

- **Generative AI and LLMs:
Background and context**

0. Today's objectives

After today's class, you should be able to:

- Give a general definition for *AI / artificial intelligence*
- List the categories of automation that are often covered under the general term "AI"
- Give a general explanation of a large-language model (LLM) as a text predictor
- Identify some of the potential problems associated with AI, along with potential benefits

1. What do we mean when we say “AI”?

Discussion

- “AI” is an initialism for artificial intelligence
- What do we mean when we say “AI”?
 - Come up with at least **3 different tools, apps, or devices** that you would say count as AI
 - What do they have in **common**?
 - Can you state any **general criteria** for what makes something be a case of AI?

1. What do we mean when we say “AI”?

- OED: The capacity of computers or other machines to exhibit or simulate intelligent behaviour; the field of study concerned with this.

In later use also: **software used to perform tasks or produce output previously thought to require human intelligence**, esp. by using **machine learning** to extrapolate from **large collections of data**.

1. What do we mean when we say “AI”?

- Wikipedia: ...the capability of computational systems to perform tasks typically associated with human intelligence, such as learning, reasoning, problem-solving, perception, and decision-making.

It is a field of research in computer science that develops and studies **methods and software** that enable **machines** to **perceive** their environment and use **learning and intelligence** to take actions that maximize their chances of **achieving defined goals**.

1. What do we mean when we say “AI”?

- Origin of the term (OED):

1955 A proposal for the Dartmouth summer research project on artificial intelligence.

[John] McCarthy et al. (typescript) 31 August (title)

- Goal: To bring together people from math, computer science, physics, and engineering to develop computerized automation
- Much of the motivation was to develop military technology for the Cold War

[*] Different from the linguist [John McCarthy](#)!

1. What do we mean when we say “AI”?

- Bender & Hanna (2025: Ch 1) take a critical view:

“To put it bluntly, AI is a marketing term. It doesn’t refer to a coherent set of technologies.

Instead, the phrase ‘artificial intelligence’ is deployed when the people building or selling a particular set of technologies will profit from getting others to believe that their technology is similar to humans,

able to do things that, in fact, intrinsically require human judgment, perception, or creativity.”

Bender, Emily M., and Alex Hanna (2025). *The AI con: How to fight big tech’s hype and create the future we want*. New York: Harper Collins.

1. What do we mean when we say “AI”?

- Bender & Hanna (2025: Ch 1):

“...the conversation becomes clearer if one speaks in terms of “automation” rather than “AI” and looks at precisely what is being automated.

1. What do we mean when we say “AI”?

- Some older types of “AI” in this sense
 - **Recommendation** — “you may also like”, “people who X also Y”, social-media feeds
 - **Transcription/translation** — text-to-speech, recognizing text in images, machine translation, turning your selfie into Rumi or Mario
 - **Classification** — image classification (personal image tagging, large-scale surveillance); customer classification for targeted advertising
 - **Decision making** — setting bail, screening resumes, allocating social benefits

1. What do we mean when we say “AI”?

- **Generative AI** is another category
- Bender & Hanna (2025: Ch 1):

“so-called generative AI or, more aptly, synthetic media machines.”

“These are systems like ChaptGPT, Gemini, or DALL-E that allow users to generate images or plausible-sounding text based on textual prompts.”

“A prompt, in generative AI terminology, is the words used to describe the desired output.”

2. How do y'all use generative AI?

- Results from prep-questions survey 10.29
[by the way, please ignore any “grading” on this survey; it’s a prep-question “pass” for completion]
 - What do you use? What is it good for / useful for?

3. Large language models

Sources for this discussion:

- Bender & Hanna (2025)
- Linzen & Baroni (2021)

Linzen, Tal, and Marco Baroni (2021). Syntactic structure from deep learning. *Annual Review of Linguistics* 7(1): 195–212. [[UNC link](#)]

3. Large language models

- Important LLM precursor: the ***n*-gram model** (1940s)
 - Does anyone know what these are?

3. Large language models

- Important LLM precursor: the ***n*-gram model** (1940s)
 - n = a **number** (1, 2, 3...)
gram means (written) **word**
 - Takes a **corpus** (collection of text) as its input
 - Produces a list of all **word strings of length n** that occur in the corpus, with their **frequencies**
 - Can be used for **prediction**:
 - What word occurs most often after *alleged*?
 - What word is this MS Word user trying to type?
- [Google Books Ngram Viewer](#): n -gram frequencies in the Google Books corpus (no predictions)

3. Large language models

- n -gram models have trouble making predictions about low-frequency (unlikely) words?
- Next step: the **neural net(work) model**
 - Consist of units called “perceptrons”
 - Each takes multiple inputs, determines an output
 - The weights (importance) each perceptron puts on inputs from other perceptrons get adjusted based on feedback

B&H (2025: Ch 2): term *neural net* was “very loosely inspired by a 1940s understanding of how neurons work in the human brain...leaving...the impression that these systems are analogous to biological brains”

3. Large language models

- Neural-net model can take information about
 - a sequence of words
 - pixels in an image
 - ...

and output probabilities of likely **labels** for the output (=classification, image identification, machine translation, ...)

- One type of output “label” = the **most likely next word** in a sequence

3. Large language models

- In principle, a neural-net model needs to be **trained** so that its perceptron weights can be adjusted for better performance
 - Image identification results are trained by data from human image identification results
 - But, collecting the human data is a big task
- B&H (2025: Ch 2)

“In the case of language modeling, the correct answer of which word came next is just whatever word happened to come next in the training corpus.

This means that it's enough to collect a bunch of text...”

3. Large language models

- What a generative language model does:
 - Pick a likely word (given some input, like the user's prompt)
 - Then pick a likely next word, given the previous word picked and the rest of the context

3. Large language models

Discussion

- Potential for problems here?

3. Large language models

Discussion

- Potential for problems here?
 - How the LLM behaves will depend on what the training corpus is like
 - Biases?
 - A role for human “correction” as a predictor
 - The LLM has no conception of “true” or “accurate,” only “a plausible next word”
 - Hallucinations?

4. GenAI and careers

- Results from prep-questions survey 10.29
- Potential for...
 - Making jobs more interesting
 - Making jobs more boring
 - Changing job opportunities, career paths

4. Y'all's concerns about genAI

- Results from prep-questions survey 10.29

	A real concern	Media hype	Haven't heard
Climate impact	20	5	6
Copyright ethics	25	5	1
Labor ethics	11	1	19
Confidentiality risk	26	3	2
Risk for humanity	11	18	2

5. For next time

- Myth: “GenAI can learn to think like a human”
- Big-picture research question:
Can an LLM learn to make the same **grammaticality judgments** about language structures as a human?

Implications:

- If yes, maybe human language acquisition also needs nothing more than lots and lots and lots of data and an ability to make predictions (acquisition is not sensitive to linguistic structure)