds
- Can be helpful to ask for a summary of answers

Meta Prompting

- Sometimes we don’t know how to even get started
  - Task is complex or the topic is unfamiliar
  - We maybe need to use a different model than a chat

- We can ask the model to write a prompt for *another* model:

> “Please help me write a prompt for a Gen AI model.
> I want to create an image of a wooden chipmunk.
> Please ask me any necessary questions you have.”

- This will give us some ideas, but probably needs context
- Can use our “Reverse Prompting” technique again!

Fun Things

- There are *many* real-world uses for text based AI
  - Analyzing a data set, writing code for a project, etc.

- We can also use AI models to generate
  - Images
  - Videos
  - Songs

- These are mostly for *entertainment value* right now
  - It’s impressive, but has limited use cases
  - This will likely change as models get better/faster

Agentic AI

- This is a new style of AI model ***use***
  - These are the same models but set up for different use

> **Agentic AI** systems are designed to operate fully *autonomously*, making *decisions* and *taking actions* to achieve specific goals for the user.

- No longer have a chat to do a task (with back and forths)

- Instead launch a small swarm of AI “bots” that do task
  - These will work on separate subtasks (i.e., files)
  - The bots will combine their efforts in the end

- Now the user just has to “manage” them

The End

Thank you!

For notes, further readings, and a full copy of the slides, just scan the QR code:

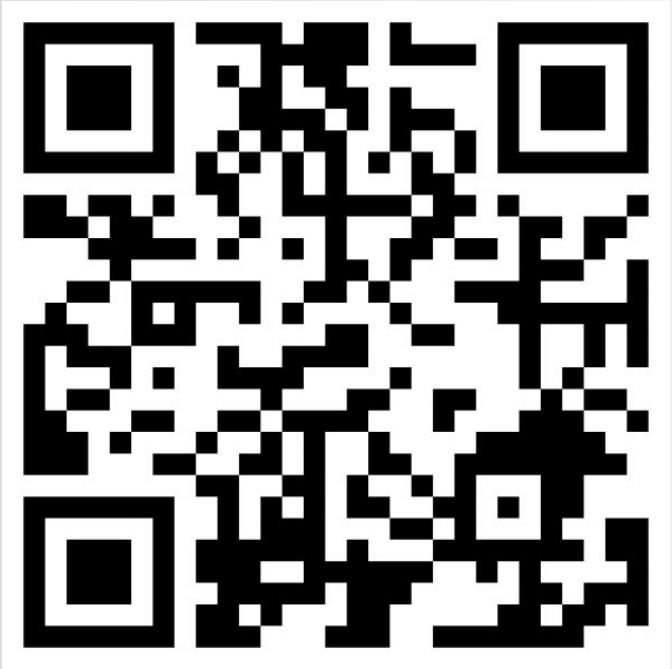

stobb.org/thursday_forum

Middle Break

For notes, further readings, and a full copy of the slides, just scan the QR code:

stobb.org/thursday_forum