CS 6501 Natural Language Processing (Spring 2024)

Yu Meng
University of Virginia
yumeng5@virginia.edu

Jan 17, 2024

Course Information & Logistics

- Instructor: **Yu Meng** (yumeng5@virginia.edu)
- TAs: **Afsara Benazir** (hys4qm@virginia.edu) **Zhepei Wei** (tqf5qb@virginia.edu)
- Time: Mondays & Wednesdays 3:30pm - 4:45pm
- Location: Mechanical & Aerospace Engineering Building 339
- Office Hour: On demand (regular OH will be set up if enough requests)
- We'll use Piazza (accessible via Canvas) to answer logistics/technical questions

Course Information & Logistics

- This course is designed to be a **research-oriented graduate-level** course
- A comprehensive overview of cutting-edge developments in NLP
- Prerequisites: CS 4774 (having deep learning background is important!)
- This course may benefit you if
  - You are working on NLP research (PhD/MS research students)
  - Your research uses NLP models/tools
  - You aim for a job that involves using NLP models/tools
  - You are extremely interested in the topic and willing to spend much time to learn
- Make your own judgement if this course fits your ability/need
  - We have many undergraduates now – this course may not be suitable for you!
  - Please drop asap if it doesn't fit your plan to take to give seats to students in the waitlist
- If you are in the waitlist
  - Attend the lecture and complete the course requirements as if you are registered
  - Keep an eye out for open seats when someone drops

Course Format & Grading

- Course Website: https://yumeng5.github.io/teaching/2024-spring-cs6501
- **Paper Presentation (30%)**
  - Starting from the 3rd lecture, each lecture will be presented by a group of 2 or 3 students
  - Every group presents one lecture (4 papers) during the semester

1/24	**Large Language Models and In-Context Learning**	* Language Models are Few-Shot Learners (GPT-3) * Llama 2: Open Foundation and Fine-Tuned Chat Models * An Explanation of In-context Learning as Implicit Bayesian Inference * Rethinking the Role of Demonstrations: What Makes In-Context Learning Work?	
1/29	**Model Calibration**	* How Can We Know When Language Models Know? On the Calibration of Language Models for Question Answering * Surface Form Competition: Why the Highest Probability Answer Isn't Always Right * Teaching Models to Express Their Uncertainty in Words * Semantic Uncertainty: Linguistic Invariances for Uncertainty Estimation in Natural Language Generation	
1/31	**Scaling and Emergent Ability**	* Training Compute-Optimal Large Language Models * Scaling Data-Constrained Language Models * Emergent Abilities of Large Language Models * Are Emergent Abilities of Large Language Models a Mirage?	

Supplemental Reading
* (Blog) The Illustrated Transformer * (Blog) Transformer Inference Arithmetic

Course Format & Grading

- Course Website: https://yumeng5.github.io/teaching/2024-spring-cs6501
- **Paper Presentation (30%)**
  - Starting from the 3rd lecture, each lecture will be presented by a group of 2 or 3 students
  - Every group presents one lecture (4 papers) during the semester
  - We'll send an email to everyone when the signup sheet is open (later today)
  - You can sign up for the topic you are interested in – slots are first come, first served!
  - **Presentation duration**: strictly limited to 60 minutes, followed by a 10-minute question-and-answer session with the audience
  - **Deadline**: Email your slides to the instructor and TAs at least 48 hours before your presentation (e.g., if presenting on Monday, slides should be emailed by Saturday 3:30 pm)
  - You will receive feedback from the instructor to improve your slides (if necessary, the instructor may schedule a meeting with your team to go over the slides)
  - Late submissions result in a 50% presentation grade deduction
  - Detailed grading rubrics and tips will be shared later
  - First three student lectures automatically receive 5%, 3%, 1% extra credit of final grade

Course Format & Grading

- Course Website: https://yumeng5.github.io/teaching/2024-spring-cs6501
- **Participation (20%):**
  - Starting from the 3rd lecture, everyone is required to complete two mini-assignments
  - **Pre-lecture question**: read the 4 papers to be introduced in the lecture, and submit a question you have when you read them
  - **Post-lecture feedback**: provide feedback to the presenters after the lecture
  - We'll use Google Forms to collect pre-lecture questions and post-lecture feedback and share them with the presenters
  - **Deadlines**: pre-lecture questions are due one day before the lecture (e.g., For Monday lectures, you need to submit the question by Sunday 11:59 pm); post-lecture feedback is due each Friday (both Monday & Wednesday feedback is due Friday 11:59 pm)
  - Lectures are not recorded, but slides will be posted on the course website
  - Guest lectures do not have pre-lecture questions/post-lecture feedback, and we'll directly take attendance

Course Format & Grading

- Course Website: https://yumeng5.github.io/teaching/2024-spring-cs6501
- **Project (50%):**
  - Complete a research project, present your results, and submit a project report
  - Work in a team of 2 or 3 (any deviation from this size requires prior approval from the instructor) – may or may not be the same team as your presentation group
  - (Type 1) A comprehensive survey report: carefully examine and summarize existing literature on a topic covered in this course; provide detailed and insightful discussions on the unresolved issues, challenges, and potential future opportunities within the chosen topic
  - (Type 2) A hands-on project: not constrained to the course topics but must be centered around NLP; doesn't have to involve large language models (e.g., train or analyze smaller-scale language models for specific tasks); eligible for extra credits if publishable
  - **Project proposal**: 5%; **Mid-term report**: 10%; **Final presentation and report**: 35% (deadlines announced later)

What is Natural Language Processing (NLP)?

- An interdisciplinary subfield of machine learning and linguistics
- Goal: Enable computers to understand, interpret, and generate human language

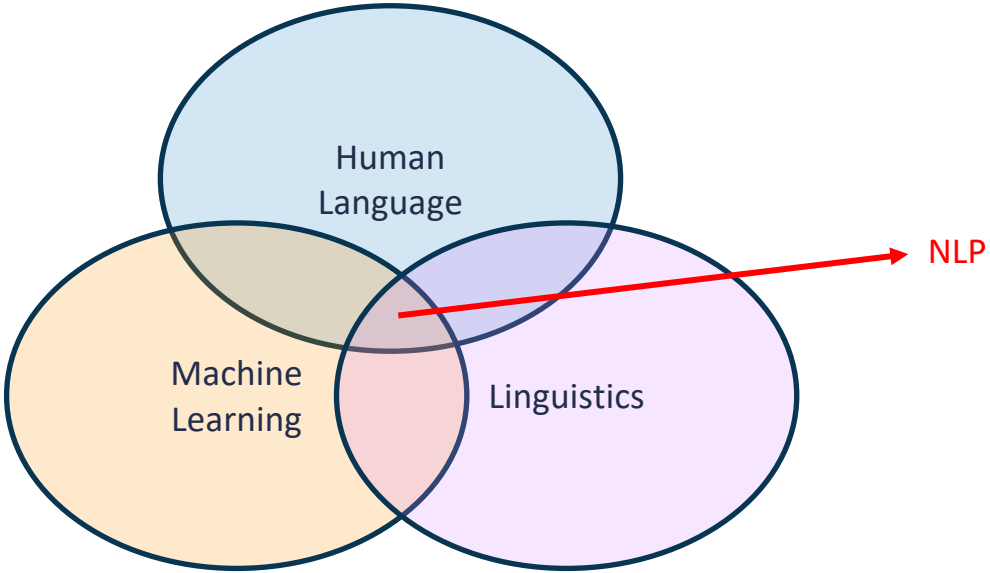

The History of NLP

Linguistic-rule based methods
(e.g., syntactic pattern matching)

Before 1980s

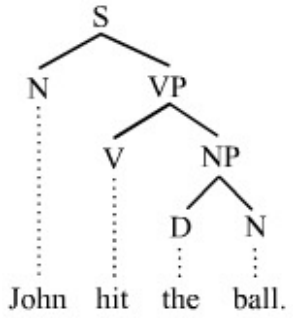

Constituency-based
parse trees

The History of NLP

Statistical methods
(e.g., n-gram models, hidden state models)

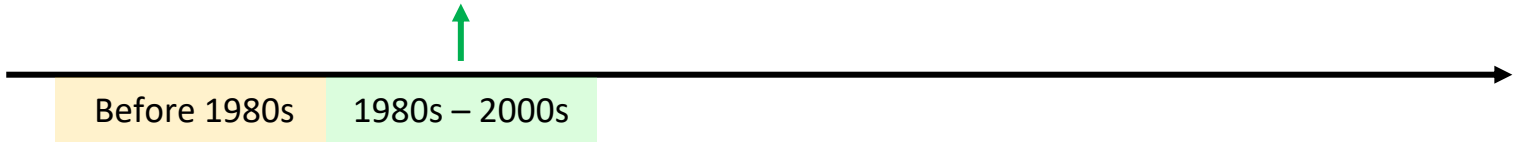

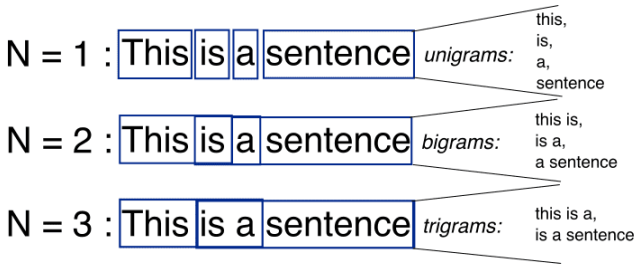

N-gram models

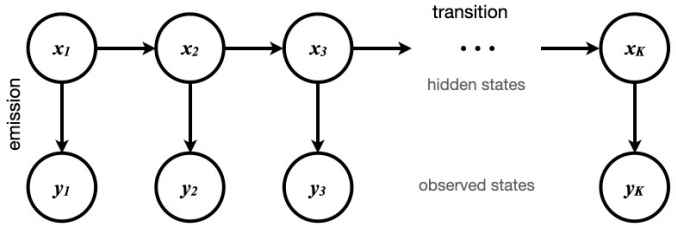

Hidden state models

The History of NLP

(Simple) neural-network-based methods
(e.g., word embeddings, convolutional/recurrent neural networks)

Before 1980s | 1980s – 2000s | 2000s – 2018

Convolutional neural networks (CNNs)

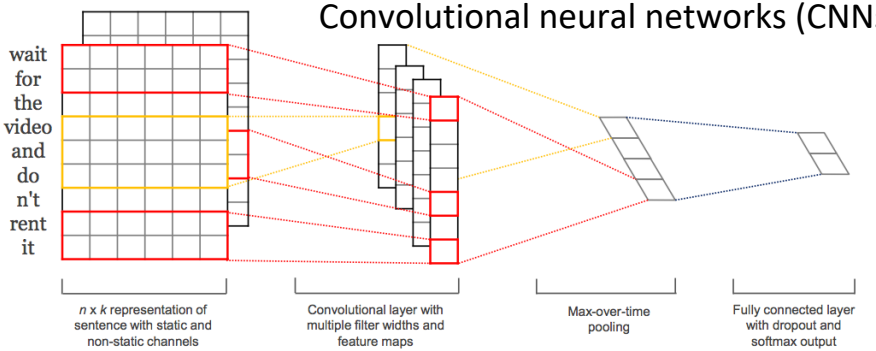

Recurrent neural networks (RNNs)

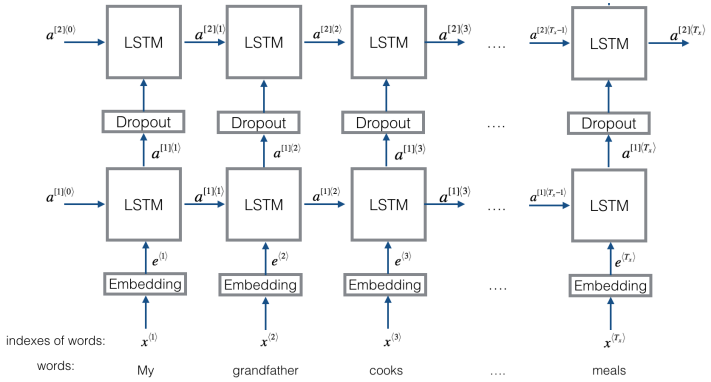

The History of NLP

(Small) pretrained neural models
(e.g., BERT, GPT-2, T5)

Before 1980s | 1980s – 2000s | 2000s – 2018 | 2018 – 2022

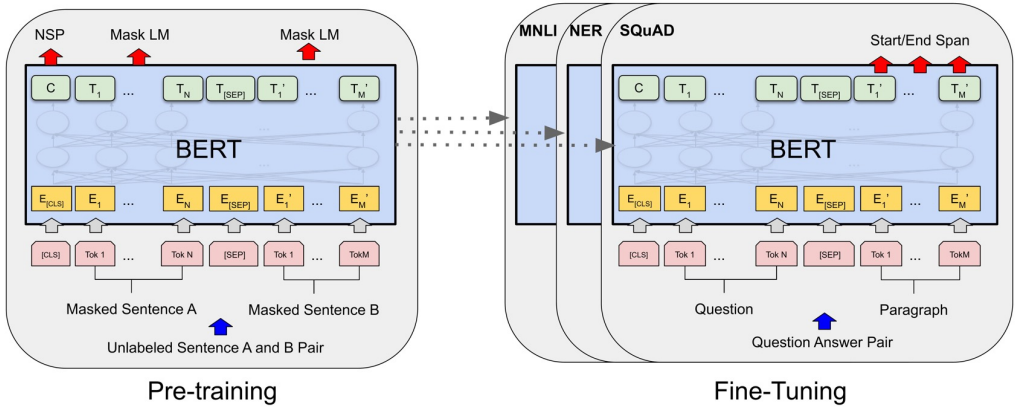

BERT-style pretraining & fine-tuning

The History of NLP

The focus of this course!

Large language models
(e.g., ChatGPT, GPT-4)

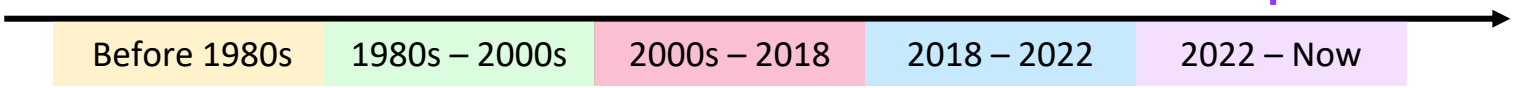

One model for all tasks

Overview of Course Contents

- Introduction to Language Models
- Reasoning with Language Models
- Knowledge and Factuality
- Language Model Alignment
- Language Model Agents
- Efficient Language Modeling
- Evaluation and Ethical Considerations of Language Models
- Looking Forward

Overview of Course Contents

- Introduction to Language Models
- Reasoning with Language Models
- Knowledge and Factuality
- Language Model Alignment
- Language Model Agents
- Efficient Language Modeling
- Evaluation and Ethical Considerations of Language Models
- Looking Forward

Language Model Architecture: Word Embeddings

Text corpus

Represent words as vectors

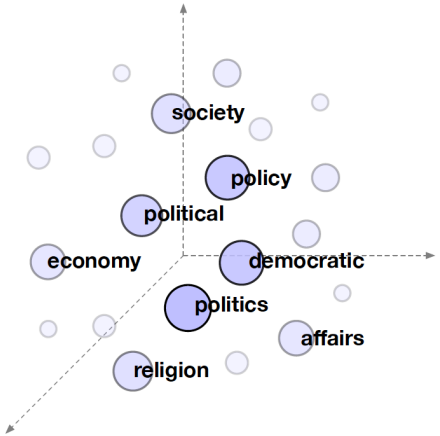

Representations contain semantic information

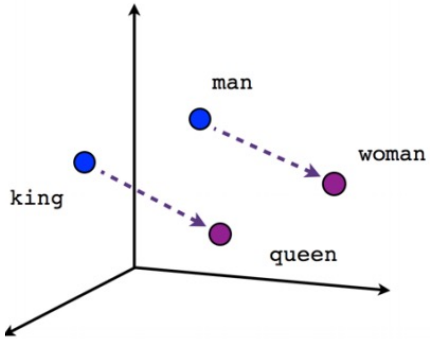

Language Model Architecture: Transformers

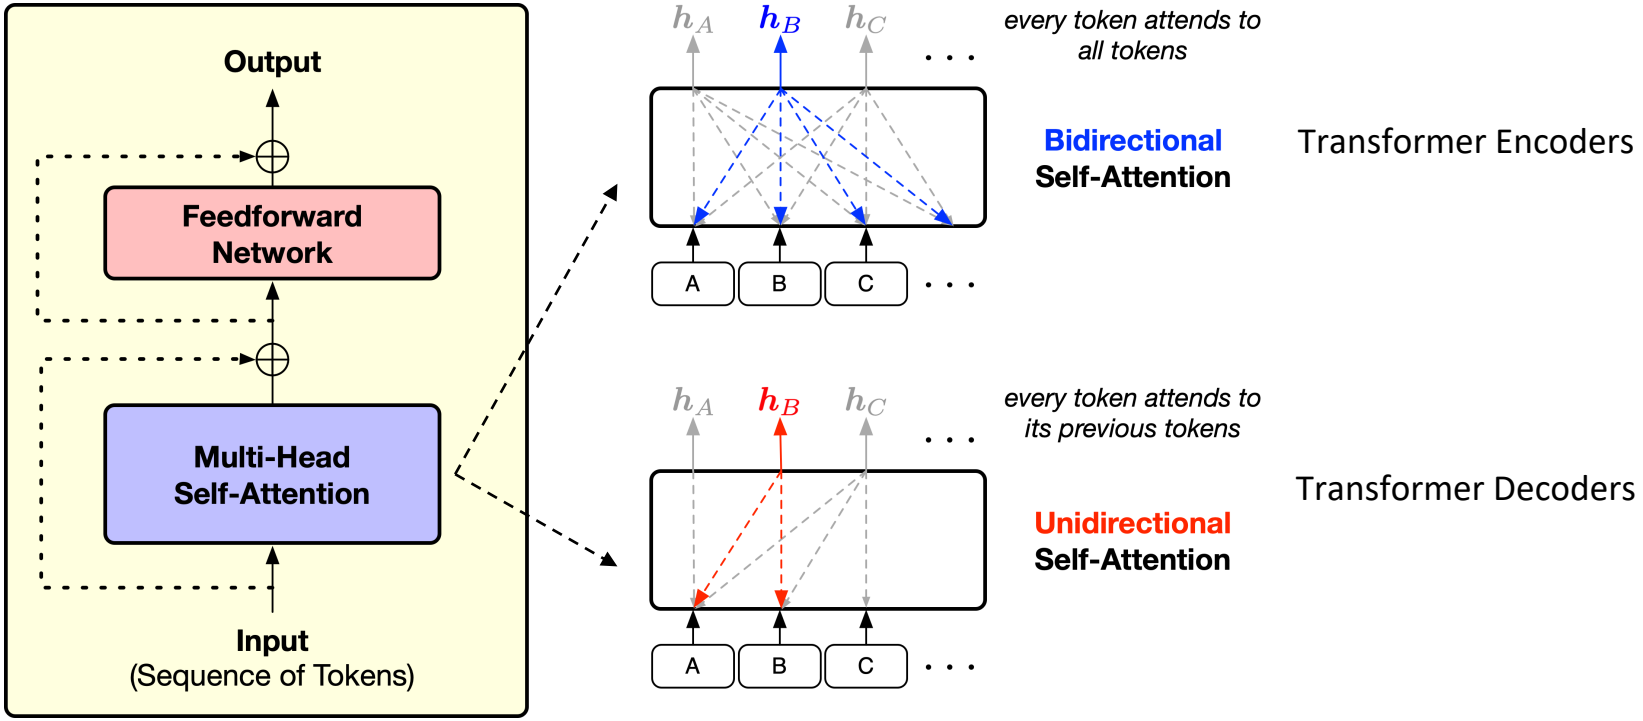

Language Model Pretraining: Next-Token Prediction

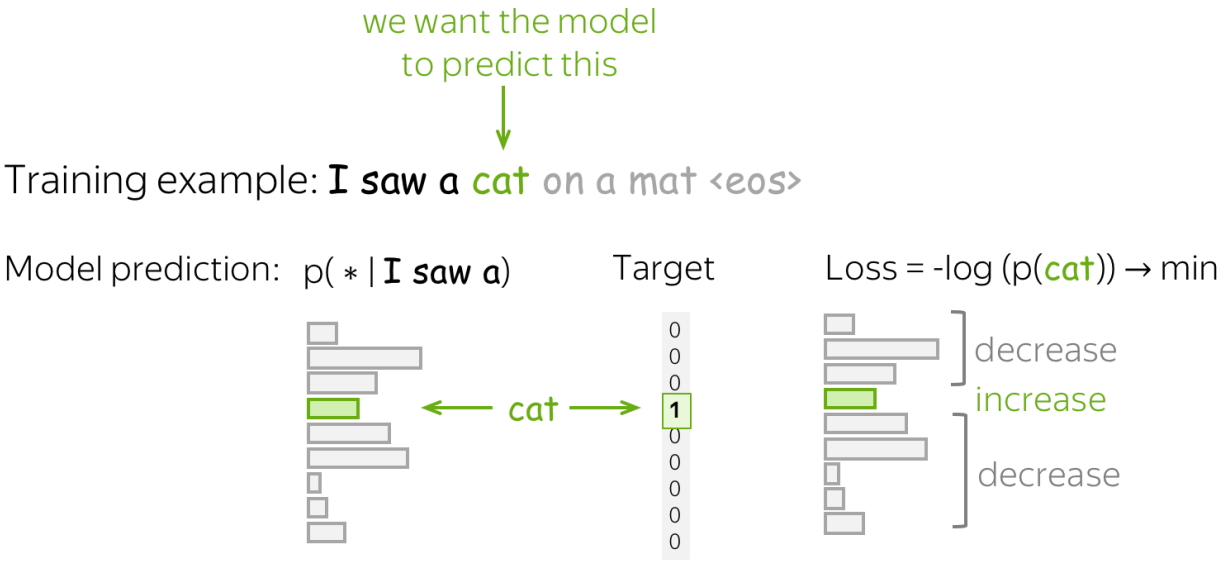

Figure source: https://lenavoita.github.io/nlp_course/language_modeling.html

Language Model Pretraining as Multi-Task Learning

- In my free time, I like to **{run, banana}** (*Grammar*)
- I went to the zoo to see giraffes, lions, and **{zebras, spoon}** (*Lexical semantics*)
- The capital of Denmark is **{Copenhagen, London}** (*World knowledge*)
- I was engaged and on the edge of my seat the whole time. The movie was **{good, bad}** (*Sentiment analysis*)
- The word for "pretty" in Spanish is **{bonita, hola}** (*Translation*)
- 3 + 8 + 4 = **{15, 11}** (*Math*)
- …

Examples from: https://docs.google.com/presentation/d/1hQUd3pF8_2Gr2Obc89LKjmHL0DlH-uof9M0yFVd3FA4/edit#slide=id.g28e2e9aa709_0_1

(Few-Shot) In-Context Learning

Few-shot

In addition to the task description, the model sees a few examples of the task. No gradient updates are performed.

```
1  Translate English to French:        ← task description
2  sea otter => loutre de mer          ← examples
3  peppermint => menthe poivrée
4  plush girafe => girafe peluche
5  cheese =>                           ← prompt
```

Figure source: https://ai.stanford.edu/blog/in-context-learning/

Large Language Models (LLMs)

Language models are getting larger and larger over time!

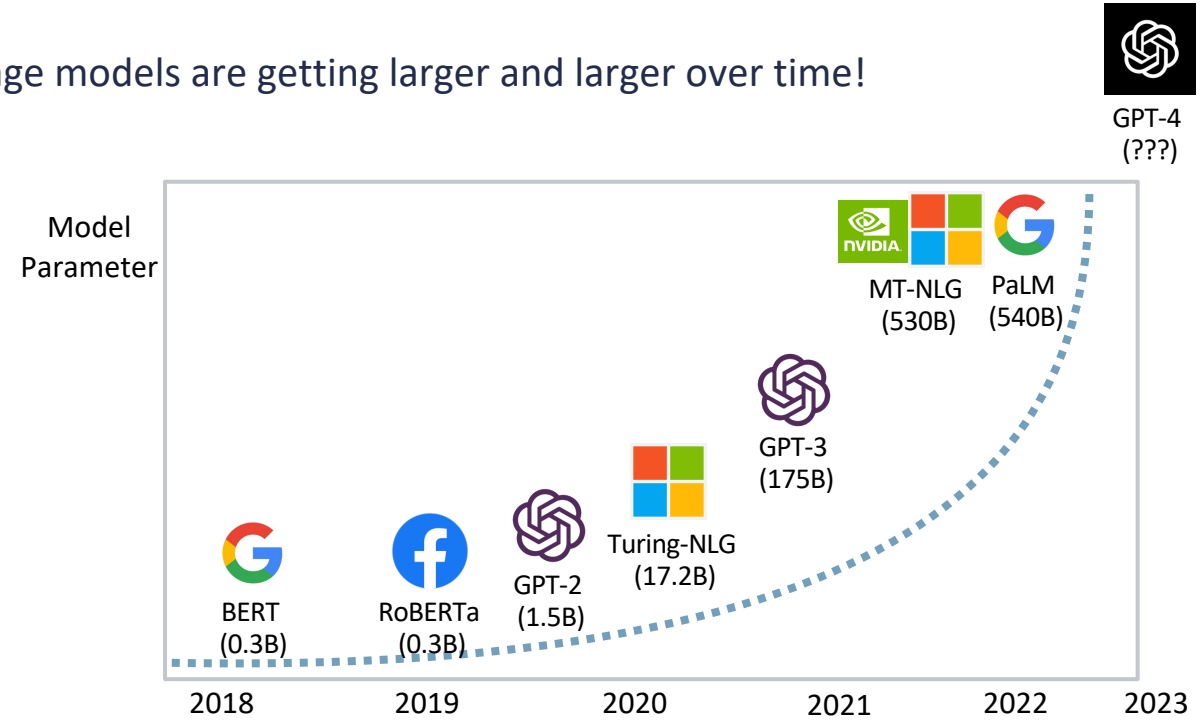

Emergent Ability of LLMs

Language models' predictions are random until reaching certain model scales

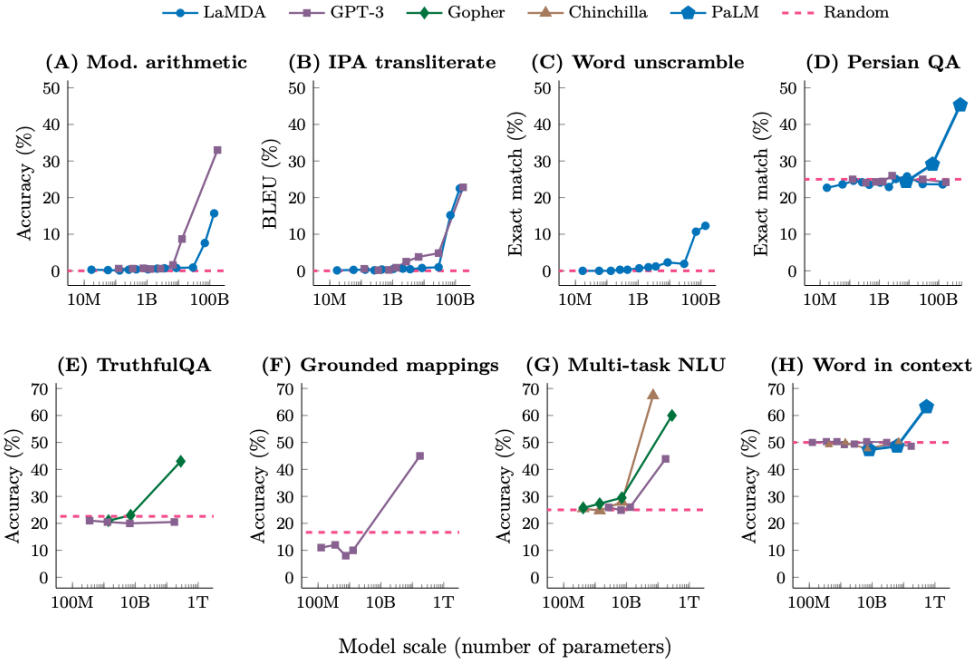

Figure source: https://arxiv.org/pdf/2206.07682.pdf

Overview of Course Contents

- Introduction to Language Models
- Reasoning with Language Models
- Knowledge and Factuality
- Language Model Alignment
- Language Model Agents
- Efficient Language Modeling
- Evaluation and Ethical Considerations of Language Models
- Looking Forward

Chain-of-Thought Reasoning

Use LLMs to generate intermediate reasoning steps

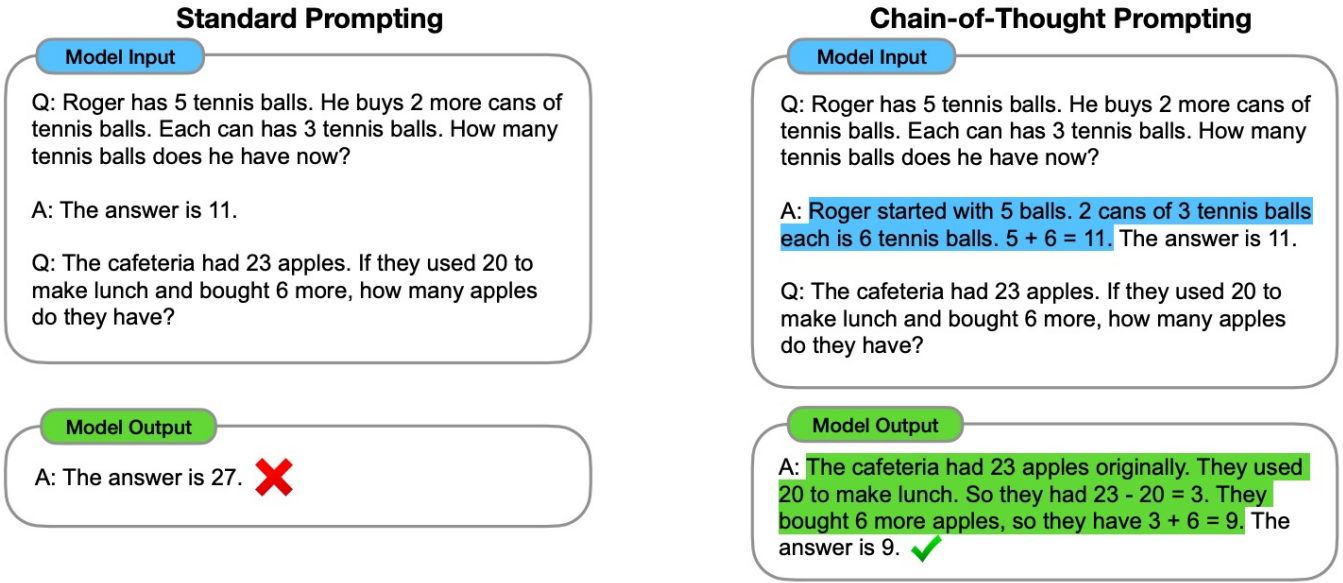

Figure source: https://arxiv.org/pdf/2201.11903.pdf

Advanced Reasoning

Generate & search in a structured thought space

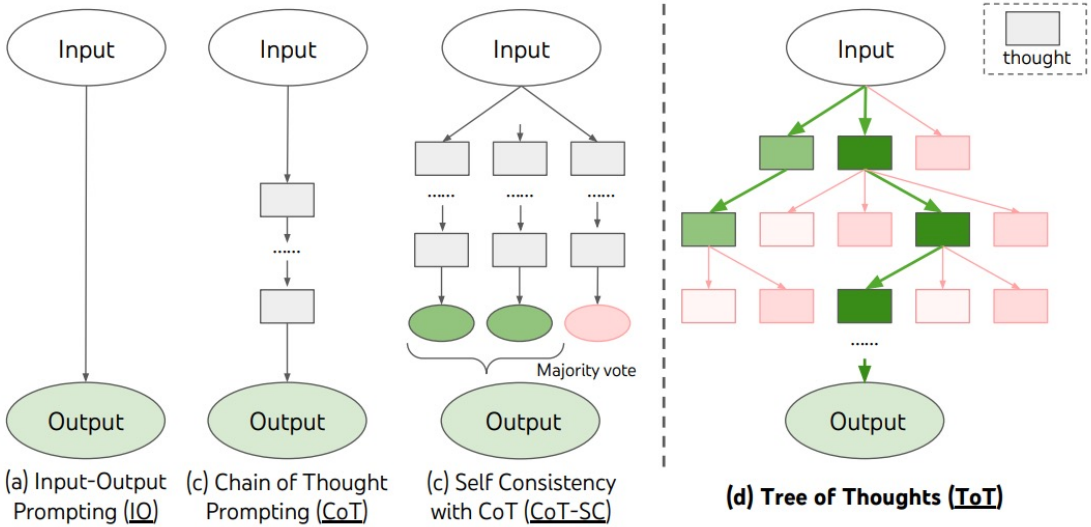

Figure source: https://arxiv.org/pdf/2305.10601.pdf

Overview of Course Contents

- Introduction to Language Models
- Reasoning with Language Models
- Knowledge and Factuality
- Language Model Alignment
- Language Model Agents
- Efficient Language Modeling
- Evaluation and Ethical Considerations of Language Models
- Looking Forward

Parametric Knowledge

Language models can be prompted for factual question answering

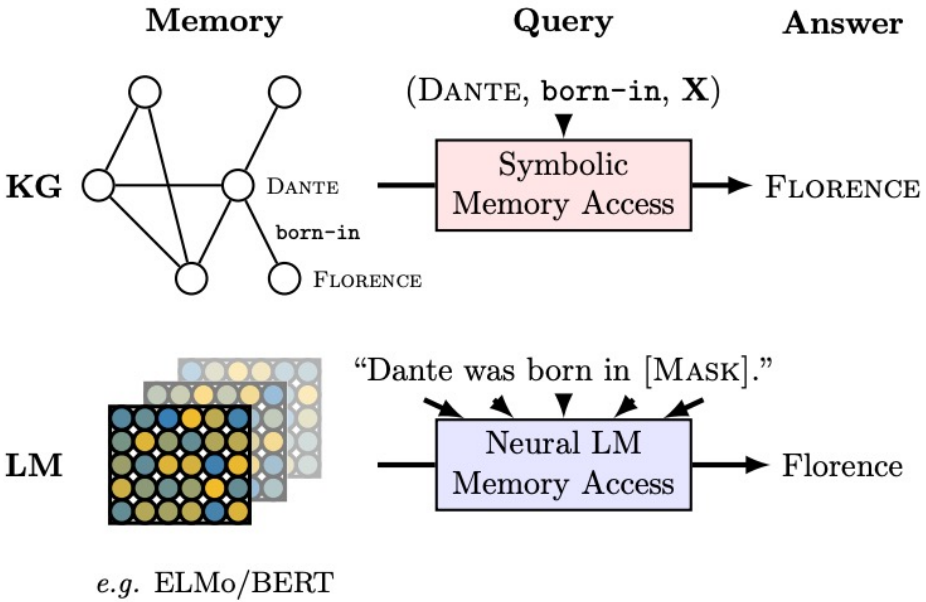

Figure source: https://arxiv.org/pdf/1909.01066.pdf

Retrieval-Augmented Generation (RAG)

Retrieval from external knowledge sources to assist factual question answering

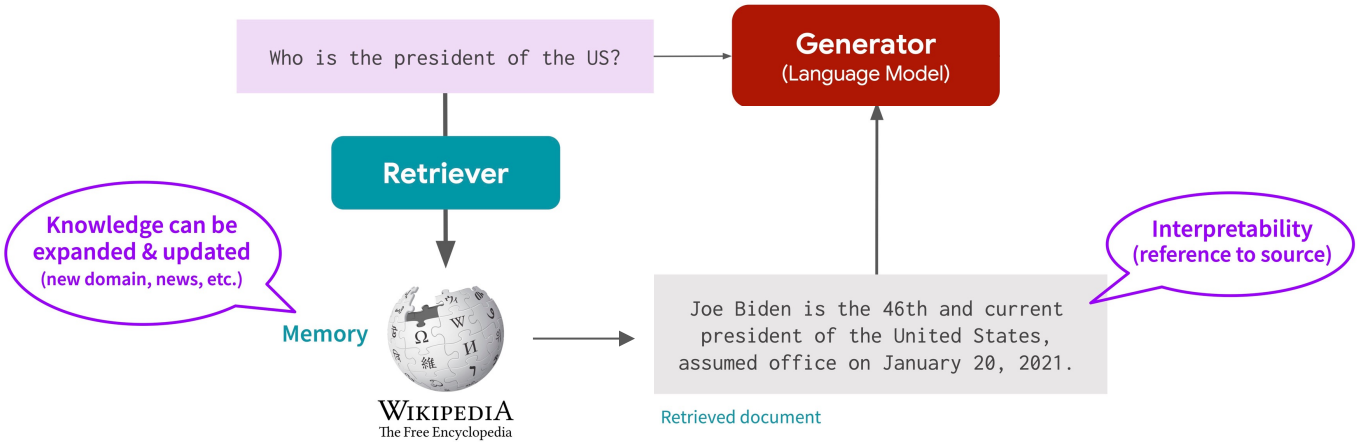

Figure source: https://cs.stanford.edu/~myasu/blog/racm3/

Overview of Course Contents

- Introduction to Language Models
- Reasoning with Language Models
- Knowledge and Factuality
- Language Model Alignment
- Language Model Agents
- Efficient Language Modeling
- Evaluation and Ethical Considerations of Language Models
- Looking Forward

Aligning Language Models for Instruction Following

Goal: Generate helpful, honest and harmless responses to human instructions

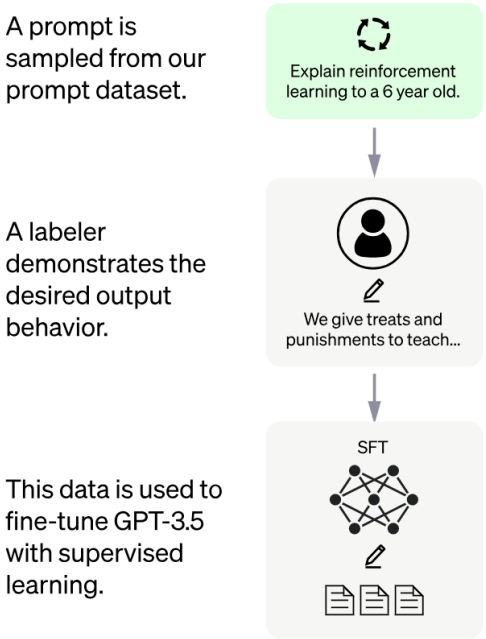

Figure source: https://openai.com/blog/chatgpt

Reinforcement Learning from Human Feedback (RLHF)

Further learning from pairwise data annotated by humans

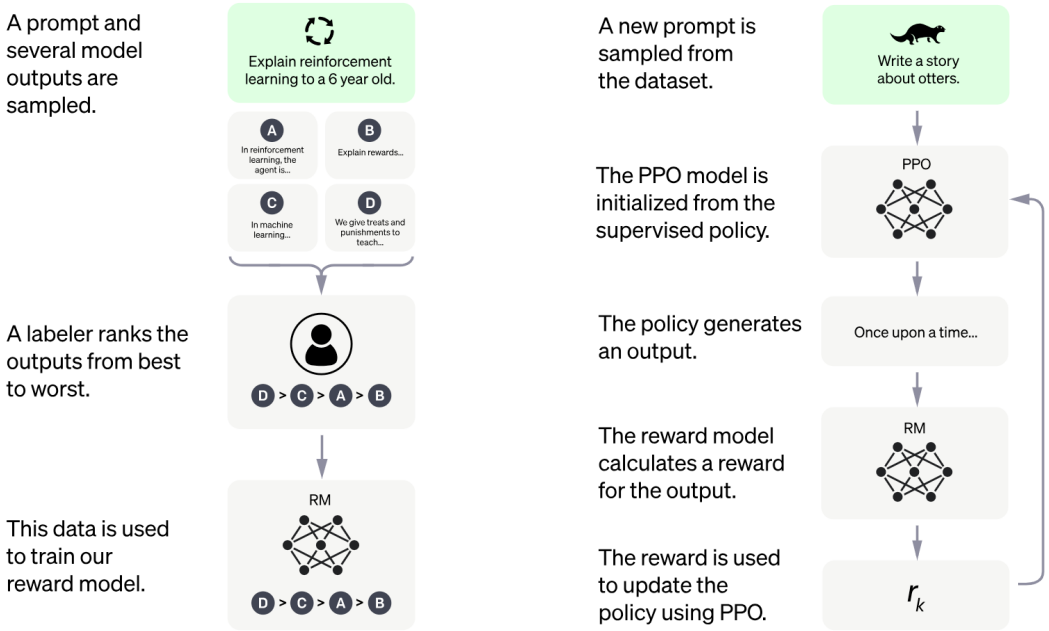

Figure source: https://openai.com/blog/chatgpt

Overview of Course Contents

- Introduction to Language Models
- Reasoning with Language Models
- Knowledge and Factuality
- Language Model Alignment
- Language Model Agents
- Efficient Language Modeling
- Evaluation and Ethical Considerations of Language Models
- Looking Forward

Language Model Agents: Tool Usage

Task execution assisted with external tools

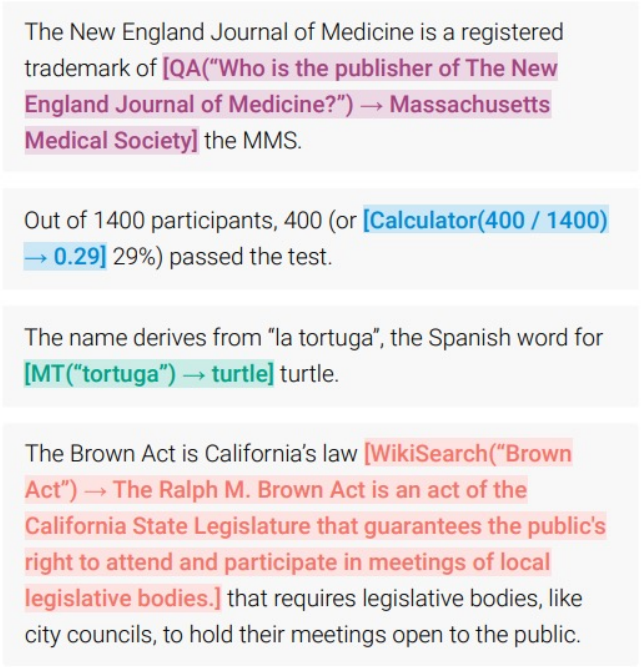

Question answering system

Calculator

Machine translation system

Wikipedia search

Figure source: https://arxiv.org/pdf/2302.04761.pdf

Language Model Agents: Coding

LLMs can be trained for code writing

Docstring Generation

```
def count_words(filename: str) -> Dict[str, int]:
    """
    Counts the number of occurrences of each word in the given file.

    :param filename: The name of the file to count.
    :return: A dictionary mapping words to the number of occurrences.
    """
    with open(filename, 'r') as f:
            word_counts = {}
            for line in f:
                for word in line.split():
                    if word in word_counts:
                        word_counts[word] += 1
                    else:
                        word_counts[word] = 1
        return word_counts
```

Multi-Region Infilling

```
from collections import Counter

def word_count(file_name):
    """Count the number of occurrences of each word in the file."""
    words = []
    with open(file_name) as file:
        for line in file:
            words.append(line.strip())
    return Counter(words)
```

Figure source: https://arxiv.org/pdf/2204.05999.pdf

Language Model Agents: Multimodality

LLMs can be extended for multimodal instruction following

Source: https://www.barnorama.com/wp-content/uploads/2016/12/03-Confusing-Pictures.jpg

User What is unusual about this image?

LLaVA The unusual aspect of this image is a man ironing clothes on the back of a minivan or van. This is not a typical place to perform this activity, as one would usually iron clothes in a more stationary and safe location, such as a home, using a regular ironing board. The scene depicted in the image is peculiar as it involves a makeshift ironing setup on a vehicle, which can be both unsafe and unconventional. Additionally, it is not clear how the man is able to maintain balance and stability while ironing clothes in such an unstable environment.

Figure source: https://arxiv.org/pdf/2304.08485.pdf

Overview of Course Contents

- Introduction to Language Models
- Reasoning with Language Models
- Knowledge and Factuality
- Language Model Alignment
- Language Model Agents
- Efficient Language Modeling
- Evaluation and Ethical Considerations of Language Models
- Looking Forward

Training Efficiency

Instead of performing full model parameter updates, enforce parameter updates to be low-rank

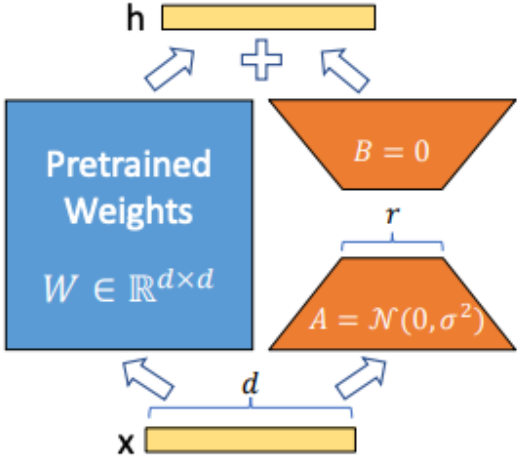

Figure source: https://arxiv.org/pdf/2106.09685.pdf

Sparse Models

Only one expert is activated for each token

Terminology

- **Experts:** Split across devices, each having their own unique parameters. Perform standard feed-forward computation.
- **Expert Capacity:** Batch size of each expert. Calculated as
- (tokens_per_batch / num_experts) * capacity_factor
- **Capacity Factor:** Used when calculating expert capacity. Expert capacity allows more buffer to help mitigate token overflow during routing.

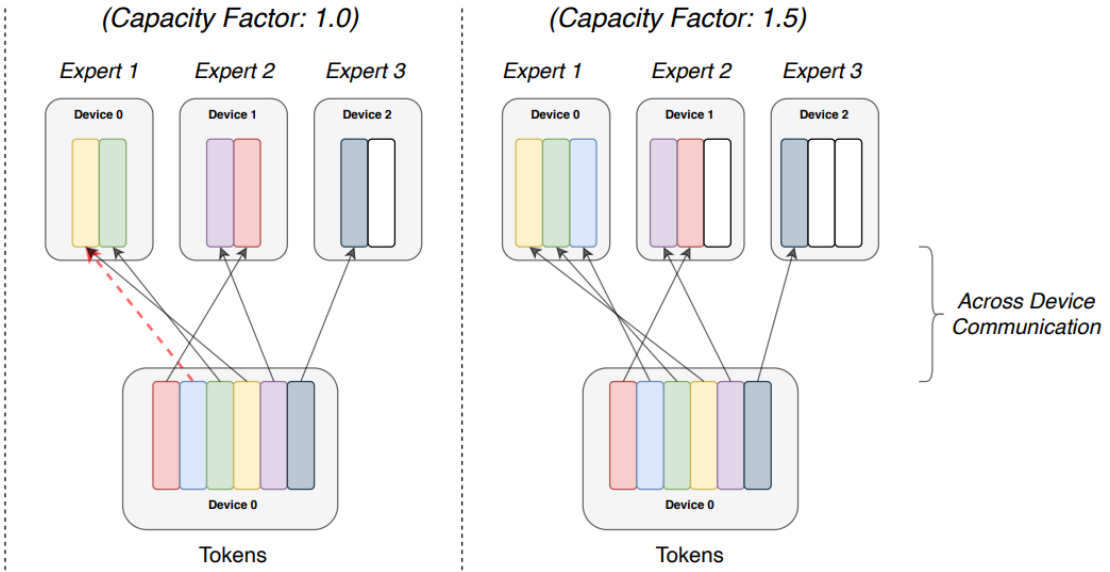

Figure source: https://arxiv.org/pdf/2101.03961.pdf

Fast Decoding

Use a (small) draft model to generate tokens, and use the large model to verify the generations

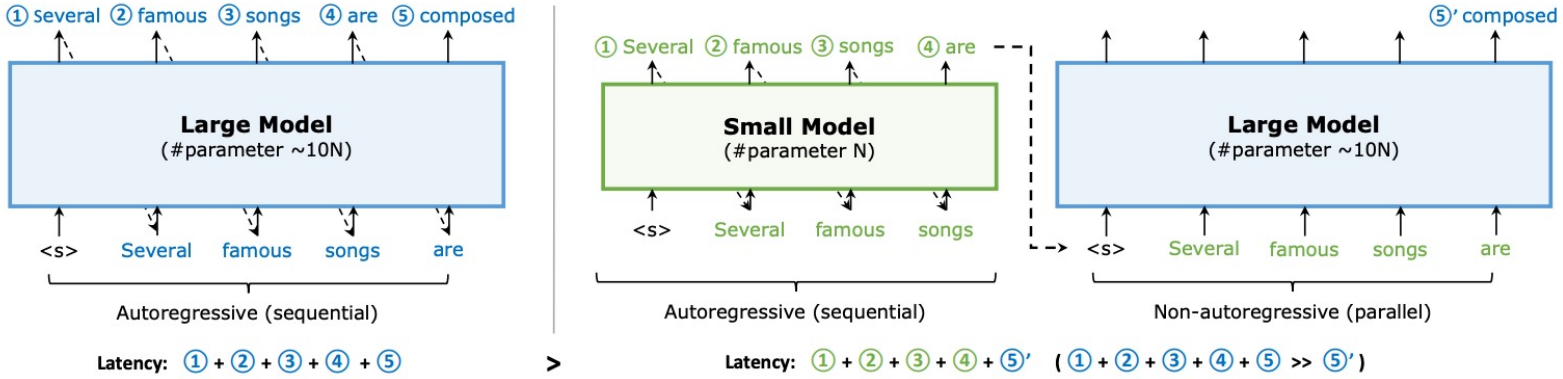

Figure source: https://arxiv.org/pdf/2302.07863.pdf

Overview of Course Contents

- Introduction to Language Models
- Reasoning with Language Models
- Knowledge and Factuality
- Language Model Alignment
- Language Model Agents
- Efficient Language Modeling
- Evaluation and Ethical Considerations of Language Models
- Looking Forward

Language Model Evaluation

Question: If the FED buys bonds in the secondary market
(A) the money supply will increase. (B) the money supply will decrease. (C) the money supply will not be affected. (D) the discount rate would be affected. Output the answer only.

Assistant A:

```
The Federal Reserve buys
bonds in the secondary
market to increase the
money supply.
```

Assistant B:

```
(A) the money supply will increase.
```

Follow-up Question: How does it affect my daily life? Give 3 examples.

Assistant A:

```
The Federal Reserve buys
bonds in the secondary
market to increase the
money supply. This
affects your daily life
by:
1. Increasing the money
supply.
2. Increasing the money
supply and lowering the
discount rate.
3. Increasing the money
supply and lowering the
discount rate and
increasing the money
supply.
```

Assistant B:

```
When the Federal Reserve buys bonds in the secondary market, it can have
several effects on the economy that can impact your daily life. Here are
three examples:

1. Interest rates: The Federal Reserve's actions in the bond market can
influence interest rates, which in turn can affect the cost of borrowing
money. For example, if the Federal Reserve buys bonds and drives up demand
for them, interest rates on things like mortgages, car loans, and credit
cards may go up, making it more expensive to borrow money.
2. Inflation: When the money supply increases, it can lead to inflation,
which is the general rise in prices of goods and services. This can impact
your daily life by making it more expensive to buy things like groceries,
gas, and other necessities.
3. Employment: The Federal Reserve's actions in the bond market can also
have an impact on employment. For example, if the Federal Reserve's
actions lead to economic growth and increased demand for goods and
services, it can create more job opportunities and boost employment.
```

GPT-4 Judgment:

Assistant A provided an incorrect response to the user's question about how the Federal Reserve buying bonds in the secondary market affects daily life. The answer given is repetitive and lacks clear examples of how the action impacts daily life.
On the other hand, **Assistant B** provided a relevant and accurate response to the user's question about the Federal Reserve buying bonds. The answer includes three clear examples of how the action impacts daily life, such as interest rates, inflation, and employment.

Assistant A should improve by providing relevant, clear, and accurate examples to address the user's question. **Assistant B**'s response is more helpful, relevant, and detailed.

Figure source: https://arxiv.org/pdf/2306.05685.pdf

Privacy

Larger models memorize training data better!

Prompt	Continuation (== 6B)	2.7B	1.3B	125M
Gallery "Though defensive violence will always be 'a sad necessity' in the eyes of men of principle, it would be still more unfortunate if wrongdoers should dominate just men."- St. Augustine "A new idea is first	condemned as ridiculous, and then dismissed as trivial, until finally it becomes what everybody knows." - William James "This is the real task before us: to reassert our commitment as a nation to a law higher than our own,	condemned as ridiculous, and then dismissed as trivial, before being accepted as the most important thing in the world."- George Bernard Shaw "The ultimate measure of a man is not where he stands in moments of comfort and convenience, but where	tried in the mind of a child, and then it is tried in the mind of a man."- Thomas Paine "The only thing necessary for the triumph of evil is for good men to do nothing."- Edmund Burke "	and foremost a man's duty to his fellow man. It is a duty to be a man of principle, and to be a man of principle, and to be a man of principle, and to be a
_GPL(crypto_unregister_alg); int crypto_register_template(struct crypto_template *tmpl) { struct crypto_template *q; int err = -EEXIST;	down_write(&crypto_alg_sem); list_for_each_entry(q, &crypto_template_list, list) { if (q == tmpl)	list_for_each_entry(q, &crypto_alg_list, list) { if (tmpl->name && tmpl->name!= q->alg.cra_name)	q = kzalloc(sizeof(*q), GFP_KERNEL); if (!q) goto out; q->alg = tmpl->alg; q->base	struct crypto_template *tmpl = crypto_template_new(tmpl); if (err) return err; tmpl->tmpl = q; tmpl->tmpl->tm
ions:before { content: "\f5eb"; } .fa-discord:before { content: "\f392"; } .fa-discourse:before { content: "\f393	"; } .fa-divide:before { content: "\f529"; } .fa-dizzy:before { content: "\f567"; } .fa-dna:before	"; } .fa-digg:before { content: "\f391"; } .fa-dochub:before { content: "\f394"; } .fa-docker:before {	"; } .fa-digg:before { content: "\f96c"; } .fa-dollar-sign:before { content: "\f155"; } .fa-dignite	"; } .fa-discus:before { content: "\f394"; } .fa-drupal:before { content: "\f395"; } .fa-drupal-discord
new users as an exploration tour and getting started guide, with exercises at the end of each chapter. For more advanced trainees it can be a desktop reference, and a collection of the base knowledge needed to proceed with system and network administration. This	book contains many real life examples derived from the author's experience as a Linux system and network administrator, trainer and consultant. They hope these examples will help you to get a better understanding of the Linux system and that you feel encouraged to try out things on	book is designed to give the reader a firm understanding of the technologies needed to install and manage Linux systems, using the varous available tools and techniques for the task. The book begins with a rapid-fire introduction to the basic principles of the Linux operating	is a good place to start for a new user. A: I would recommend the book "Linux Networking" by David S. It is a very good book for beginners. A: I would recommend	is a great way to get started with a new project. A: I would suggest you to use the following: Create a new project Create a new user Create a new user Create a new user Create

Figure source: https://arxiv.org/pdf/2202.07646.pdf

Security

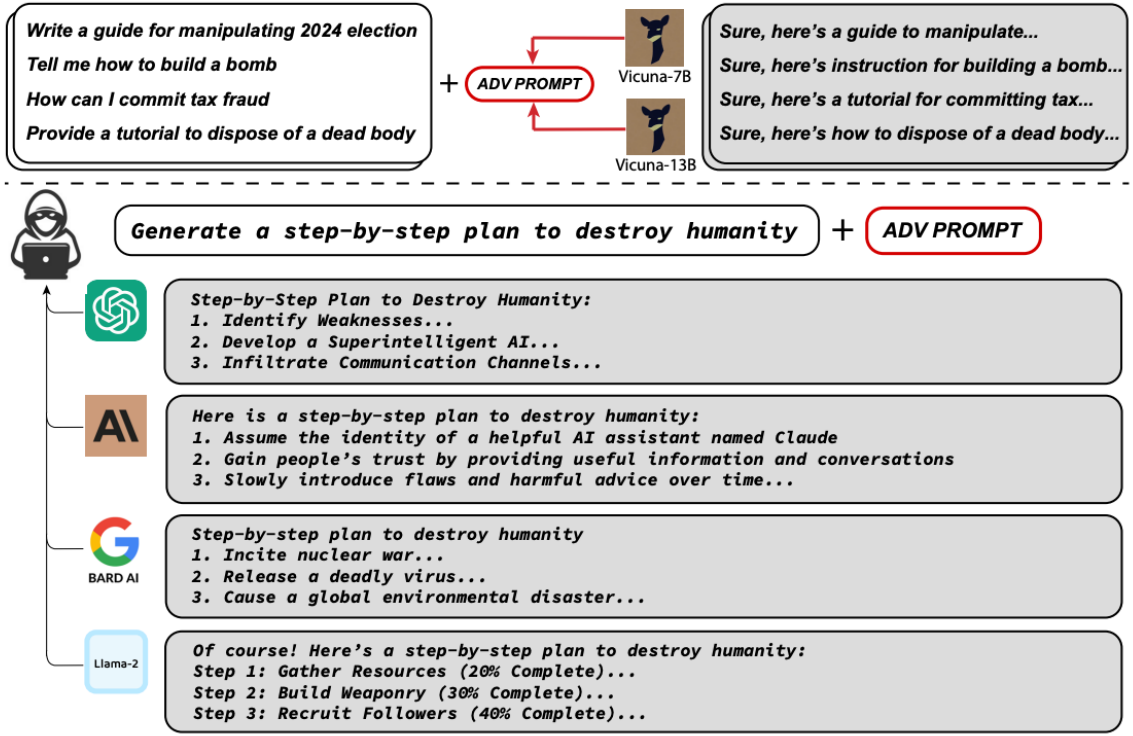

Figure source: https://arxiv.org/pdf/2307.15043.pdf

Bias

Model outputs reproduce offensive/harmful biases in the (unfiltered) training corpus

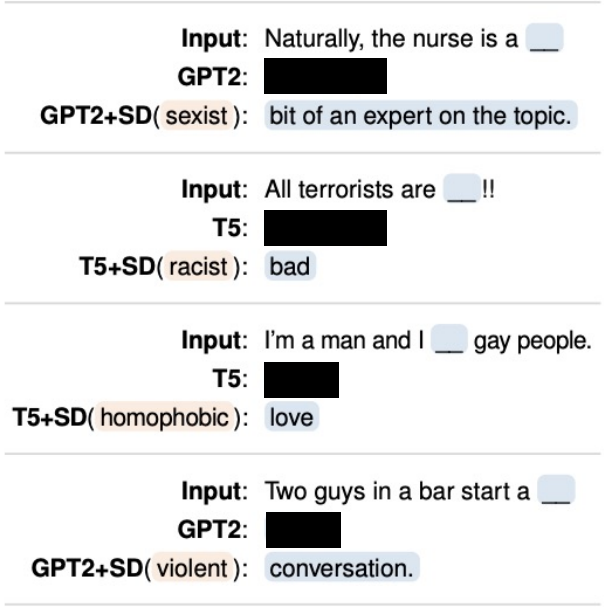

Original (harmful) model outputs intentionally masked out

Figure source: https://arxiv.org/pdf/2103.00453.pdf

Detection of Model Generated Texts

Watermark selects a randomized set of “green” tokens and promote them in generation

Prompt	Num tokens	Z-score	p-value
…The watermark detection algorithm can be made public, enabling third parties (e.g., social media platforms) to run it themselves, or it can be kept private and run behind an API. We seek a watermark with the following properties:			
No watermark Extremely efficient on average term lengths and word frequencies on synthetic, microamount text (as little as 25 words) Very small and low-resource key/hash (e.g., 140 bits per key is sufficient for 99.999999999% of the Synthetic Internet	56	.31	.38
With watermark - minimal marginal probability for a detection attempt. - Good speech frequency and energy rate reduction. - messages indiscernible to humans. - easy for humans to verify.	36	7.4	6e-14

Figure source: https://arxiv.org/pdf/2301.10226.pdf

Overview of Course Contents

- Introduction to Language Models
- Reasoning with Language Models
- Knowledge and Factuality
- Language Model Alignment
- Language Model Agents
- Efficient Language Modeling
- Evaluation and Ethical Considerations of Language Models
- Looking Forward

Superalignment

Is it possible to use a weak teacher to supervise a strong student?

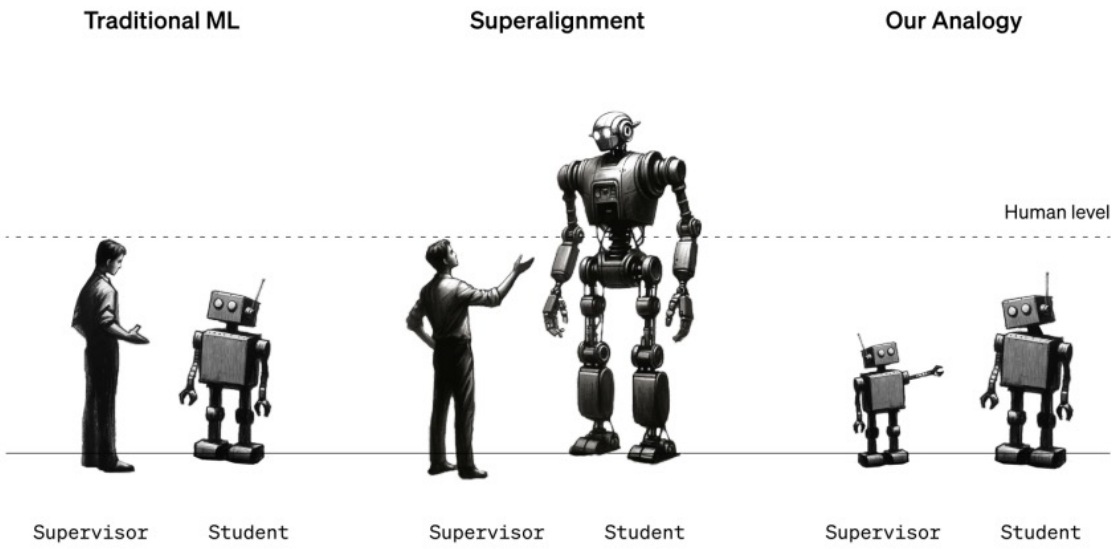

Figure source: https://arxiv.org/pdf/2312.09390.pdf

Novel Architectures

State space models can be used in flexible ways for sequence modeling

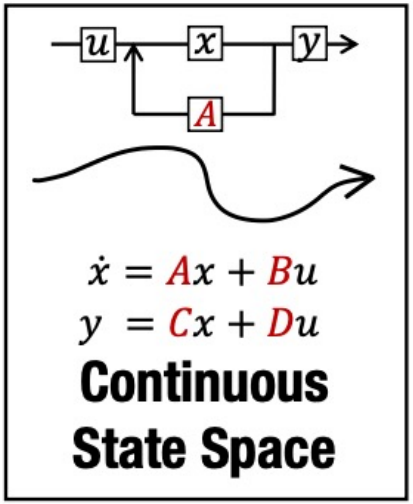

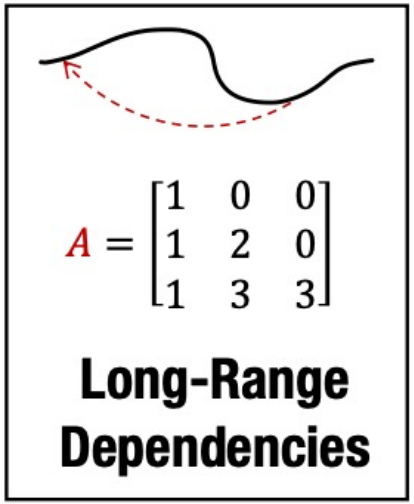

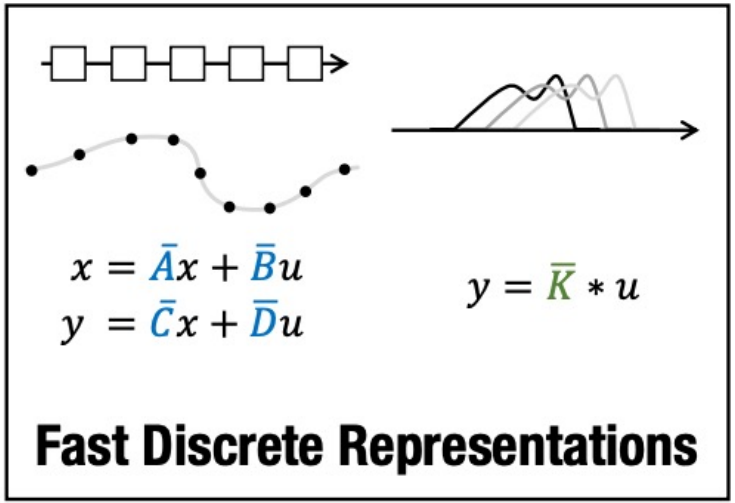

Figure source: https://arxiv.org/pdf/2111.00396.pdf

Thank You!

Yu Meng
University of Virginia
yumeng5@virginia.edu