

Large Language Models

Jian Tang

HEC Montreal

Mila-Quebec AI Institute


Email: jian.tang@hec.ca



Some of the slides are borrowed from: <https://phontron.com/class/anlp-fall2024/>

Overview of LLMs Training

- Pretraining -> Supervised Fine-tuning (SFT) -> Reinforcement Learning Human Feedback (RLHF)




Outline

- Fine-tuning
- Instruction Tuning
- Prompting


Standard Multi-task Learning

- Train a neural network for multiple different tasks




Pre-train and Fine-tune Framework

- Pre-train a model on a set of tasks and then fine-tune it on new tasks
 - E.g. Pretrain a LLM and then fine-tune it with a classification task




Prompting

- Fix the parameters of LLM, change the instructions
 - Some task-specific labeled examples can be inserted into the prompts



Instruction Tuning

- Fine-tuning LLMs with many different tasks, with instructions specifying each task
 - The parameters of LLMs will be updated



Outline

- Fine-tuning
- Instruction Tuning
- Prompting

Different NLP Downstream Tasks

Context-Free Question Answering

- Also called “open-book QA”
- Answer a question without any specific grounding into documents
- Example dataset: MMLU (Hendrycks et al. 2020)

Professional Law

As Seller, an encyclopedia salesman, approached the grounds on which Hermit's house was situated, he saw a sign that said, "No salesmen. Trespassers will be prosecuted. Proceed at your own risk."

Although Seller had not been invited to enter, he ignored the sign and drove up the driveway toward the house. As he rounded a curve, a powerful explosive charge buried in the driveway exploded, and Seller was injured. Can Seller recover damages from Hermit for his injuries?

- (A) Yes, unless Hermit, when he planted the charge, intended only to deter, not harm, intruders. 
- (B) Yes, if Hermit was responsible for the explosive charge under the driveway. 
- (C) No, because Seller ignored the sign, which warned him against proceeding further. 
- (D) No, if Hermit reasonably feared that intruders would come and harm him or his family. 

Contextual Question Answering

- Also called “machine reading”, “closed-book QA”
- Answer a question about a document or document collection
- *Example:* Natural Questions (Kwiatkowski et al. 2019) is grounded in a Wikipedia document, or the Wikipedia document collection

Question: what color was john wilkes booth’s hair

Wikipedia Page: John_Wilkes_Booth

Long answer: Some critics called Booth “the handsomest man in America” and a “natural genius”, and noted his having an “astonishing memory”; others were mixed in their estimation of his acting. He stood 5 feet 8 inches (1.73 m) tall, had jet-black hair , and was lean and athletic. Noted Civil War reporter George Alfred Townsend described him as a “muscular, perfect man” with “curling hair, like a Corinthian capital”.

Short answer: jet-black

Code Generation

- Generate code (e.g. Python, SQL, etc.) from a natural language command and/or input+output examples
- *Example:* HumanEval (Chen et al. 2021) has evaluation questions for Python standard library

```
def incr_list(l: list):  
    """Return list with elements incremented by 1.  
    >>> incr_list([1, 2, 3])  
    [2, 3, 4]  
    >>> incr_list([5, 3, 5, 2, 3, 3, 9, 0, 123])  
    [6, 4, 6, 3, 4, 4, 10, 1, 124]  
    """  
    return [i + 1 for i in l]
```

Text Summarization

- Single-document: Compress a longer document to shorter
- Multi-document: Compress multiple documents into one
- Example: WikiSum compresses the references in a Wikipedia article into the first paragraph

References

1. ^ "Barack Hussein Obama Takes The Oath Of Office" [on YouTube](#). January 20, 2009.
2. ^ "American Presidents: Greatest and Worst – Siena College Research Institute" [Archived](#) from the original on July 15, 2022. Retrieved February 12, 2023.
3. ^ "Barack Obama I C-SPAN Survey on Presidents 2017" [Archived](#) from the original on February 12, 2023. Retrieved February 12, 2023.
4. ^ "Siena's 6th Presidential Expert Poll 1982–2018 – Siena College Research Institute" [Archived](#) from the original on July 19, 2019. Retrieved February 13, 2023.
5. ^ "President Barack Obama" [The White House](#). 2008. Archived from the original [on October 26, 2009](#). Retrieved December 12, 2008.
6. ^ "President Obama's Long Form Birth Certificate" [whitehouse.gov](#). April 27, 2011. [Archived](#) from the original on July 31, 2023. Retrieved August 4, 2023.
7. ^ "Certificate of Live Birth: Barack Hussein Obama II, August 4, 1961, 7:24 pm, Honolulu" [\(PDF\)](#). [whitehouse.gov](#). April 27, 2011. Archived from the original [\(PDF\)](#) on March 3, 2017. Retrieved March 11, 2017 – via [National Archives](#).



Barack Obama

Article Talk

F

From Wikipedia, the free encyclopedia

"Barack" and "Obama" redirect here. For other uses, see [Barack \(disambiguation\)](#), [Obama \(disambiguation\)](#)

Barack Hussein Obama II (/bəˈrɑːk huːseɪn oʊˈbɑːmə/ [ⓘ] [ⓘ] *bə-RAHK hoo-SAYN oh-BAH-mə*^[1] born August 4, 1961) is an American politician who served as the 44th [president of the United States](#) from 2009 to 2017. A member of the [Democratic Party](#), he was the first [African-American president](#) in U.S. history. Obama previously served as a U.S. senator representing Illinois from 2005 to 2008, as an [Illinois state senator](#) from 1997 to 2004, and as a civil rights lawyer and university lecturer.

Obama was born in [Honolulu, Hawaii](#). He graduated from [Columbia University](#) in 1983 with a B.A. in political science and later worked as a [community organizer](#) in Chicago. In 1988, Obama enrolled in [Harvard Law School](#), where he was the first black president of the [Harvard Law Review](#). He became a civil rights attorney and an academic, teaching [constitutional law](#) at the [University of Chicago Law School](#) from 1992 to 2004. He also went into elective politics. Obama represented the [13th district in the Illinois Senate](#) from 1997 until 2004, when he [successfully ran for the U.S. Senate](#). In 2008, after a close primary campaign against [Hillary Clinton](#), he was nominated by the Democratic Party for president and chose Delaware Senator [Joe Biden](#) as his running mate. Obama was elected president, defeating [Republican Party](#) nominee [John McCain](#) in the [presidential election](#) and was inaugurated on January 20, 2009. Nine months later he was named the [2009 Nobel Peace Prize](#) laureate, a decision that drew a mixture of praise and criticism.

Information Extraction

- *Entity recognition*: identify which words are entities
- *Entity linking*: link entities to a knowledge base (e.g. Wikipedia)
- *Entity co-reference*: find which entities in an input correspond to each-other
- *Event recognition/linking/co-reference*: identify what events occurred
- Example: OntoNotes (Weischedel et al. 2013) annotates many types of information like this on various domains

Machine Translation



- Translate from one language to another
- Quality assessment done using similarity to reference translation
- Example: FLORES dataset (Goyal et al. 2021) — translations of Wikipedia articles into 101 languages

Mismatch between LLM Pretraining and Downstream Applications

- LLMs are pretrained on the task of next token prediction
- Downstream applications: perform specific tasks
 - E.g., sentiment classification
 - Question answering
 - Summarization
 - Machine translation

Instruction Tuning

- Pre-train an LLM, then fine-tune it on many different tasks, and generalize to new tasks




Instruction Tuning

- Finetuning LLMs using (INSTRUCTION, OUTPUT) pairs
 - INSTRUCTION: human instruction for the model
 - OUTPUT: the desired output that follows the INSTRUCTION
- Objectives:
 - Bridge the gap between pretraining objectives and downstream applications
 - Steer the model behavior towards more controllable and predictable
 - Quickly adapt to a new domain without extensive retraining

Instruction Dataset Construction

- Each instance is composed of three elements:
 - An instruction: a natural language text sequence to specify the task
 - An optional input: supplementary information for context
 - Output: an anticipated output based on the instruction and the input



Instruction Dataset Construction

- Two approaches
 - Data integration from annotated natural language datasets
 - Generating outputs from LLMs





Figure 1: General pipeline of instruction tuning.

Automatically generating instruction tuning data from LLMs (Self-Instruct, Wang et al. 2022)




Outline

- Fine-tuning
- Instruction Tuning
- Prompting

What is Prompting?

- Encouraging a pre-trained model to make particular predictions by providing a textual “prompt” specifying the task to be done
 - providing LLMs with more context and information
- With prompts, the model can better understand what kind of output is expected and produce more accurate and relevant results.



Basic Prompting (Radford et al. 2018)

- Append a textual string to the beginning of the sequence and complete

x = **When a dog sees a squirrel, it will usually**

(GPT-2 Small) be afraid of anything unusual. As an exception, that's when
a squirrel is usually afraid to bite.

(GPT-2 XL) lick the squirrel. It will also touch its nose to the squirrel on
the tail and nose if it can.

Standard Prompting Workflow

- Fill a prompt template
- Predict the answer
- Post-process the answer

Prompt Templates

- A template where you fill in with an actual input

Input: $x = \text{"I love this movie"}$



Template: $[x]$ Overall, it was $[z]$



Prompting: $x' = \text{"I love this movie. Overall it was } [z]\text{"}$

Chat Prompts

- Recently, many models are trained as chatbots
- Usually inputs are specified in OpenAI messages format

```
messages=[  
  {  
    "role": "system",  
    "content": "Please classify movie reviews as 'positive' or 'negative'."  
  },  
  {  
    "role": "user",  
    "content": "This movie is a banger."  
  },  
]
```

- Roles:
 - **“system”**: message provided to the system to influence behavior
 - **“user”**: message input by the user
 - **“assistant”**: message output by the system

Chat Prompts Behind the Scenes

- Behind the scenes, messages are converted to token strings

LLaMa

```
Sys. | [INST]
      | <<SYS>>
      | You are an assistant that ...
      | <</SYS>>
      | [/INST]
User | [INST]This movie is great.[/INST]
Asst. | Positive.
```

Alpaca

```
Sys. | ### Instruction:
      | You are an assistant that ...
User | ### Instruction:
      | This movie is great.
Asst. | ### Response:
      | Positive.
```

Answer Prediction

- Given a prompt, predict the answer

Prompting: $x' =$ "I love this movie. Overall it was [z]"



Predicting: $x' =$ "I love this movie. Overall it was fantastic"

- Use any inference algorithms

Post-processing

- Based on the answer, select the actual output
- For instance:
 - Taking the output as-is
 - Formatting the output for easy visualization
 - Selecting only parts of the output that you want to use
 - Mapping the outputs to other actions

Explicit Instructions

Stylization

Explain this to me like a topic on a children's educational network show teaching elementary students.

I'm a software engineer using large language models for summarization.
Summarize the following text in under 250 words:

Give your answer like an old timey private investigator hunting down a case step by step.

Explicit Instructions

Formatting

Use bullet points.

Return as a JSON object.

Use less technical terms and help me apply it in my work in communications.



Explicit Instructions

- More specific results on recently created resources



```
Explain the latest advances in large language models to me.  
# More likely to cite sources from 2017
```

```
Explain the latest advances in large language models to me. Always cite your  
sources.
```

```
Never cite sources older than 2020.
```

```
# Gives more specific advances and only cites sources from 2020
```

In-context Learning/Few-shot Prompting (Brown+ 2021)

- Provide a few labeled examples of the downstream task in the instruction
 - Zero-shot
 - One-shot
 - Few-shot

Instruction | Please classify movie reviews as 'positive' or 'negative'.

Examples

Input: I really don't like this movie.
Output: negative

Input: This movie is great!
Output: positive

Chain of Thought Prompting (Wei et al. 2022)

- Push the model to explain its reasoning before generating a response




Figure 1: Chain-of-thought prompting enables large language models to tackle complex arithmetic, commonsense, and symbolic reasoning tasks. Chain-of-thought reasoning processes are highlighted.

Unsupervised Chain-of-thought Prompting (Kojima et al. 2022)

- Adding a prompt to encourage the model to generate explanations before generating the response

(a) Few-shot

Q: Roger has 5 tennis balls. He buys 2 more cans of tennis balls. Each can has 3 tennis balls. How many tennis balls does he have now?

A: The answer is 11.

Q: A juggler can juggle 16 balls. Half of the balls are golf balls, and half of the golf balls are blue. How many blue golf balls are there?

A:

(Output) The answer is 8. ✗

(b) Few-shot-CoT

Q: Roger has 5 tennis balls. He buys 2 more cans of tennis balls. Each can has 3 tennis balls. How many tennis balls does he have now?

A: Roger started with 5 balls. 2 cans of 3 tennis balls each is 6 tennis balls. $5 + 6 = 11$. The answer is 11.

Q: A juggler can juggle 16 balls. Half of the balls are golf balls, and half of the golf balls are blue. How many blue golf balls are there?

A:

(Output) The juggler can juggle 16 balls. Half of the balls are golf balls. So there are $16 / 2 = 8$ golf balls. Half of the golf balls are blue. So there are $8 / 2 = 4$ blue golf balls. The answer is 4. ✓

(c) Zero-shot

Q: A juggler can juggle 16 balls. Half of the balls are golf balls, and half of the golf balls are blue. How many blue golf balls are there?

A: The answer (arabic numerals) is

(Output) 8 ✗

(d) Zero-shot-CoT (Ours)

Q: A juggler can juggle 16 balls. Half of the balls are golf balls, and half of the golf balls are blue. How many blue golf balls are there?

A: **Let's think step by step.**

(Output) There are 16 balls in total. Half of the balls are golf balls. That means that there are 8 golf balls. Half of the golf balls are blue. That means that there are 4 blue golf balls. ✓

Guidelines of Crating Effective Prompts


- 1.Be clear and concise:** easy to understand, provide enough information for the model to generate relevant output.
- 2.Use specific examples:** Providing specific examples to help the model better understand expected output.
- 3.Vary the prompts:** Try different prompts to produce more diverse and creative output..
- 4.Test and refine:** Refine the prompts according to the outputs of the tested prompts
- 5.Use feedback:** Use feedback from users or other sources to continually improve your prompts

OpenAI o1

- OpenAI o1: a new large language model trained with reinforcement learning to perform complex reasoning. o1 thinks before it answers—it can produce a long internal chain of thought before responding to the user.

Performance w.r.t. train/test-time compute

- o1 performance smoothly improves with both train-time and test-time compute



Prompt Engineering with Llama 3

- https://github.com/amitsangani/Llama/blob/main/Llama_3_Prompt_Engineering.ipynb

The overall training process of ChatGPT


Step 1

Collect demonstration data, and train a supervised policy.

A prompt is sampled from our prompt dataset.

A labeler demonstrates the desired output behavior.

This data is used to fine-tune GPT-3 with supervised learning.




Step 2

Collect comparison data, and train a reward model.

A prompt and several model outputs are sampled.

A labeler ranks the outputs from best to worst.

This data is used to train our reward model.



Step 3


Optimize a policy against the reward model using reinforcement learning.

A new prompt is sampled from the dataset.

The policy generates an output.

The reward model calculates a reward for the output.

The reward is used to update the policy using PPO.



Thanks!