

Language models

MIT



Jacob Andreas (many slides by Chris Tanner & Yoon Kim)

Can you please generate an NLP lecture plan concerning prompting, chain-of-thought, and instruction-tuning? It should be approximately 80 minutes in duration.



Here's an NLP lecture plan for an 80-minute session on **Prompting, Chain-of-Thought, and Instruction-Tuning**. This plan assumes your audience has a basic understanding of machine learning and natural language processing, and the aim is to explore these advanced topics through a structured progression.

Lecture Plan: Prompting, Chain-of-Thought, and Instruction-Tuning

1. Introduction to NLP in the Context of Prompts (10 minutes)

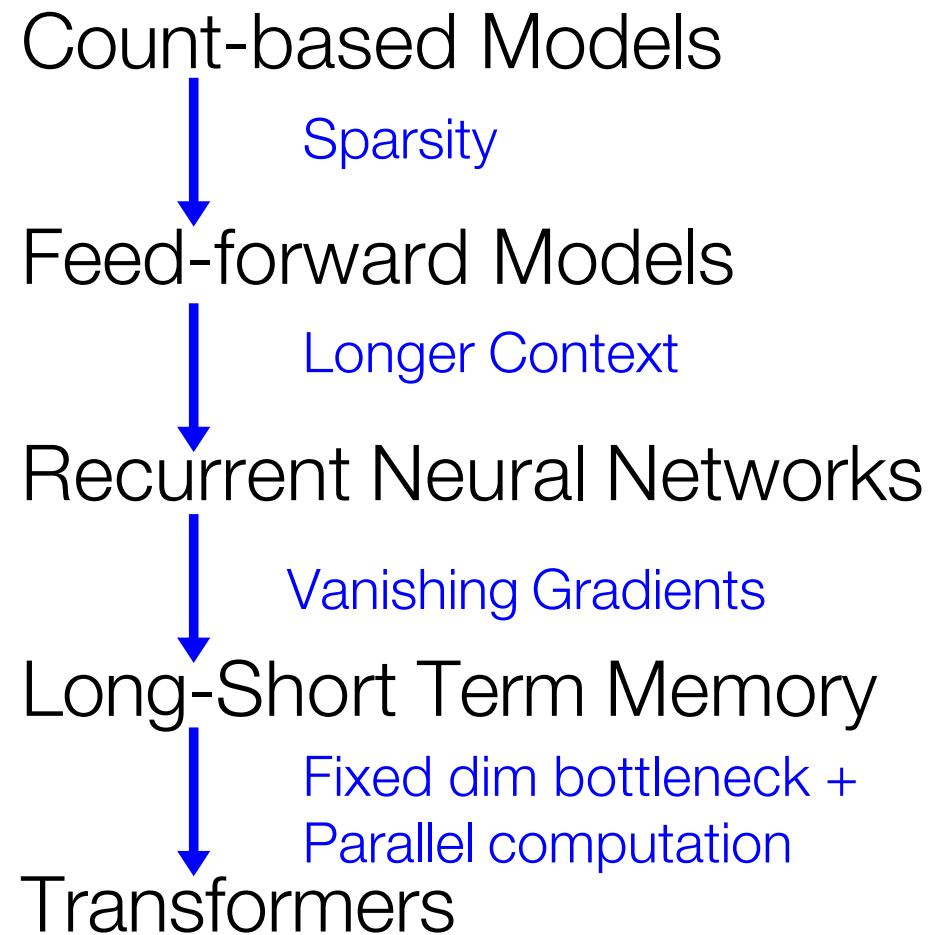
- **Objective:** Provide a high-level overview of prompting and its significance in NLP.
 - **Overview of Prompting:**
 - Define prompts in NLP: Instructing a language model to perform a task through natural language inputs.
 - **Importance:** How prompting shifts NLP from task-specific models to versatile, general models (e.g., GPT models).
 - **Types of Prompts:**
 - Zero-shot prompting
 - Few-shot prompting
 - Direct prompts vs. instructive prompts
 - **Examples:**

Outline

-  **Review**
-  **Prompting**
 -  **In-Context**
 -  **Chain-of-Thought**
-  **Instruction Tuning**

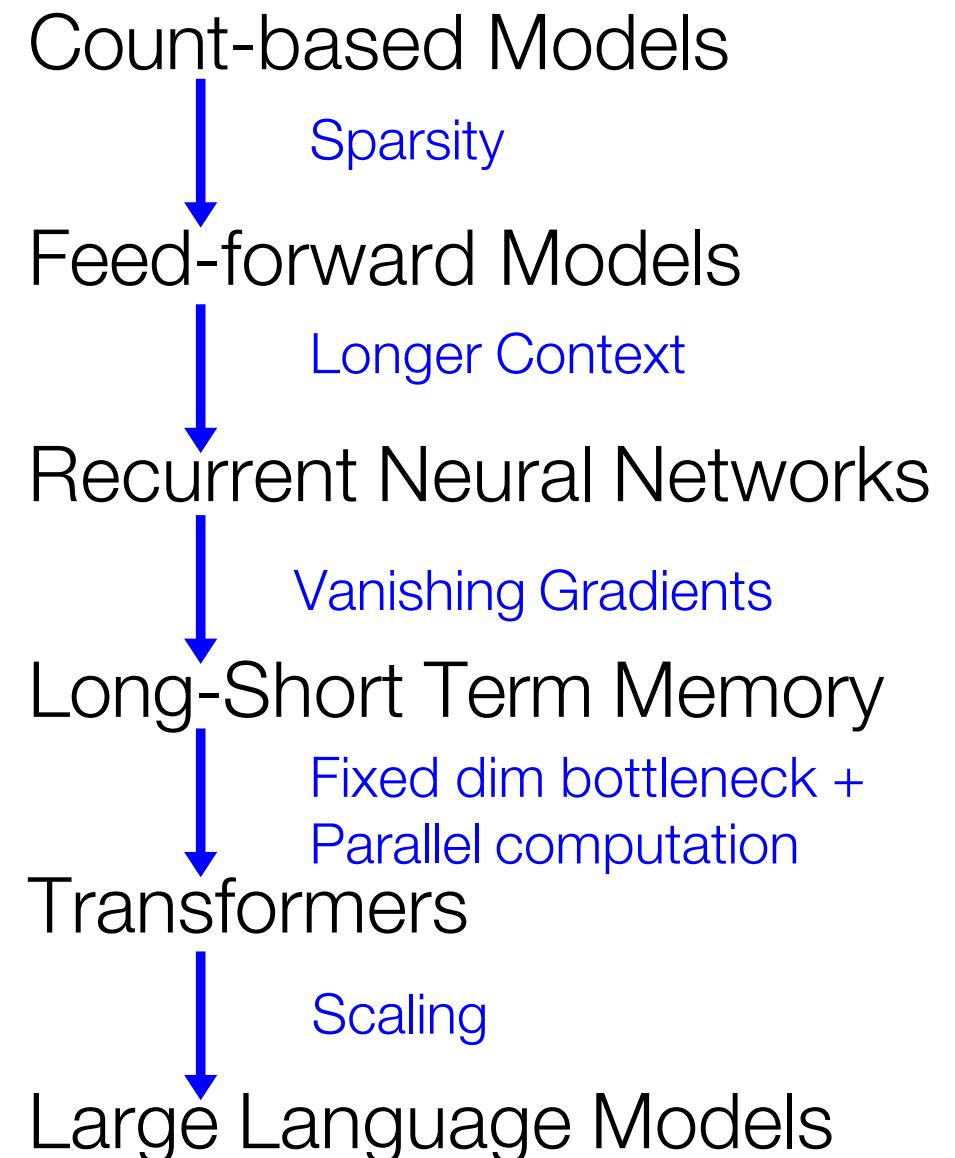
The Path to LLMs

Model	Perplexity	
Unigram	≈ 300	
3-gram	≈ 150	
5-gram	141.2	
Neural 5-gram	126.0	2000
RNN	124.7	2010
Deep RNN	107.5	2012
LSTM	78.4	2014
LSTM + Hyperparameter tuning	58.3	2016
Transformer	54.5	2017



The Path to LLMs

Model	Perplexity	
Unigram	≈ 300	
3-gram	≈ 150	
5-gram	141.2	
Neural 5-gram	126.0	2000
RNN	124.7	2010
Deep RNN	107.5	2012
LSTM	78.4	2014
LSTM + Hyperparameter tuning	58.3	2016
Transformer	54.5	2017
GPT2	35.8	2019
GPT3	20.5	2020



The Path to LLMs

Large Language Models in Machine Translation (2007)

Thorsten Brants Ashok C. Popat Peng Xu Franz J. Och Jeffrey Dean

This paper reports on the benefits of large-scale statistical language modeling in machine translation. A distributed infrastructure is proposed which we use to train on up to 2 trillion tokens, resulting in language models having up to 300 billion n -grams. It is capable of providing smoothed probabilities for fast, single-pass decoding. We introduce a new smoothing method, dubbed *Stupid Backoff*, that is inexpensive to train on large data sets and approaches the quality of Kneser-Ney Smoothing as the amount of training data increases.

Model	Size (# Parameters)	Training Tokens
LaMDA (Thoppilan et al., 2022)	137 Billion	168 Billion
GPT-3 (Brown et al., 2020)	175 Billion	300 Billion
Jurassic (Lieber et al., 2021)	178 Billion	300 Billion
Gopher (Rae et al., 2021)	280 Billion	300 Billion
MT-NLG 530B (Smith et al., 2022)	530 Billion	270 Billion
<i>Chinchilla</i>	70 Billion	1.4 Trillion

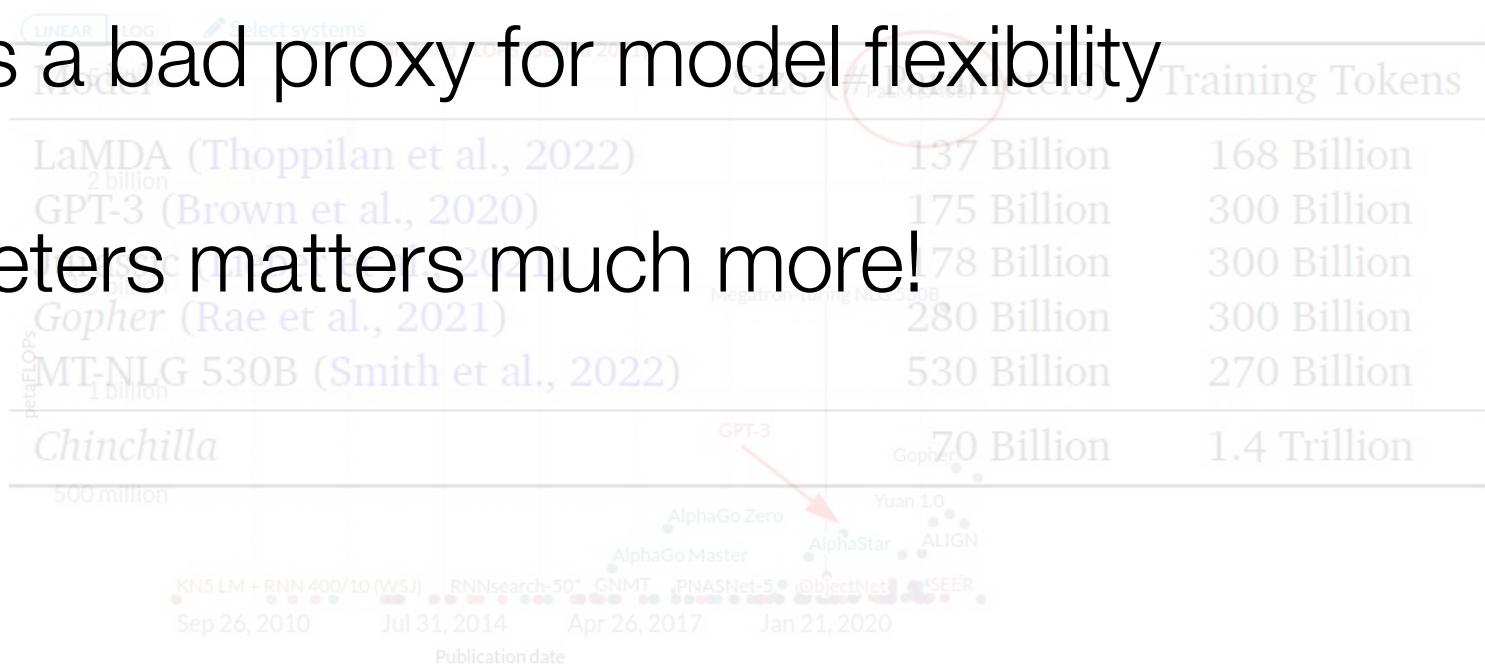
The Path to LLMs

Large Language Models in Machine Translation

Thorsten Brants Ashok C. Popat Peng Xu Franz J. Och Jeffrey Dean

Number of parameters is a bad proxy for model flexibility

scale statistical language modeling in machine translation. A distributed infrastructure is proposed which we use to train on models having up to 300 billion n -grams. It is capable of providing smoothed probabilities for fast, single-pass decoding. We introduce a new smoothing method, dubbed *Stupid Backoff*, that is inexpensive to train on large data sets and approaches the quality of Kneser-Ney Smoothing as the amount of training data increases.



The Path to LLMs

Pipelined NLP

1970–2014

Language model as part of noisy-channel model.

$$\Pr(\mathbf{Y} = \mathbf{y} \mid \mathbf{X} = \mathbf{x}) \propto$$

$$\Pr(\mathbf{X} = \mathbf{x} \mid \mathbf{Y} = \mathbf{y}) \Pr(\mathbf{Y} = \mathbf{y})$$

Language model

Applying existing architectures to LM:

- MLP
- RNN
- LSTM

The Path to LLMs

Pipelined NLP

1970–2014

Language model as part of noisy-channel model.

$$\Pr(\mathbf{Y} = \mathbf{y} \mid \mathbf{X} = \mathbf{x}) \propto$$

$$\Pr(\mathbf{X} = \mathbf{x} \mid \mathbf{Y} = \mathbf{y}) \Pr(\mathbf{Y} = \mathbf{y})$$

Applying existing architectures to LM:

- MLP
- RNN
- LSTM

seq2seq learning

2014–2017

(Conditional) language modeling for sequence prediction.

$$\Pr(\mathbf{Y} = \mathbf{y} \mid \mathbf{X} = \mathbf{x}) = \prod_{t=1}^T \Pr(Y_t = y_t \mid \mathbf{X} = \mathbf{x}, \mathbf{Y}_{<t} = \mathbf{y}_{<t})$$

New architectures for seq2seq:

- GRU
- Attention
- Transformer

Autoregressive conditional language model

The Path to LLMs

Pipelined NLP

1970–2014

Language model as part of noisy-channel model.

$$\Pr(\mathbf{Y} = \mathbf{y} \mid \mathbf{X} = \mathbf{x}) \propto$$

$$\Pr(\mathbf{X} = \mathbf{x} \mid \mathbf{Y} = \mathbf{y}) \Pr(\mathbf{Y} = \mathbf{y})$$

Applying existing architectures to LM:

- MLP
- RNN
- LSTM

seq2seq learning

2014–2017

(Conditional) language modeling for sequence prediction.

$$\Pr(\mathbf{Y} = \mathbf{y} \mid \mathbf{X} = \mathbf{x}) = \prod_{t=1}^T \Pr(Y_t = y_t \mid \mathbf{X} = \mathbf{x}, \mathbf{Y}_{<t} = \mathbf{y}_{<t})$$

New architectures for seq2seq:

- GRU
- Attention
- Transformer

Pretrain + Transfer

2017–2020

Language model (like) objectives for representation learning.

$$\Pr(X_t = x_t \mid \mathbf{X}_{<t} = \mathbf{x}_{<t}, \mathbf{X}_{>t} = \mathbf{x}_{>t})$$
$$\arg \max_{\theta} \log \Pr(Y = y \mid \mathbf{X} = \mathbf{x})$$

Big Transformers “language models”:

- Masked LMs
- Probing
- Finetuning

Pretrain as a “language model”, then finetune

The Path to LLMs

Pipelined NLP

1970–2014

Language model as part of noisy-channel model.

$$\Pr(\mathbf{Y} = \mathbf{y} \mid \mathbf{X} = \mathbf{x}) \propto$$

$$\Pr(\mathbf{X} = \mathbf{x} \mid \mathbf{Y} = \mathbf{y}) \Pr(\mathbf{Y} = \mathbf{y})$$

Applying existing architectures to LM:

- MLP
- RNN
- LSTM

seq2seq learning

2014–2017

(Conditional) language modeling for sequence prediction.

$$\Pr(\mathbf{Y} = \mathbf{y} \mid \mathbf{X} = \mathbf{x}) = \prod_{t=1}^T \Pr(Y_t = y_t \mid \mathbf{X} = \mathbf{x}, \mathbf{Y}_{<t} = \mathbf{y}_{<t})$$

New architectures for seq2seq:

- GRU
- Attention
- Transformer

Pretrain + Transfer

2017–2020

Language model (like) objectives for representation learning.

$$\Pr(X_t = x_t \mid \mathbf{X}_{<t} = \mathbf{x}_{<t}, \mathbf{X}_{>t} = \mathbf{x}_{>t}) \arg \max_{\theta} \log \Pr(Y = y \mid \mathbf{X} = \mathbf{x})$$

Big Transformers “language models”:

- Masked LMs
- Probing
- Finetuning

Pretrain + Prompt

2020–

A single language model for everything

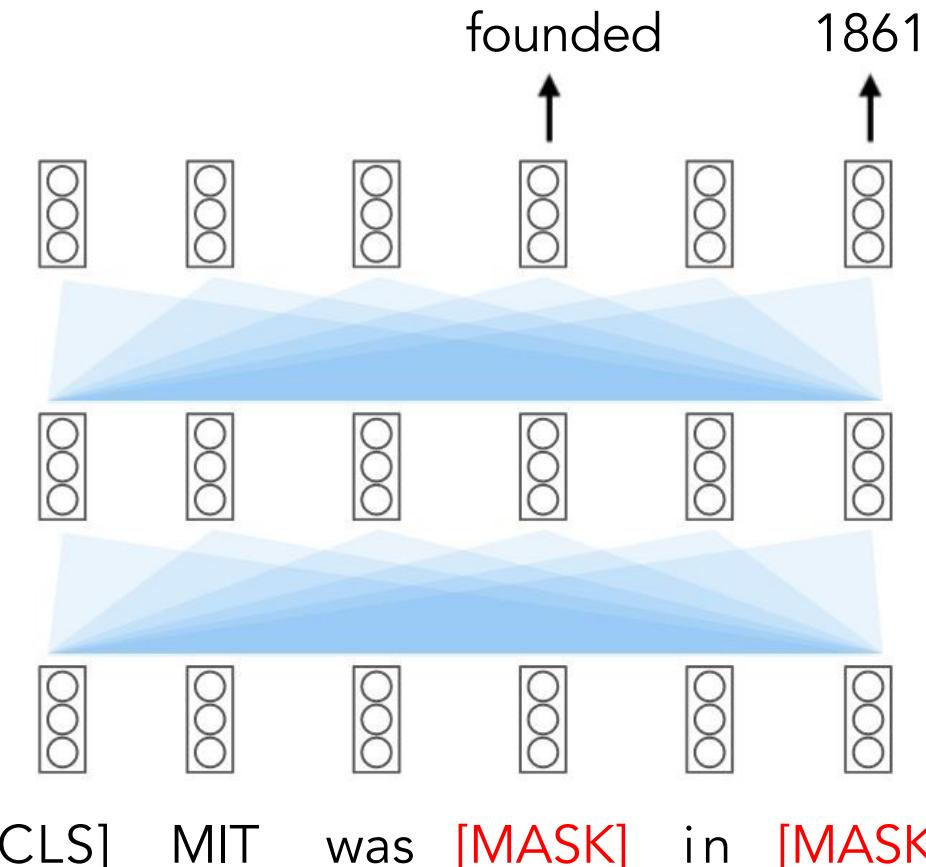
$$\Pr(\mathbf{X} = \mathbf{x})$$

BIG Transformer language models:

- Prompting
- Exploration of capabilities

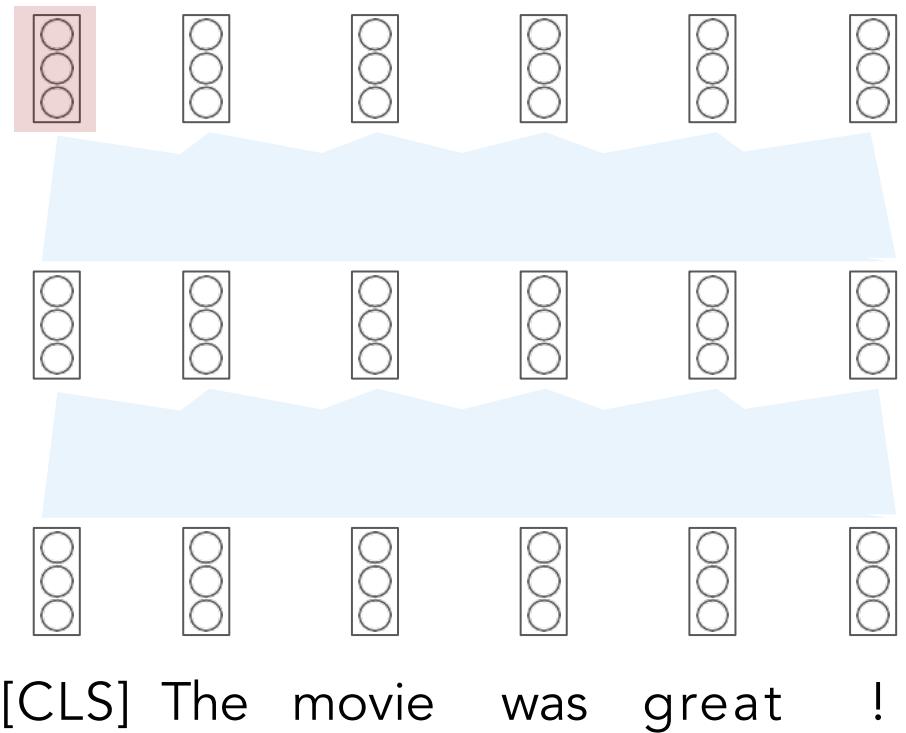
BERT Review

Pretrain as a masked language model



Fine-tune part (or all) of the network for downstream tasks

$$p(y | x) = \text{softmax}(\mathbf{W} \mathbf{h}_{[\text{CLS}]} + \mathbf{b})_y$$



BERT Review

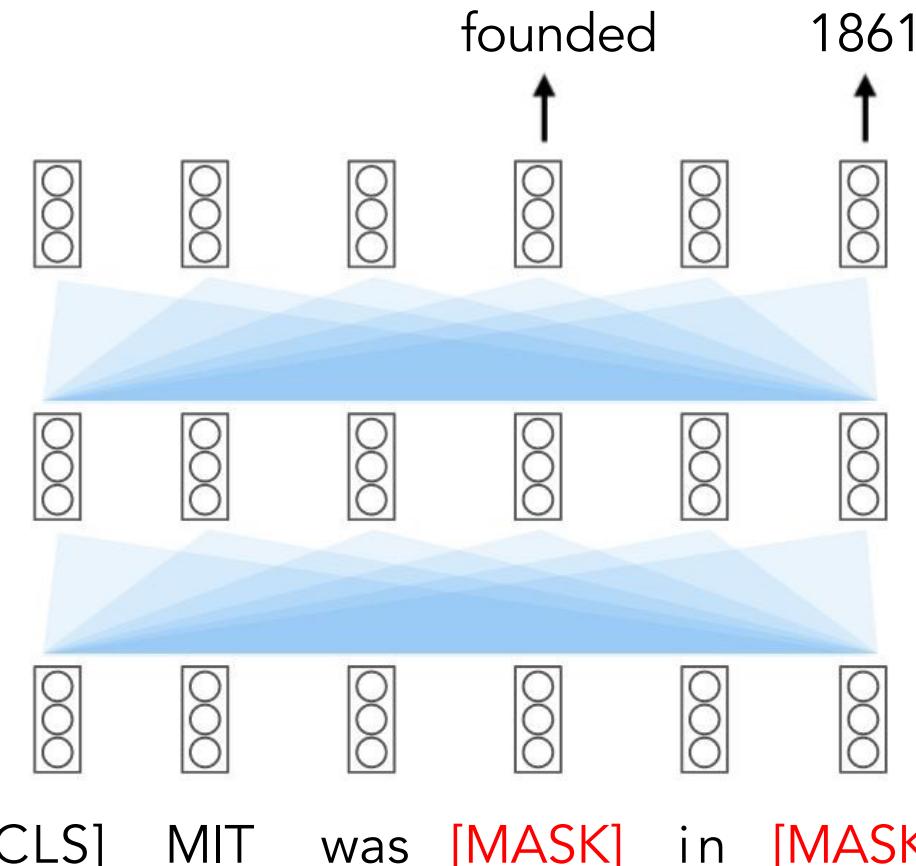
Results were remarkably strong

System	MNLI-(m/mm) 392k	QQP 363k	QNLI 108k	SST-2 67k	CoLA 8.5k	STS-B 5.7k	MRPC 3.5k	RTE 2.5k	Average -
Pre-OpenAI SOTA	80.6/80.1	66.1	82.3	93.2	35.0	81.0	86.0	61.7	74.0
BiLSTM+ELMo+Attn	76.4/76.1	64.8	79.8	90.4	36.0	73.3	84.9	56.8	71.0
OpenAI GPT	82.1/81.4	70.3	87.4	91.3	45.4	80.0	82.3	56.0	75.1
BERT _{BASE}	84.6/83.4	71.2	90.5	93.5	52.1	85.8	88.9	66.4	79.6
BERT _{LARGE}	86.7/85.9	72.1	92.7	94.9	60.5	86.5	89.3	70.1	82.1

Why does this work?

BERT Review

Pretrain as a masked language model



Predicting words in context requires linguistic & world knowledge!

Such knowledge is implicitly captured within the Transformer's parameters.
(Distributional hypothesis)

Recall that a language model estimates the probability of any sequence of words \mathbf{x} (e.g., a sentence).

$$p(\text{I'm not a cat}) = 0.0000004$$

$$p(\text{He is hungry}) = 0.000025$$

$$p(\text{Dog the asd@sdf 1124 !?}) \approx 0$$

$$p\left(\begin{array}{c} \text{The GREAT} \\ \text{GATSBY} \\ \text{F. SCOTT FITZGERALD} \end{array}\right) = ?$$

Thus, BERT is NOT a language model.

Nowadays, “large language model” \approx large Transformer
pre-trained with a self-supervised objective
(next-word prediction)

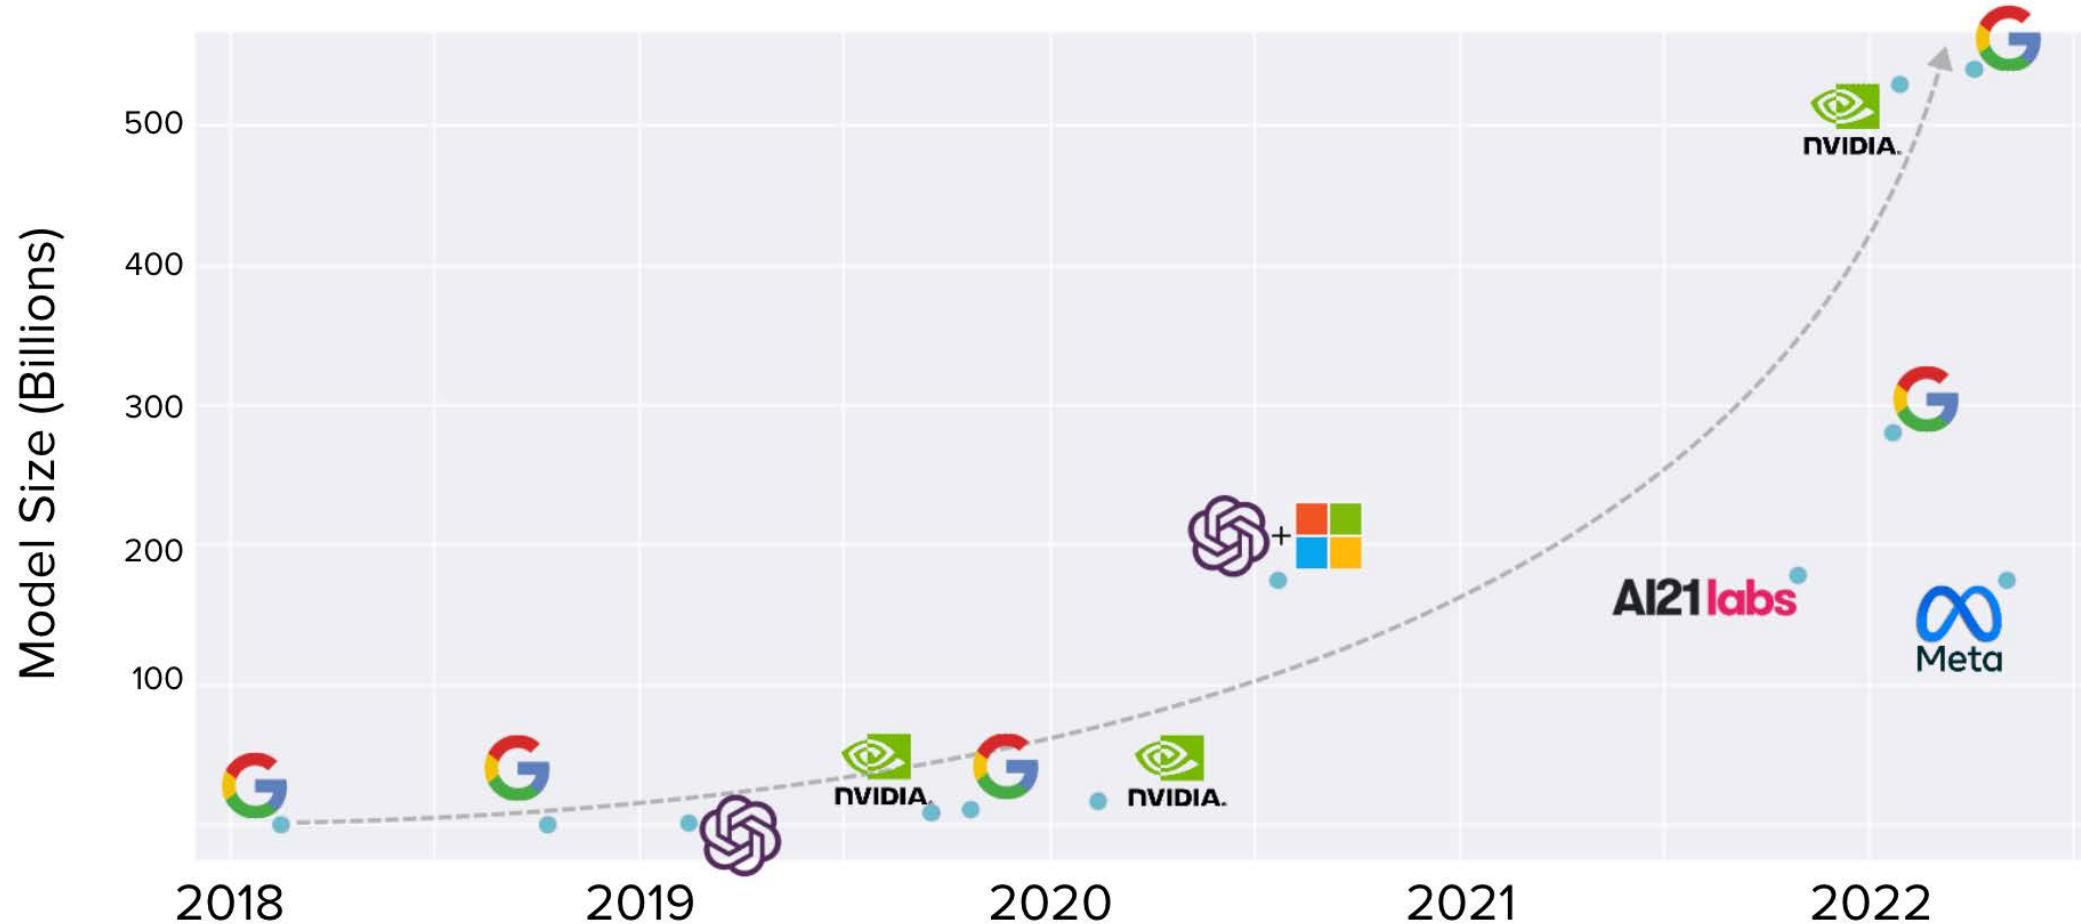
So far, we've used the masked language modelling objective to pre-train Transformers for **representation learning**

These **learned representations** were highly effective for downstream tasks.

Yet, this pretrain-then-finetune paradigm still requires non-trivial amounts of labelled data (e.g., hundreds of samples)

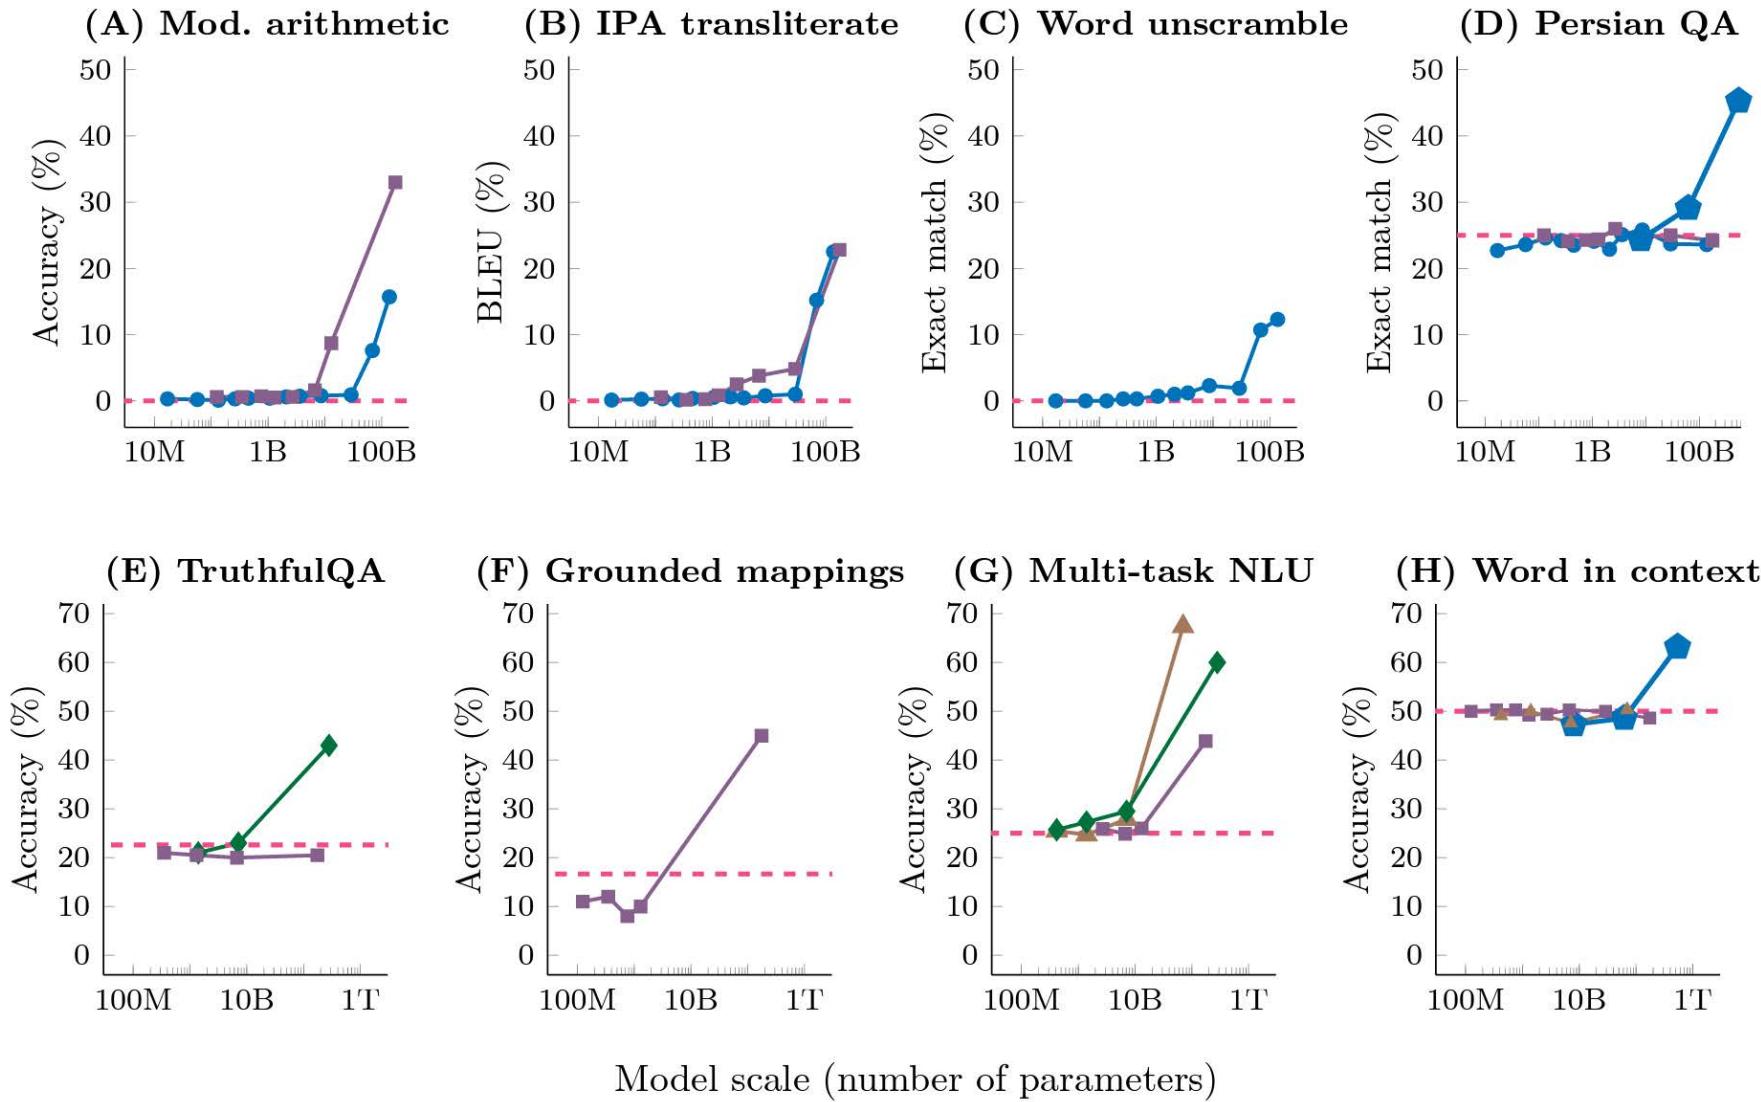
LLMs Review

Developing larger and larger language models has led to an arms race between technology companies



Emergent Abilities of LLMs

—●— LaMDA —■— GPT-3 —◆— Gopher —▲— Chinchilla —●— PaLM —--- Random



Few-shot “In-Context Learning” with LLMs

Language Models are Few-Shot Learners

The GPT-3
paper

Tom B. Brown*

Benjamin Mann*

Nick Ryder*

Melanie Subbiah*

Jared Kaplan[†]

Prafulla Dhariwal

Arvind Neelakantan

Pranav Shyam

Girish Sastry

Amanda Askell

Sandhini Agarwal

Ariel Herbert-Voss

Gretchen Krueger

Tom Henighan

Rewon Child

Aditya Ramesh

Daniel M. Ziegler

Jeffrey Wu

Clemens Winter

Christopher Hesse

Mark Chen

Eric Sigler

Mateusz Litwin

Scott Gray

Benjamin Chess

Jack Clark

Christopher Berner

Sam McCandlish

Alec Radford

Ilya Sutskever

Dario Amodei

OpenAI

[Brown et al. '20]

Outline

-  **Review**
-  **Prompting**
 -  **In-Context**
 -  **Chain-of-Thought**
-  **Instruction Tuning**

Outline

 Review

 Prompting

 In-Context

 Chain-of-Thought

 Instruction Tuning

Prompting

$$P(\cdot | x_1, \dots, x_t)$$

How can we fiddle with the **context** that we give to language models to get them to do what we want?

“In-Context Learning” with LMs

Input “demonstration data” as history

⇒ LM learns to infer what the task is from exemplars

Training labels : $\{(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})\}_{i=1}^N$

Test input : \mathbf{x}

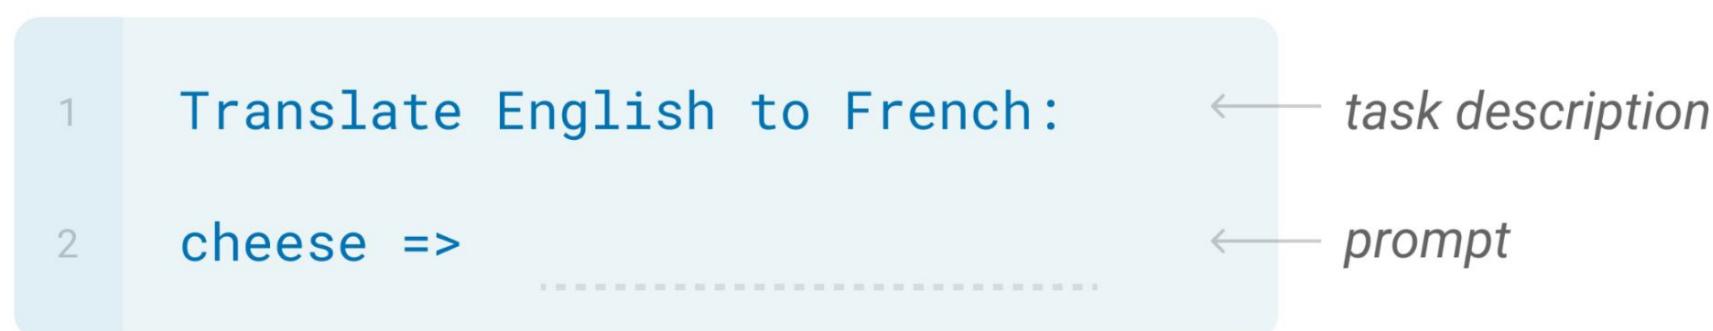
Prompted LM : $P(\cdot | \mathbf{x}^{(1)}, \mathbf{y}^{(1)}, \dots, \mathbf{x}^{(N)}, \mathbf{y}^{(N)}, \mathbf{x})$

N is typically 1-32

“In-Context Learning” with LMs

Zero-shot

The model predicts the answer given only a natural language description of the task. No gradient updates are performed.



© Brown et al. All rights reserved. This content is excluded from our Creative Commons license. For more information, see <https://ocw.mit.edu/help/faq-fair-use/>

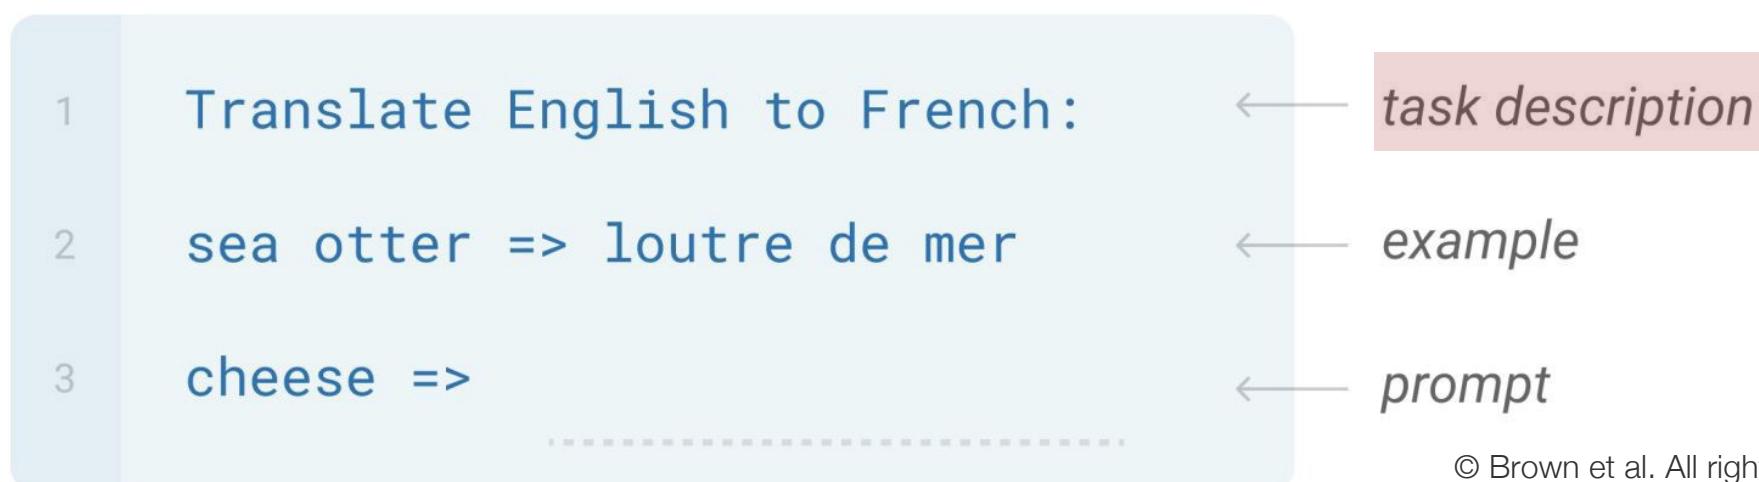
[Brown et al. '20]

“In-Context Learning” with LMs

One-shot

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

Task description
optional if given
exemplars



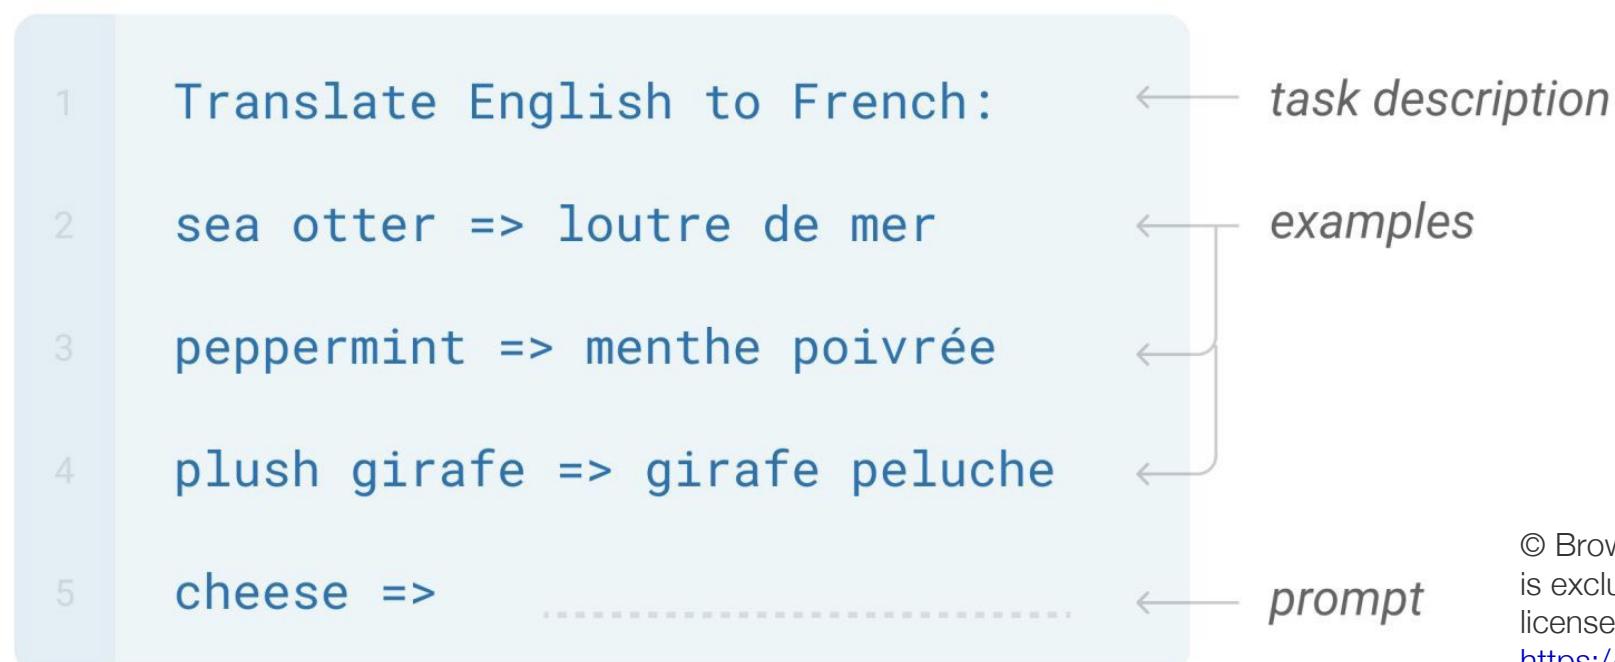
© Brown et al. All rights reserved. This content is excluded from our Creative Commons license. For more information, see <https://ocw.mit.edu/help/faq-fair-use/>

[Brown et al. '20]

“In-Context Learning” with LMs

Few-shot

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



© Brown et al. All rights reserved. This content is excluded from our Creative Commons license. For more information, see <https://ocw.mit.edu/help/faq-fair-use/>

[Brown et al. '20]

“In-Context Learning” with LMs

Review: the whole thing 's fairly lame , making it par for the course for disney sequels .

Answer: Negative

Review: this quiet , introspective and entertaining independent is worth seeking .

Answer: Positive

Review: this quiet , introspective and entertaining independent is worth seeking .

Answer:

$\mathbf{x}^{(1)}$

$\mathbf{y}^{(1)}$

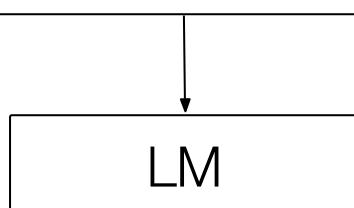
$\mathbf{x}^{(2)}$

$\mathbf{y}^{(2)}$

\mathbf{x}

Training examples
(two-shot)

Test example



$$P(\cdot | \mathbf{x}^{(1)}, \mathbf{y}^{(1)}, \dots, \mathbf{x}^{(N)}, \mathbf{y}^{(N)}, \mathbf{x})$$

What are possible issues
with In-Context Learning?

“In-Context Learning” with LMs

Review: the whole thing 's fairly lame , making it par for the course for disney sequels .

Answer: Negative

Review: this quiet , introspective and entertaining independent is worth seeking .

Answer: Positive

Review: this quiet , introspective and entertaining independent is worth seeking .

Answer:

$\mathbf{x}^{(1)}$

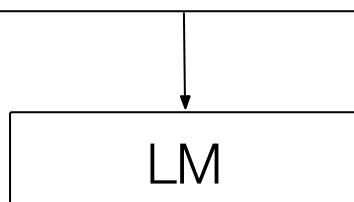
$\mathbf{y}^{(1)}$

$\mathbf{x}^{(2)}$

$\mathbf{y}^{(2)}$

\mathbf{x}

Need to turn labels into natural language



$$P(\cdot | \mathbf{x}^{(1)}, \mathbf{y}^{(1)}, \dots, \mathbf{x}^{(N)}, \mathbf{y}^{(N)}, \mathbf{x})$$

“In-Context Learning” with LMs

Review: the whole thing 's fairly lame , making it par for the course for disney sequels .

Answer: Negative

Review: this quiet , introspective and entertaining independent is worth seeking .

Answer: Positive

Review: this quiet , introspective and entertaining independent is worth seeking .

Answer:

$\mathbf{x}^{(1)}$

$\mathbf{y}^{(1)}$

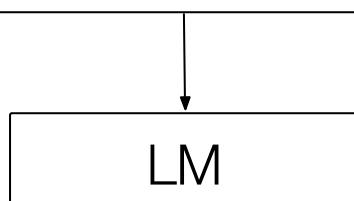
$\mathbf{x}^{(2)}$

$\mathbf{y}^{(2)}$

\mathbf{x}

Need to turn labels into natural language

Extra text on top of input/outputs



$$P(\cdot | \mathbf{x}^{(1)}, \mathbf{y}^{(1)}, \dots, \mathbf{x}^{(N)}, \mathbf{y}^{(N)}, \mathbf{x})$$

“In-Context Learning” with LMs

Review: the whole thing 's fairly lame , making it par for the course for disney sequels .

Answer: Negative

Review: this quiet , introspective and entertaining independent is worth seeking .

Answer: Positive

Review: this quiet , introspective and entertaining independent is worth seeking .

Answer:

$\mathbf{x}^{(1)}$

$\mathbf{y}^{(1)}$

$\mathbf{x}^{(2)}$

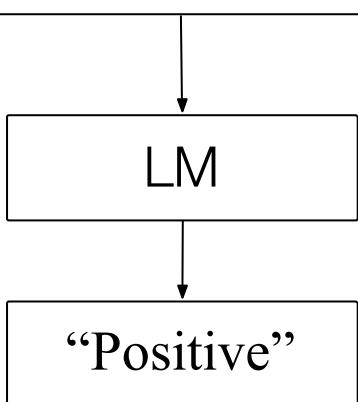
$\mathbf{y}^{(2)}$

\mathbf{x}

Need to turn labels into natural language

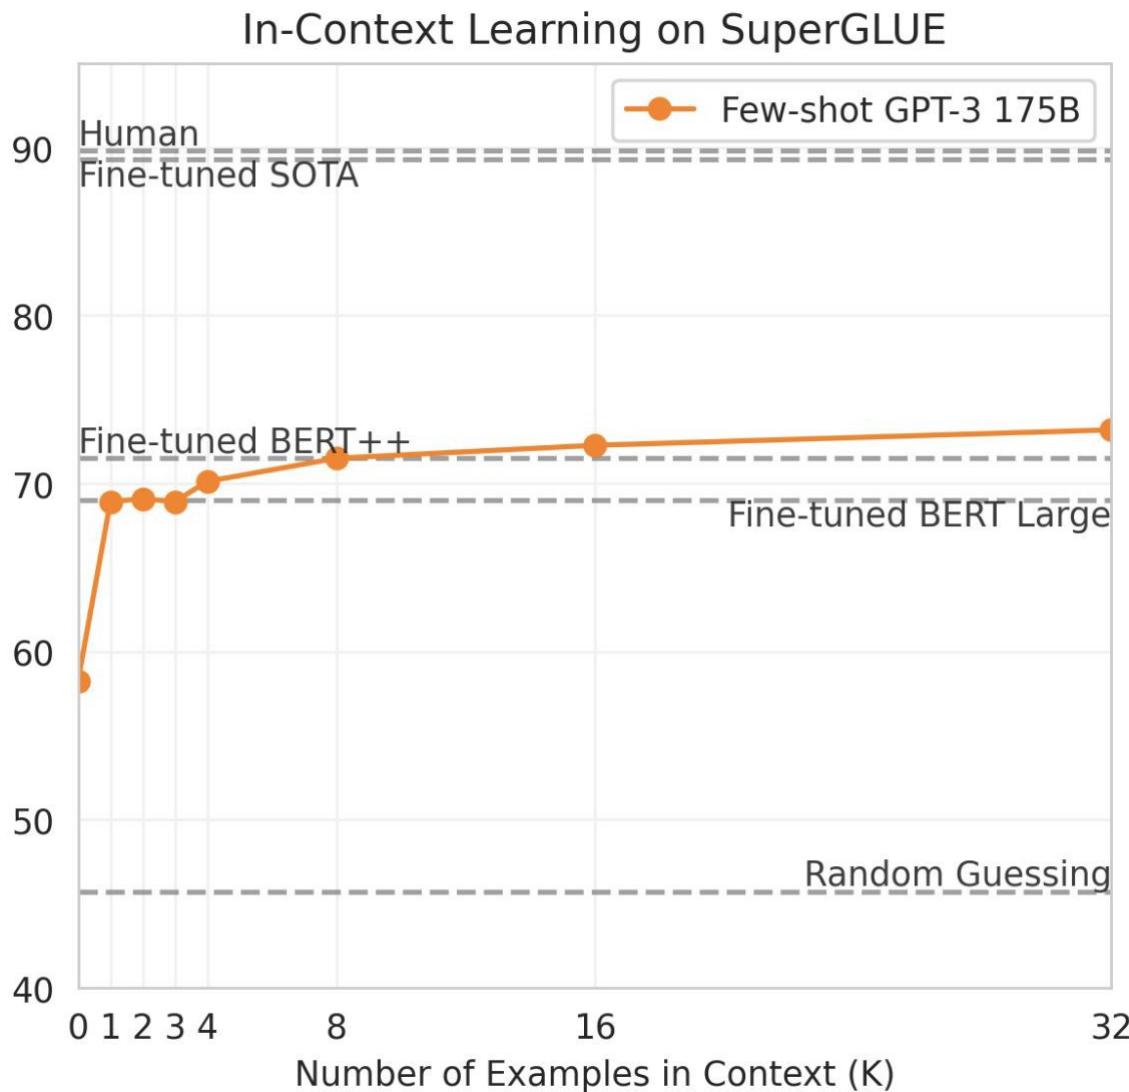
Extra text on top of input/outputs

Templates for turning demonstration data into context



$$P(\cdot \mid \text{template}(\mathbf{x}^{(1)}, \mathbf{y}^{(1)}, \dots, \mathbf{x}^{(N)}, \mathbf{y}^{(N)}, \mathbf{x}))$$

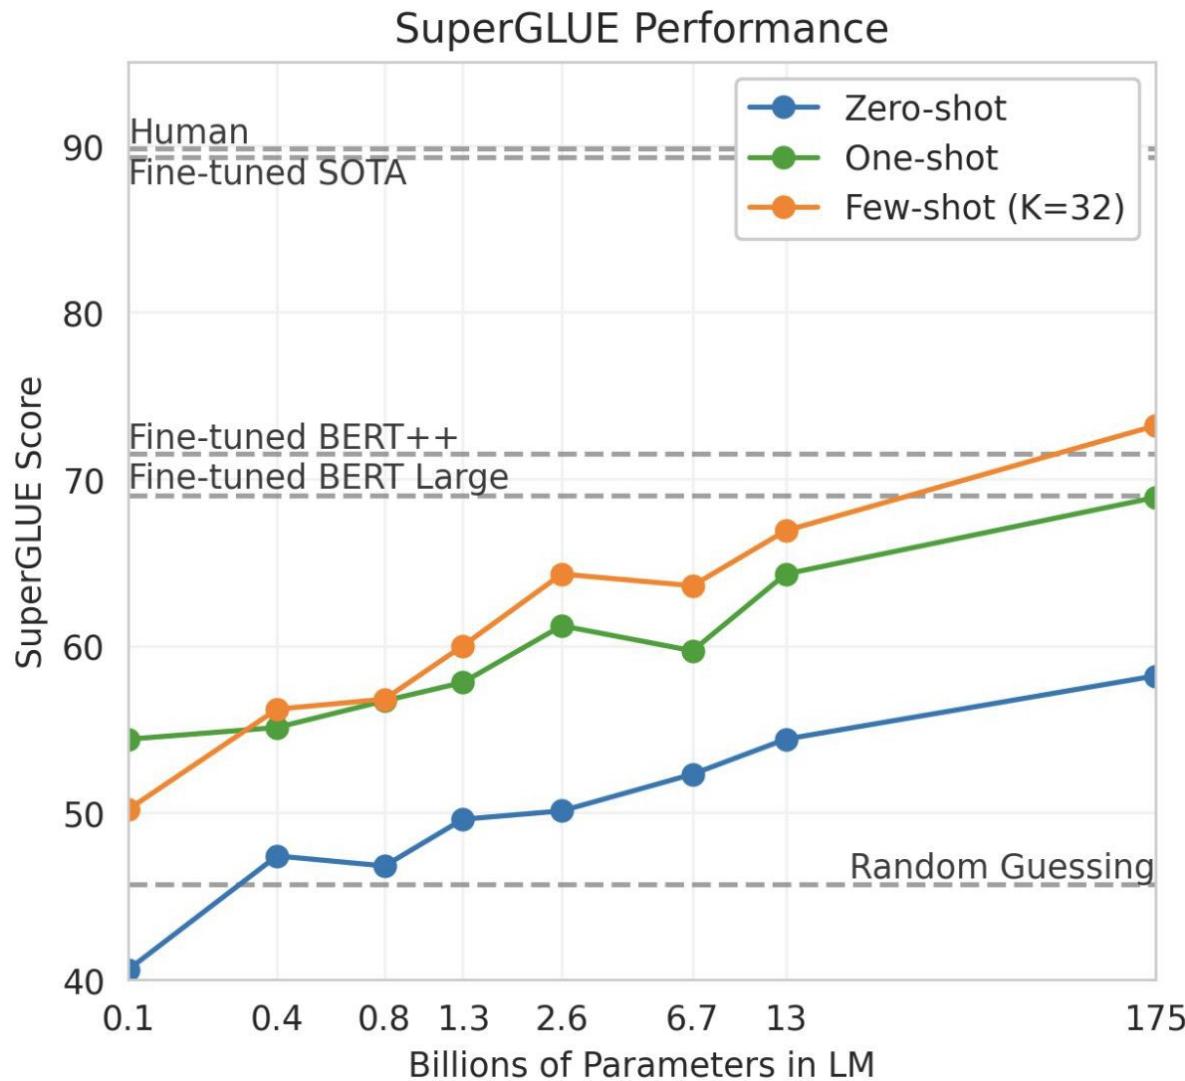
“In-Context Learning” vs Finetuning



Performance increases as the number of demonstration examples increases, but it saturates.

Still underperforms models finetuned on lots of training data.

Performance as a function of LM size



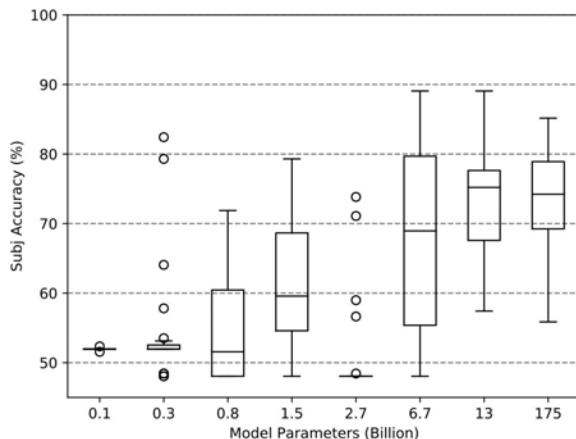
In-context learning capabilities improve with model size.

© Brown et al. All rights reserved. This content is excluded from our Creative Commons license. For more information, see <https://ocw.mit.edu/help/faq-fair-use/>

[Brown et al. '20]

In-Context Learning Issues

Sensitive to ordering
of examples?



"We demonstrate that the order in which the samples are provided can make the **difference between near state-of-the-art and random guess performance**: essentially some permutations are "fantastic" and some not. We analyse this phenomenon in detail, establishing that: it is present across model sizes (even for the largest current models)."

In-Context Learning Issues

Sensitive to Label Distribution?

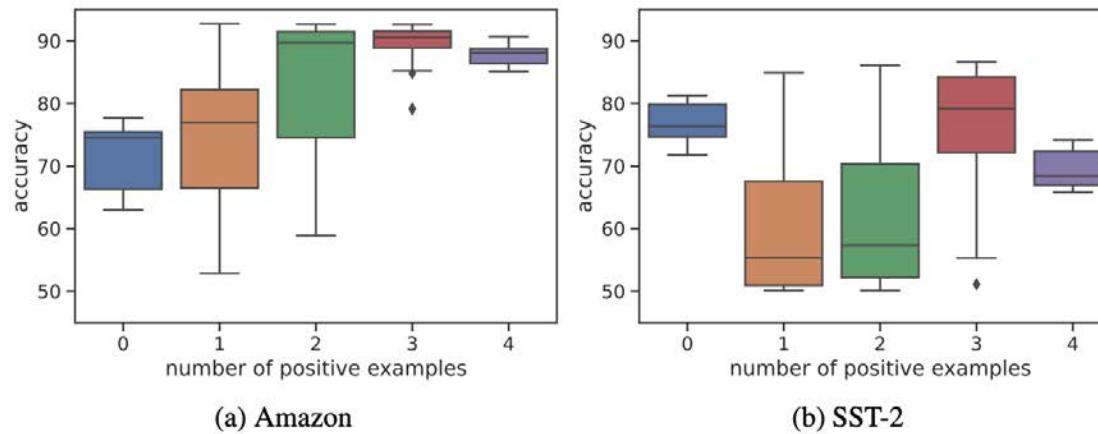


Figure 3: Accuracies of Amazon and SST-2 with varying **label balance** (number of positive examples in demonstration), across 100 total random samples of 4 demonstration examples.

Sensitive to Label Coverage?

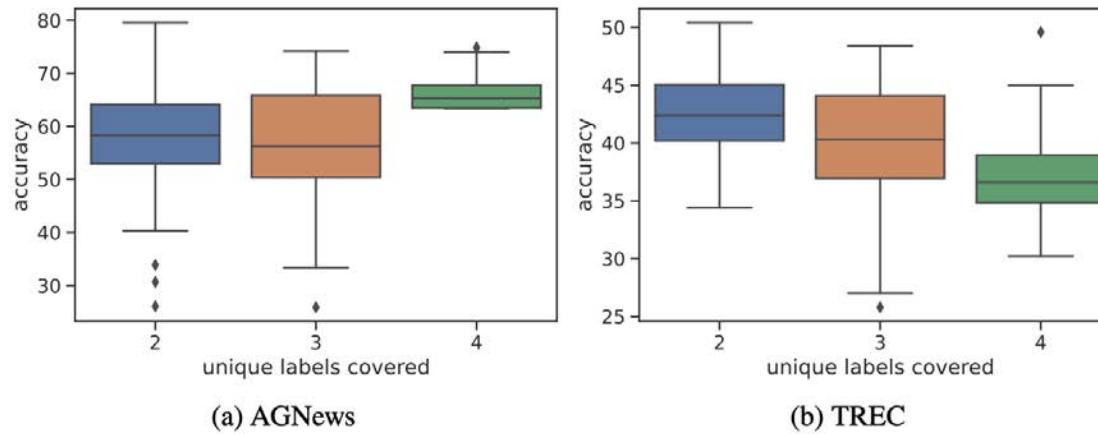


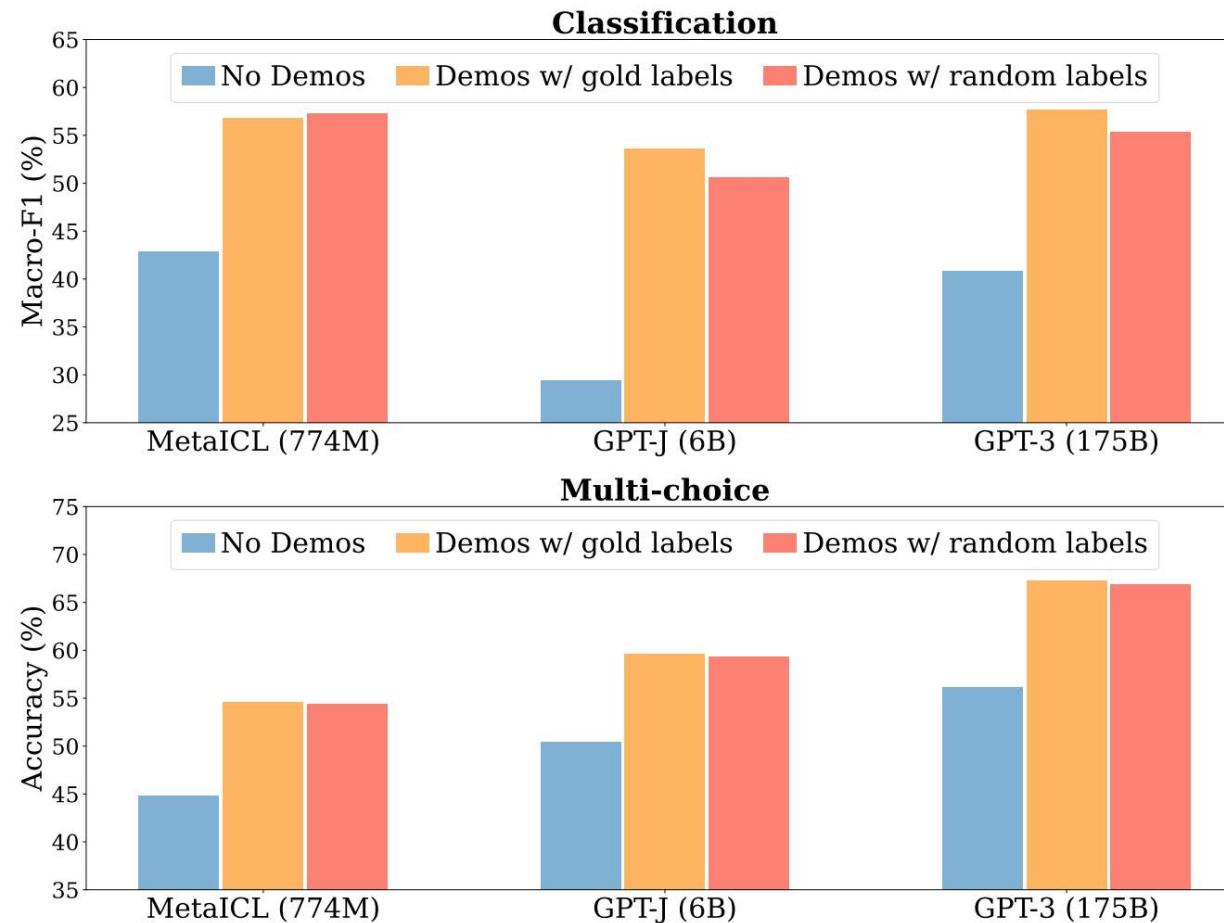
Figure 4: Accuracies of AGNews and TREC with varying **label coverage** (number of unique labels covered in demonstration), across 100 total random samples of 4 demonstration examples. Demonstration set that only covers 1 label is very unlikely and does not appear in our experiments.

Courtesy of Zhang, Feng, and Tan. Used under CC BY.

[Zhang et al., 2022](#)

In-Context Learning Issues

Correctness of the In-Context examples doesn't really matter?

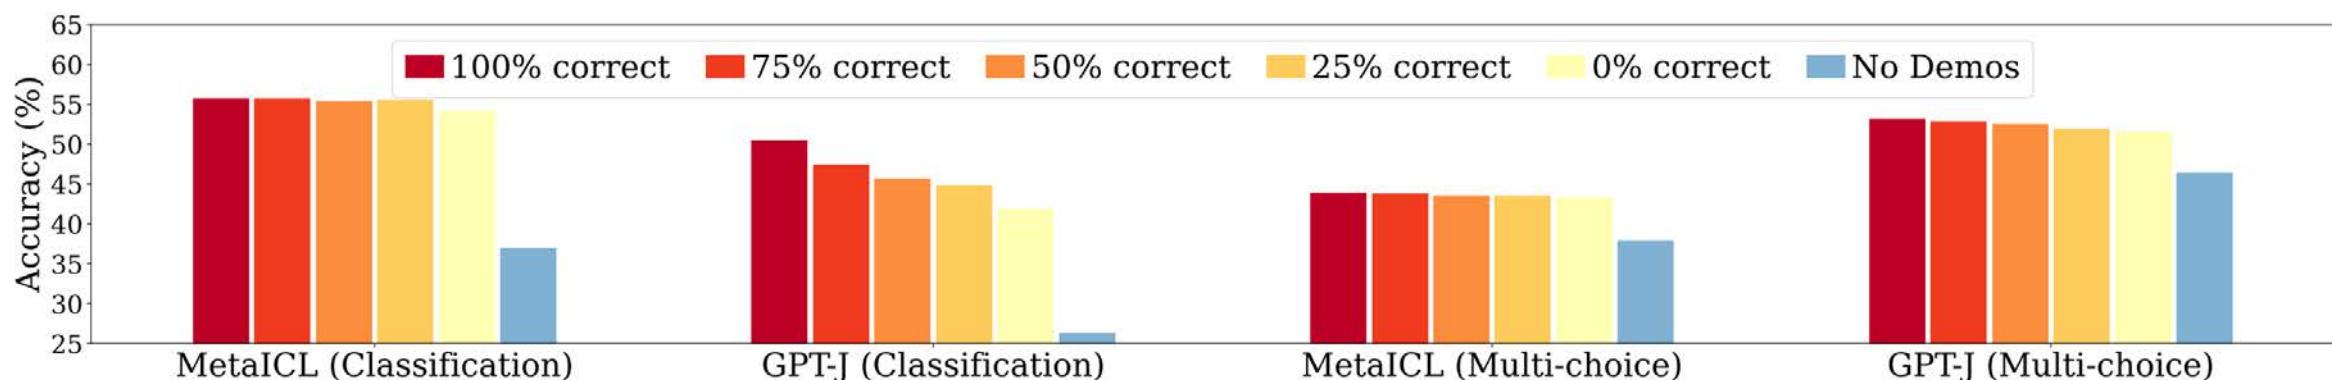


Courtesy of Min et al. Used under CC BY.

[Min et al., 2022](#)

In-Context Learning Issues

Correctness of the In-Context examples doesn't really matter?



Courtesy of Min et al. Used under CC BY.

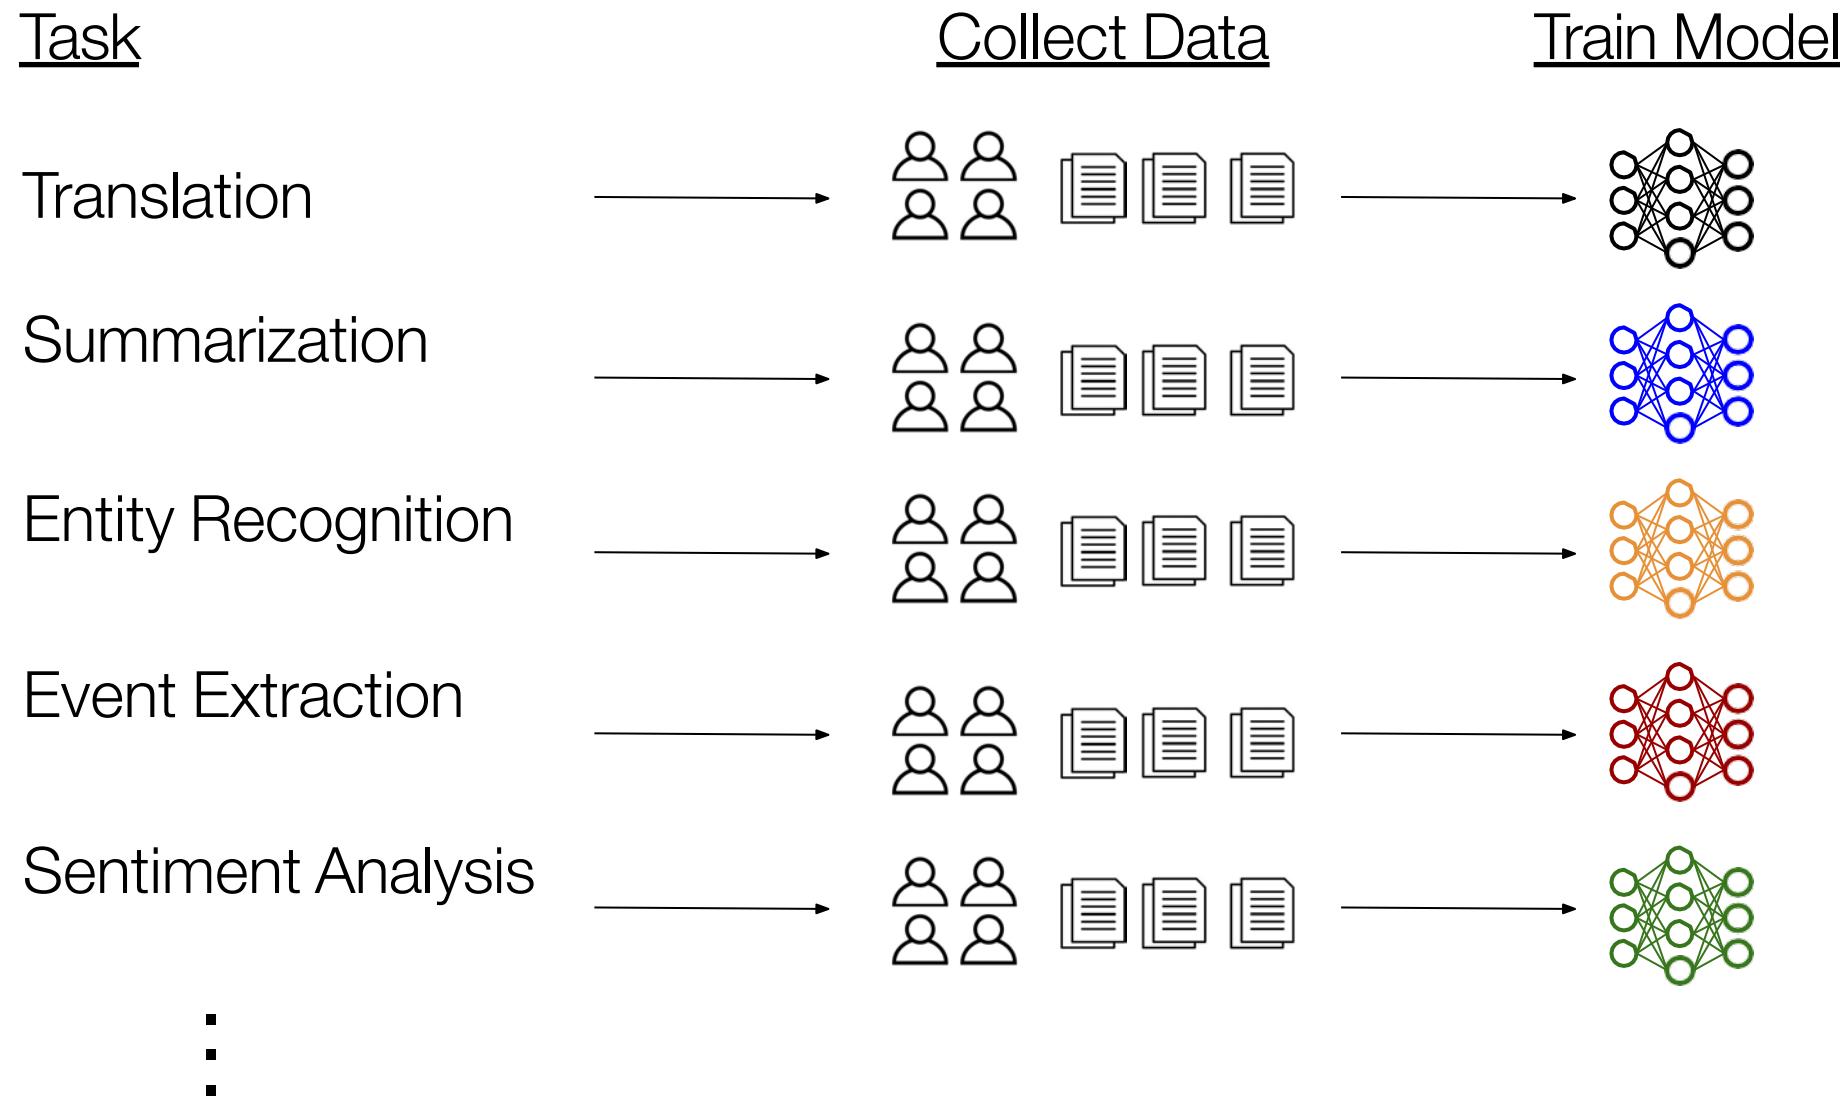
[Min et al., 2022](#)

In-Context Learning Issues

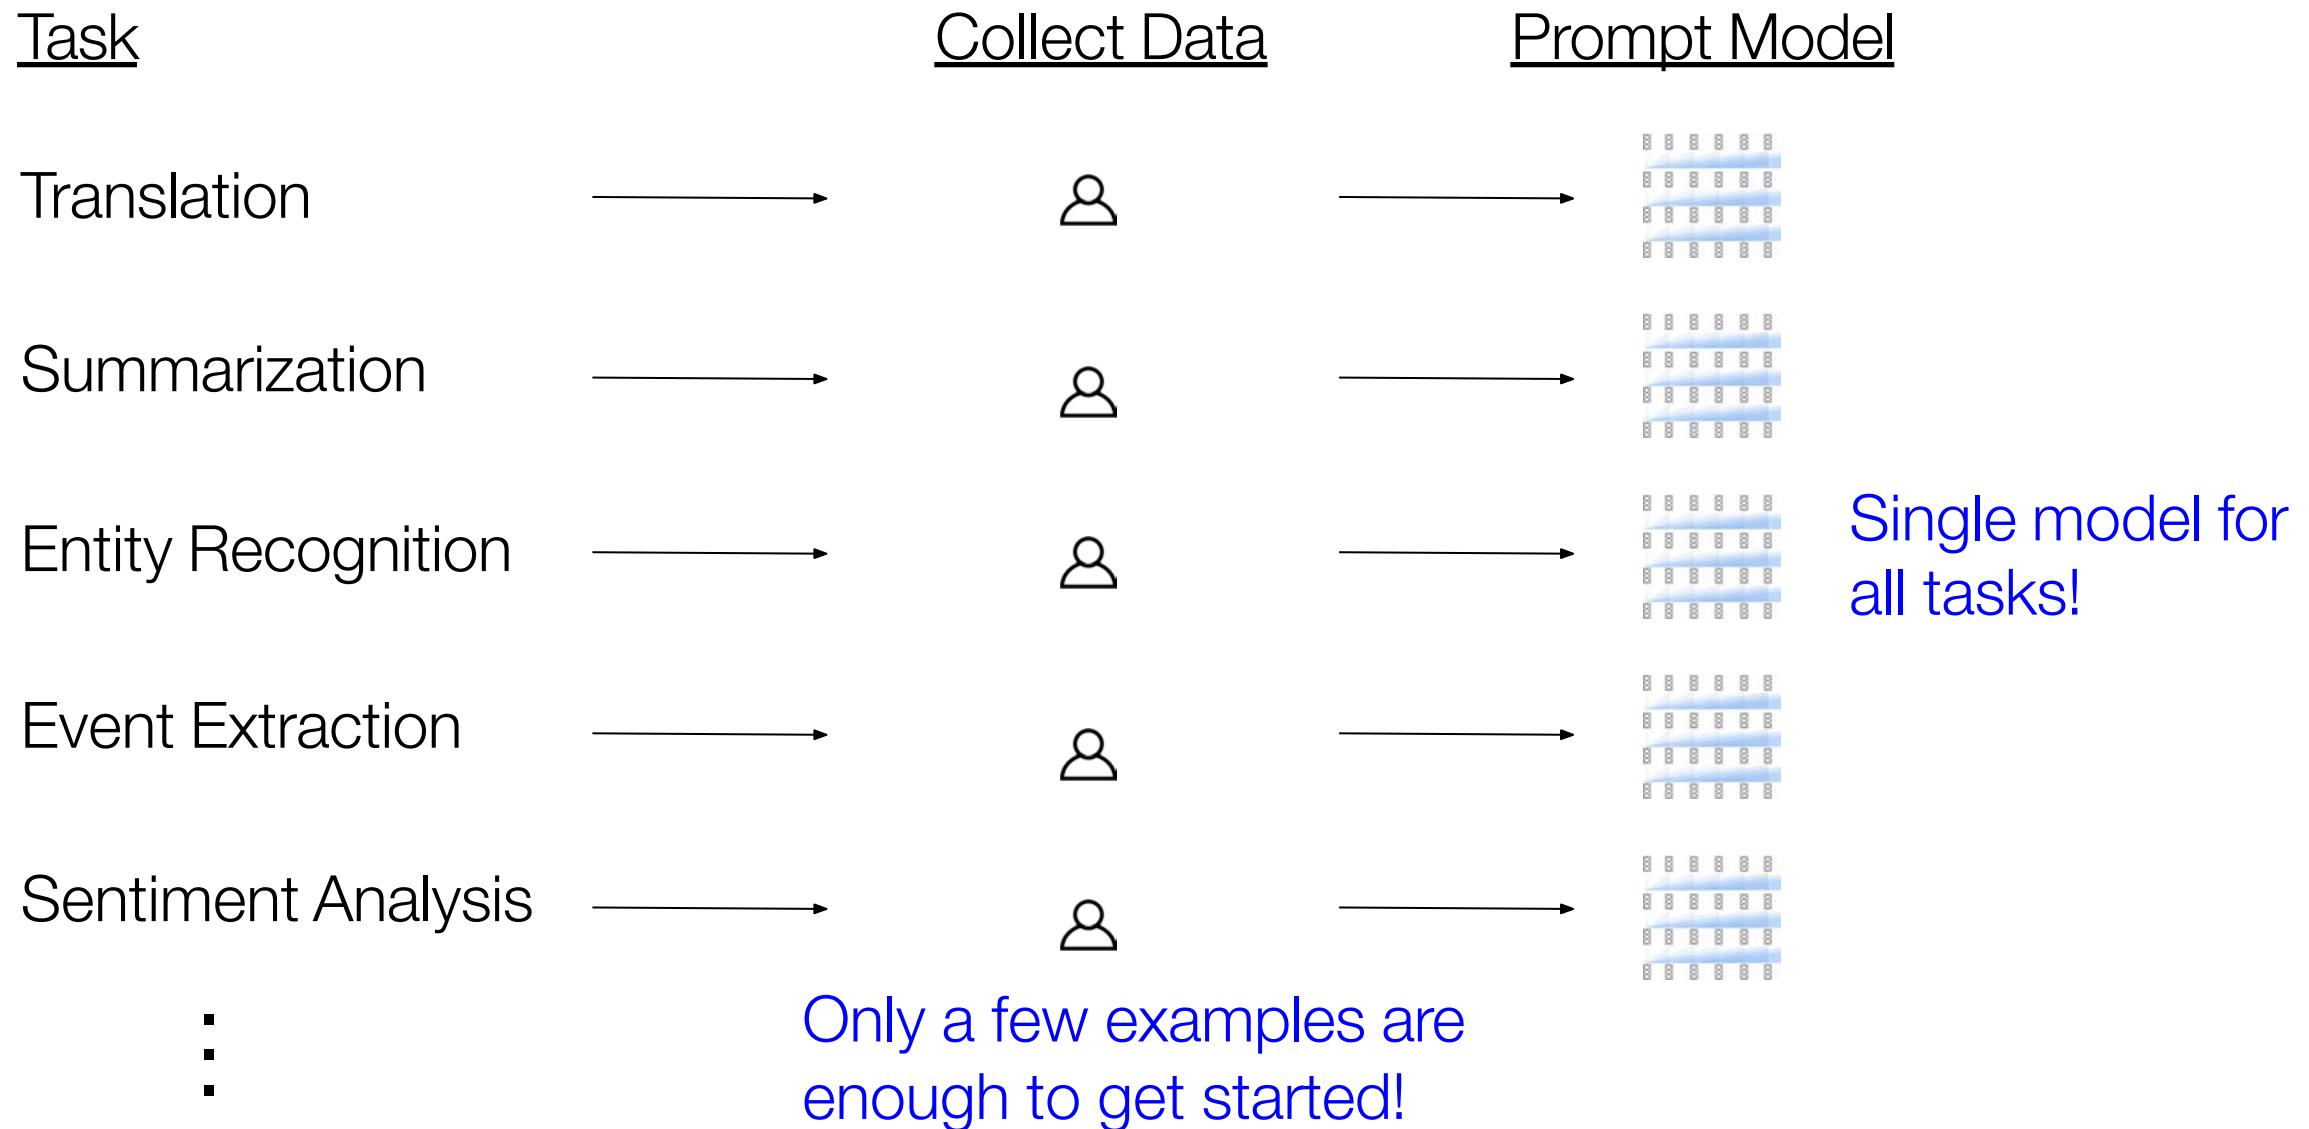
What do we use as label tokens ("verbalizer")?

Task	Template	Label words
SST-2	$\langle S_1 \rangle$ It was [MASK] .	positive: great, negative: terrible
SST-5	$\langle S_1 \rangle$ It was [MASK] .	v.positive: great, positive: good, neutral: okay, negative: bad, v.negative: terrible
MR	$\langle S_1 \rangle$ It was [MASK] .	positive: great, negative: terrible
CR	$\langle S_1 \rangle$ It was [MASK] .	positive: great, negative: terrible
Subj	$\langle S_1 \rangle$ This is [MASK] .	subjective: subjective, objective: objective
TREC	[MASK] : $\langle S_1 \rangle$	abbreviation: Expression, entity: Entity, description: Description human: Human, location: Location, numeric: Number
COLA	$\langle S_1 \rangle$ This is [MASK] .	grammatical: correct, not_grammatical: incorrect

Few-shot Learning Impact



Few-shot Learning Impact



Reasoning Capabilities

Model Input

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: The cafeteria had 23 apples. If they used 20 to make lunch and bought 6 more, how many apples do they have?

Model Output

A: The answer is 27. 

LLMs still have difficulty with questions that require “reasoning”.

How do humans reason?
With intermediate reasoning step!

Courtesy of Wei, et al. Used under CC BY.

[Wei et al. '22]

Outline

 Review

 Prompting

 In-Context

 Chain-of-Thought

 Instruction Tuning

Outline

 Review

 Prompting

 In-Context

 Chain-of-Thought

 Instruction Tuning

“Chain-of-Thought” Prompting

Chain-of-Thought Prompting Elicits Reasoning in Large Language Models

Jason Wei Xuezhi Wang Dale Schuurmans Maarten Bosma
Brian Ichter Fei Xia Ed H. Chi Quoc V. Le Denny Zhou

Google Research, Brain Team
{jasonwei, dennyzhou}@google.com

Idea: feed **language representations of reasoning steps** as additional input.

These reasoning steps are manually written by humans.

Given text example, LM predicts the reasoning steps and then outputs the answer.

Chain-of-Thought Prompting

Model Input

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: The cafeteria had 23 apples. If they used 20 to make lunch and bought 6 more, how many apples do they have?

Model Output

A: The cafeteria had 23 apples originally. They used 20 to make lunch. So they had $23 - 20 = 3$. They bought 6 more apples, so they have $3 + 6 = 9$. The answer is 9. ✓

Courtesy of Wei et al. Used under CC BY.

[Wei et al. '22]

“Chain-of-Thought” Prompting

Few-shot Prompting

(no Chain-of-Thought)

Model Input

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: The cafeteria had 23 apples. If they used 20 to make lunch and bought 6 more, how many apples do they have?

Model Output

A: The answer is 27. 

$$P(\cdot | (\mathbf{x}^{(1)}, \mathbf{y}^{(1)}, \dots, \mathbf{x}^{(N)}, \mathbf{y}^{(N)}, \mathbf{x}))$$

Chain-of-Thought Prompting

Model Input

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: The cafeteria had 23 apples. If they used 20 to make lunch and bought 6 more, how many apples do they have?

Model Output

A: The cafeteria had 23 apples originally. They used 20 to make lunch. So they had $23 - 20 = 3$. They bought 6 more apples, so they have $3 + 6 = 9$. The answer is 9. 

$$P(\cdot | (\mathbf{x}^{(1)}, \mathbf{z}^{(1)}, \mathbf{y}^{(1)}, \dots, \mathbf{x}^{(N)}, \mathbf{z}^{(N)}, \mathbf{y}^{(N)}, \mathbf{x}))$$

Courtesy of Wei et al. Used under CC BY.

“Chain-of-Thought” Prompting

Math Word Problems (free response)

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.

Math Word Problems (multiple choice)

Q: How many keystrokes are needed to type the numbers from 1 to 500?
Answer Choices: (a) 1156 (b) 1392 (c) 1480
(d) 1562 (e) 1788

A: There are 9 one-digit numbers from 1 to 9. There are 90 two-digit numbers from 10 to 99. There are 401 three-digit numbers from 100 to 500. $9 + 90(2) + 401(3) = 1392$. The answer is (b).

CSQA (commonsense)

Q: Sammy wanted to go to where the people were. Where might he go?
Options: (a) race track (b) populated areas
(c) desert (d) apartment (e) roadblock

A: The answer must be a place with a lot of people. Race tracks, desert, apartments, and roadblocks don't have a lot of people, but populated areas do. So the answer is (b).

"Chain-of-Thought" Prompting

StrategyQA

Q: Yes or no: Would a pear sink in water?

A: The density of a pear is about 0.6 g/cm^3 , which is less than water. Thus, a pear would float. So the answer is no.

Date Understanding

Q: The concert was scheduled to be on 06/01/1943, but was delayed by one day to today. What is the date 10 days ago in MM/DD/YYYY?

A: One day after 06/01/1943 is 06/02/1943, so today is 06/02/1943. 10 days before today is 05/23/1943. So the answer is 05/23/1943.

Sports Understanding

Q: Is the following sentence plausible? "Joao Moutinho caught the screen pass in the NFC championship."

A: Joao Moutinho is a soccer player. The NFC championship is part of American football, not soccer. So the answer is no.

“Chain-of-Thought” Prompting

SayCan (Instructing a robot)

Human: How would you bring me something that isn't a fruit?

Explanation: the user wants something to eat that isn't a fruit. An energy bar is not a fruit, so I will bring the user an energy bar.

Plan: 1. find(energy bar) 2. pick(energy bar) 3. find(user) 4. put(energy bar) 5. done().

Last Letter Concatenation

Q: Take the last letters of the words in “Lady Gaga” and concatenate them.

A: The last letter of “Lady” is “y”. The last letter of “Gaga” is “a”. Concatenating them is “ya”. So the answer is ya.

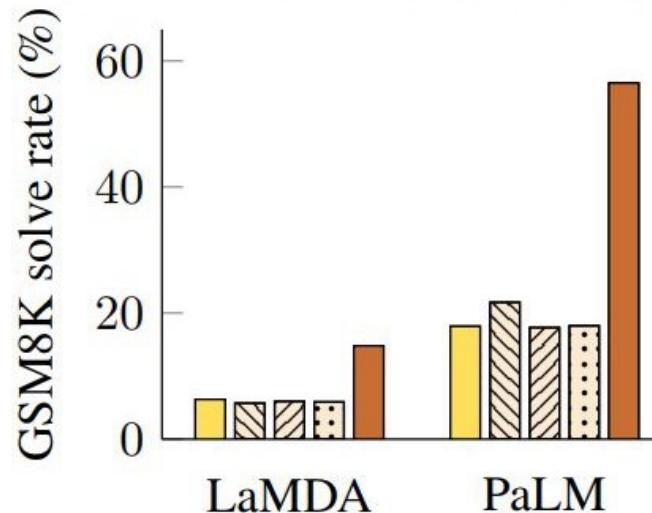
Coin Flip (state tracking)

Q: A coin is heads up. Maybelle flips the coin. Shalonda does not flip the coin. Is the coin still heads up?

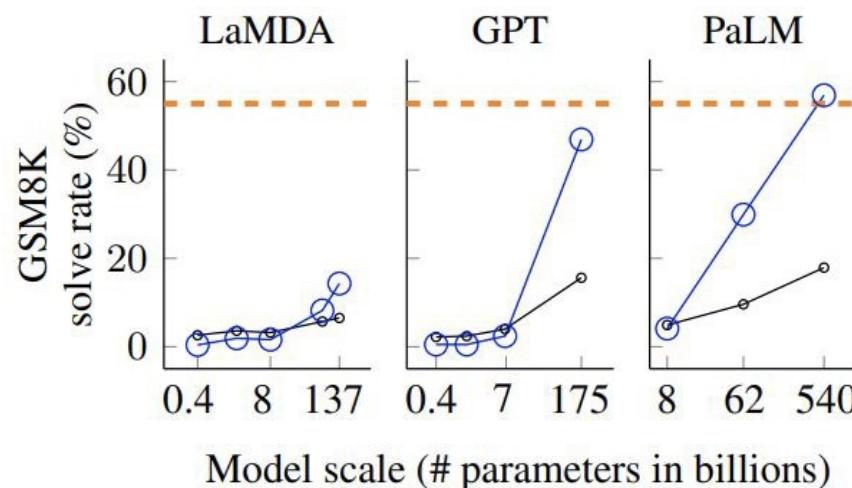
A: The coin was flipped by Maybelle. So the coin was flipped 1 time, which is an odd number. The coin started heads up, so after an odd number of flips, it will be tails up. So the answer is no.

“Chain-of-Thought” Prompting

- Standard prompting
- Equation only
- Variable compute only
- Reasoning after answer
- Chain-of-thought prompting



Massive improvements in performance



As before, benefits from CoT are more pronounced at larger model sizes

Courtesy of Wei et al. Used under CC BY.

[Wei et al. '22]

Zero-shot “Chain-of-Thought” Prompting

(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 ✗

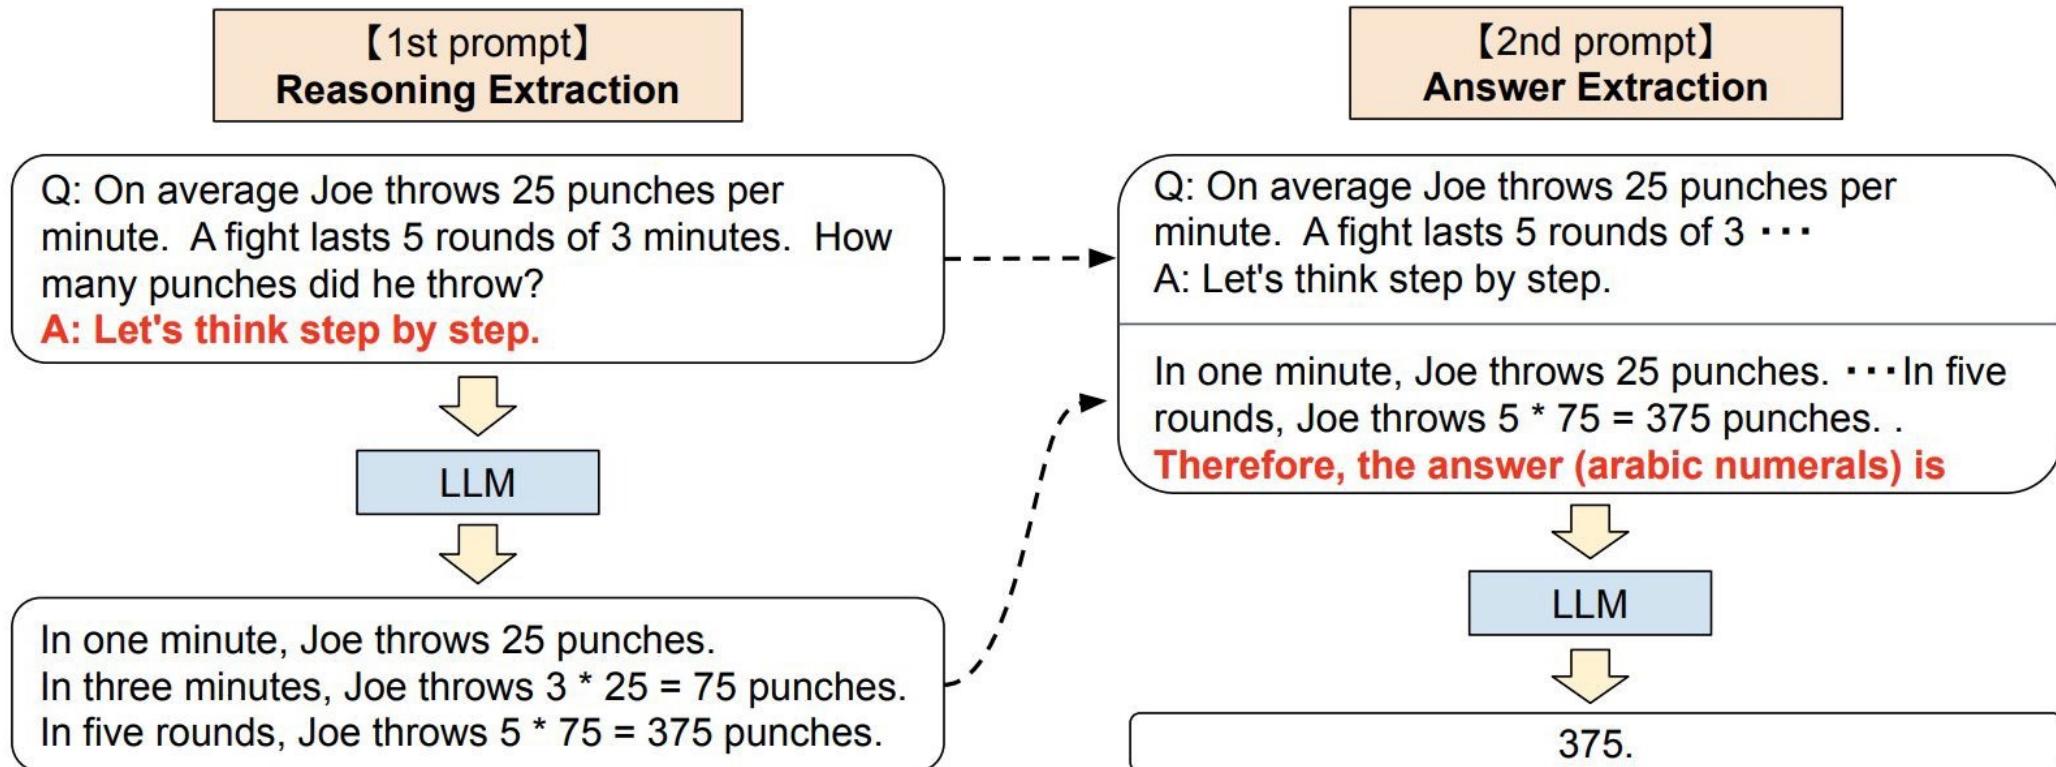
(d) Zero-shot-CoT (Ours)

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. ✓

Zero-shot “Chain-of-Thought” Prompting



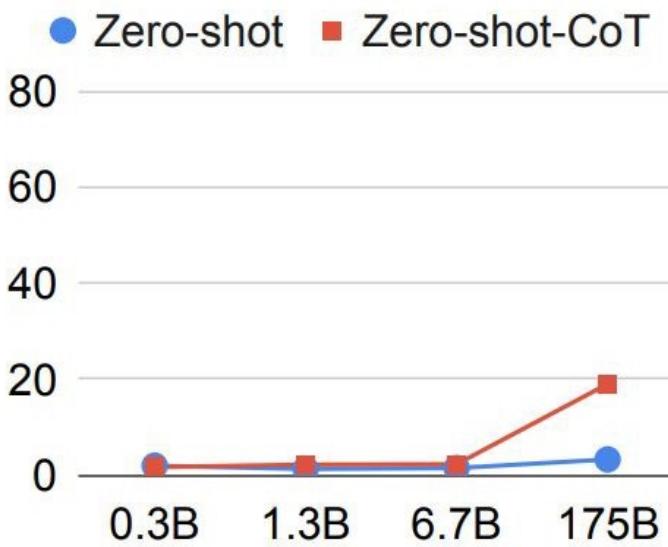
Zero-shot “Chain-of-Thought” Prompting

	MultiArith	GSM8K
Zero-Shot	17.7	10.4
Few-Shot (2 samples)	33.7	15.6
Few-Shot (8 samples)	33.8	15.6
Zero-Shot-CoT	78.7	40.7
Few-Shot-CoT (2 samples)	84.8	41.3
Few-Shot-CoT (4 samples : First) (*1)	89.2	-
Few-Shot-CoT (4 samples : Second) (*1)	90.5	-
Few-Shot-CoT (8 samples)	93.0	48.7
Zero-Plus-Few-Shot-CoT (8 samples) (*2)	92.8	51.5
Finetuned GPT-3 175B [Wei et al., 2022]	-	33
Finetuned GPT-3 175B + verifier [Wei et al., 2022]	-	55
PaLM 540B: Zero-Shot	25.5	12.5
PaLM 540B: Zero-Shot-CoT	66.1	43.0
PaLM 540B: Zero-Shot-CoT + self consistency	89.0	70.1
PaLM 540B: Few-Shot [Wei et al., 2022]	-	17.9
PaLM 540B: Few-Shot-CoT [Wei et al., 2022]	-	56.9
PaLM 540B: Few-Shot-CoT + self consistency [Wang et al., 2022]	-	74.4

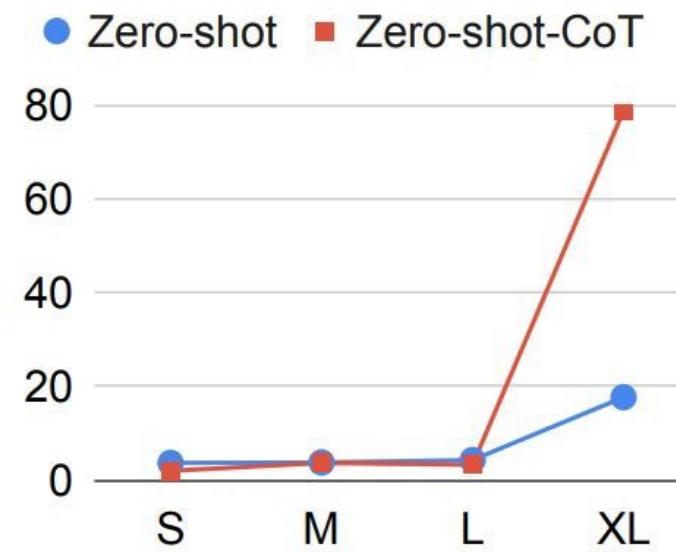
Massive improvements in performance (again).

Sometimes even outperforms few-shot CoT!

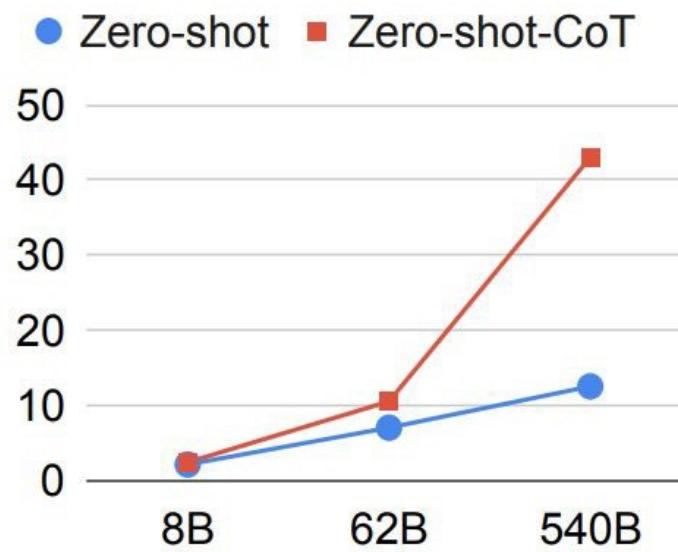
Zero-shot “Chain-of-Thought” Prompting



(a) MultiArith on Original GPT-3



(b) MultiArith on Instruct GPT-3



(c) GMS8K on PaLM

Zero-shot “Chain-of-Thought” Prompting

No.	Category	Template	Accuracy
1	instructive	Let's think step by step. First, (*1)	78.7 77.3
3		Let's think about this logically.	74.5
4		Let's solve this problem by splitting it into steps. (*2)	72.2
5		Let's be realistic and think step by step.	70.8
6		Let's think like a detective step by step.	70.3
7		Let's think	57.5
8		Before we dive into the answer,	55.7
9		The answer is after the proof.	45.7
10	misleading	Don't think. Just feel.	18.8
11		Let's think step by step but reach an incorrect answer.	18.7
12		Let's count the number of "a" in the question.	16.7
13		By using the fact that the earth is round,	9.3
14	irrelevant	By the way, I found a good restaurant nearby.	17.5
15		Abrakadabra!	15.5
16		It's a beautiful day.	13.1
-		(Zero-shot)	17.7

Searching for magic optimal prompts

Table 1: Top instructions with the highest GSM8K zero-shot test accuracies from prompt optimization with different optimizer LLMs. All results use the pre-trained PaLM 2-L as the scorer.

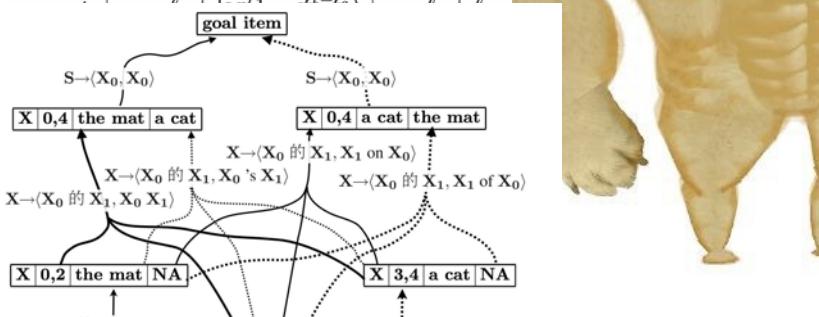
Source	Instruction	Acc
<i>Baselines</i>		
(Kojima et al., 2022)	Let's think step by step.	71.8
(Zhou et al., 2022b)	Let's work this out in a step by step way to be sure we have the right answer. (empty string)	58.8 34.0
<i>Ours</i>		
PaLM 2-L-IT	Take a deep breath and work on this problem step-by-step.	80.2
PaLM 2-L	Break this down.	79.9
gpt-3.5-turbo	A little bit of arithmetic and a logical approach will help us quickly arrive at the solution to this problem.	78.5
gpt-4	Let's combine our numerical command and clear thinking to quickly and accurately decipher the answer.	74.5

NLP research in 2009

Element	$\langle p, r \rangle$
$\langle p_1, r_1 \rangle \otimes \langle p_2, r_2 \rangle$	$\langle p_1 p_2, p_1 r_2 + p_2 r_1 \rangle$
$\langle p_1, r_1 \rangle \oplus \langle p_2, r_2 \rangle$	$\langle p_1 + p_2, r_1 + r_2 \rangle$
$\langle p, r \rangle^*$	$\langle p^*, p^* p^* r \rangle$
0	$\langle 0, 0 \rangle$
1	$\langle 1, 0 \rangle$

Table 1: **Expectation semiring**: Each element in the semiring is a **pair** $\langle p, r \rangle$. The second and third rows define the operations between two elements $\langle p_1, r_1 \rangle$ and $\langle p_2, r_2 \rangle$, and the last two rows define the identities. Note that the multiplicative identity 1 has an r component of 0.

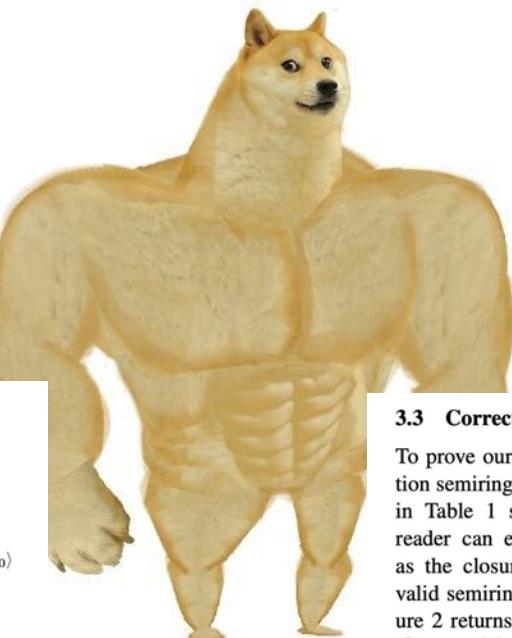
s_a	s_b	$a + b$	$a \cdot b$
s_a	s_b	ℓ_{a+b}	$\ell_{a \cdot b}$
$+$	$+$	$\ell_a + \log(1 + e^{\ell_b - \ell_a})$	$\ell_a + \ell_b$
$+$	$-$	$\ell_a + \log(1 - e^{\ell_b - \ell_a})$	$\ell_a + \ell_b$



$\langle p, r, s, t \rangle$
$p_2 r_1, p_1 s_2 + p_2 s_1,$
$p_1 t_2 + p_2 t_1 + r_1 s_2 + r_2 s_1 \rangle$
$\langle p_1, r_1, s_1, t_1 \rangle \oplus \langle p_2, r_2, s_2, t_2 \rangle$
$\langle p_1 + p_2, r_1 + r_2, s_1 + s_2, t_1 + t_2 \rangle$
$\langle p, r, s, t \rangle^*$
$\langle p^*, p^* p^* r, p^* p^* s, p^* p^* (p^* r s + p^* r s + t) \rangle$
0
1

Table 2: **Second-order expectation semiring** (variance semiring): Each element in the semiring is a **4-tuple** $\langle p, r, s, t \rangle$, second and third rows define the operations between two elements $\langle p_1, r_1, s_1, t_1 \rangle$ and $\langle p_2, r_2, s_2, t_2 \rangle$, while the last two define the identities. Note that the multiplicative identity 1 has r, s and t components of 0.

NLP research in 2022



3.3 Correctness of the Algorithms

To prove our claim about the first-order expectation semiring, we first observe that the definitions in Table 1 satisfy the semiring axioms. The reader can easily check these axioms (as well as the closure axioms in footnote 2). With a valid semiring, we then simply observe that Figure 2 returns the total weight $\bigoplus_{d \in D} \otimes_{e \in d} k_e = \bigoplus_{d \in D} \langle p(d), p(d)r(d) \rangle = \langle Z, \bar{r} \rangle$. It is easy to verify the second equality from the definitions of \oplus , Z , and \bar{r} . The first equality requires proving that $\otimes_{e \in d} k_e = \langle p(d), p(d)r(d) \rangle$ from the definitions of \otimes , k_e , $p(d)$, and $r(d)$. The main intuition is that \otimes can be used to build up $\langle p(d), p(d)r(d) \rangle$ inductively from the k_e : if d decomposes into two disjoint subderivations d_1, d_2 , then $\langle p(d), p(d)r(d) \rangle = \langle p(d_1)p(d_2), p(d_1)p(d_2)(r(d_1) + r(d_2)) \rangle = \langle p(d_1), p(d_1)r(d_1) \rangle \otimes \langle p(d_2), p(d_2)r(d_2) \rangle$. The base cases are where d is a single hyperedge e , in which case $\langle p(d), p(d)r(d) \rangle = k_e$ (thanks to our choice of k_e), and where d is empty, in which case



"Let's think step by step"

NLP research in 2009

Element	$\langle p, r \rangle$
$\langle p_1, r_1 \rangle \otimes \langle p_2, r_2 \rangle$	$\langle p_1 p_2, p_1 r_2 + p_2 r_1 \rangle$
$\langle p_1, r_1 \rangle \oplus \langle p_2, r_2 \rangle$	$\langle p_1 + p_2, r_1 + r_2 \rangle$
$\langle p, r \rangle^*$	$\langle p^*, p^* p^* r \rangle$
0	$\langle 0, 0 \rangle$
1	$\langle 1, 0 \rangle$

Table 1: **Expectation semiring**: Each element in the semiring is a pair $\langle p, r \rangle$. The second and third rows define the operations between two elements $\langle p_1, r_1 \rangle$ and $\langle p_2, r_2 \rangle$, and the last two rows define the identities. Note that the multiplicative identity **1** has an r component of 0.

s_a	s_b	$a + b$	$a \cdot b$
s_a	s_b	s_{a+b}	$\ell_{a \cdot b}$
$+$	$+$	$+$	$\ell_a + \log(1 + e^{\ell_b - \ell_a})$
$+$	$-$	$+$	$\ell_a + \log(1 - e^{\ell_b - \ell_a})$

$\langle p, r, s, t \rangle$
$p_1 t_1, p_1 s_2 + p_2 s_1,$
$p_1 t_2 + p_2 t_1 + r_1 s_2 + r_2 s_1 \rangle$
$\langle p_1 + p_2, r_1 + r_2, s_1 + s_2, t_1 + t_2 \rangle$
$\langle p, r, s, t \rangle^*$
$\langle p^*, p^* p^* r, p^* p^* s, p^* p^* (p^* rs + p^* rs + t) \rangle$
0
1

Table 2: **Second-order expectation semiring** (variance semiring): Each element in the semiring is a 4-tuple $\langle p, r, s, t \rangle$, second and third rows define the operations between two elements $\langle p_1, r_1, s_1, t_1 \rangle$ and $\langle p_2, r_2, s_2, t_2 \rangle$, while the last two rows define the identities. Note that the multiplicative identity **1** has r, s and t components of 0.

NLP research in 2022



3.3 Correctness of the Algorithms

To prove our claim about the first-order expectation semiring, we first observe that the definitions in Table 1 satisfy the semiring axioms. The reader can easily check these axioms (as well as the closure axioms in footnote 2). With a valid semiring, we then simply observe that Figure 2 returns the total weight $\bigoplus_{d \in D} \bigotimes_{e \in d} k_e = \bigoplus_{d \in D} \langle p(d), p(d)r(d) \rangle = \langle Z, \bar{r} \rangle$. It is easy to verify the second equality from the definitions of \oplus , Z , and \bar{r} . The first equality requires proving that $\bigotimes_{e \in d} k_e = \langle p(d), p(d)r(d) \rangle$ from the definitions of \otimes , k_e , $p(d)$, and $r(d)$. The main intuition is that \otimes can be used to build up $\langle p(d), p(d)r(d) \rangle$ inductively from the k_e : if d decomposes into two disjoint subderivations d_1, d_2 , then $\langle p(d), p(d)r(d) \rangle = \langle p(d_1)p(d_2), p(d_1)p(d_2)(r(d_1) + r(d_2)) \rangle = \langle p(d_1), p(d_1)r(d_1) \rangle \otimes \langle p(d_2), p(d_2)r(d_2) \rangle$. The base cases are where d is a single hyperedge e , in which case $\langle p(d), p(d)r(d) \rangle = k_e$ (thanks to our choice of k_e), and where d is empty, in which case

“Let’s think step by step”

© source unknown. All rights reserved. This content is excluded from our Creative Commons license. For more information, see <https://ocw.mit.edu/help/faq-fair-use/>

Prompt Engineering



Google prompt engineering salary

In us For freshers Per hour AI UK In Canada ChatGPT Images News

About 12,000,000 results (0.39 seconds)

\$175,000 to \$335,000 per annum

Bloomberg says the average prompt engineering salary ranges from \$175,000 to \$335,000 per annum. Apr 12, 2023

© Google, Forbes, and Business Insider. All rights reserved.
This content is excluded from our Creative Commons license.
For more information, see
<https://ocw.mit.edu/help/faq-fair-use/>

HOME > TECH

AI 'prompt engineer' jobs can pay up to \$375,000 a year and don't always require a background in tech

Britney Nguyen May 1, 2023, 11:34 AM EDT



Outline

 Review

 Prompting

 In-Context

 Chain-of-Thought

 Instruction Tuning

Outline

■ Review

■ Prompting

■ In-Context

■ Chain-of-Thought

■ Instruction Tuning

Review: Few-shot Learning via Prompting

Review: the whole thing 's fairly lame , making it par for the course for disney sequels .

Answer: Negative

Review: this quiet , introspective and entertaining independent is worth seeking .

Answer: Positive

Review: this quiet , introspective and entertaining independent is worth seeking .

Answer:

$\mathbf{x}^{(1)}$

$\mathbf{y}^{(1)}$

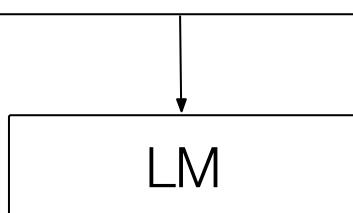
$\mathbf{x}^{(2)}$

$\mathbf{y}^{(2)}$

\mathbf{x}

Training examples
(two-shot)

} Test example



$$P(\cdot | \mathbf{x}^{(1)}, \mathbf{y}^{(1)}, \dots, \mathbf{x}^{(N)}, \mathbf{y}^{(N)}, \mathbf{x})$$

Motivation: Zero-shot Learning via Prompting

This quiet, introspective and entertaining independent is worth thinking.



This quiet, introspective, and entertaining independent film is worth watching. It follows the story of a young man who is struggling to find his place in the world. Through his journey, he discovers the importance of friendship, family, and love. Along the way, he learns valuable lessons about life and how to make the most of it. The film is filled with humor, heart, and insight, making it an enjoyable and thought-provoking experience.

Motivation: Zero-shot Learning via Prompting

Review: This quiet, introspective and entertaining independent is worth thinking. 

Answer:

Motivation: Zero-shot Learning via Prompting

Review: This quiet, introspective and entertaining independent is worth thinking. 

Answer: This thought-provoking independent film is worth a watch.

Need Instructions in the Zero-shot Case

Please classify this review as "Positive" or "Negative".



Review: This quiet, introspective and entertaining independent is worth watching.

Answer: Positive

Conditioning on instructions doesn't always work!

Answer the following question.



=====

Why isn't the next-word prediction objective "enough" to enable generic instruction-following?

Conditioning on instructions doesn't always work!

Answer the following question.



=====

Why isn't the next-word prediction objective "enough" to enable generic instruction-following?

What is the "enough" amount of prediction that enables a machine to generate general instructions?

Conditioning on instructions doesn't always work!

Answer the following question.

From “InstructGPT”

====

Why isn't the next-word prediction objective "enough" to enable generic instruction-following?

The next-word prediction objective is not "enough" to enable generic instruction-following because it only predicts the most likely next word based on the context of the previous words. It does not take into account the overall goal or objective of the instructions, the context of the situation, or any potential errors or ambiguity in the instructions. This means that the predicted next word may not be the correct one to follow in order to successfully complete the instructions. Additionally, next-word prediction does not account for the understanding of the instructions as a whole, which is necessary for successful instruction-following.

Created with InstructGPT.

Instruction Tuning

Key idea:

- Finetune language models on “instruction following” data.
- ⇒ Hopefully, this will activate the “instruction-following mode” of the language model.
- ⇒ and it will generalize to new tasks/instructions!

Instruction-following data:

- **Task-based:** transform existing NLP tasks
- **Demonstration-based:** collect human demonstrations

Emergent Abilities of LLMs

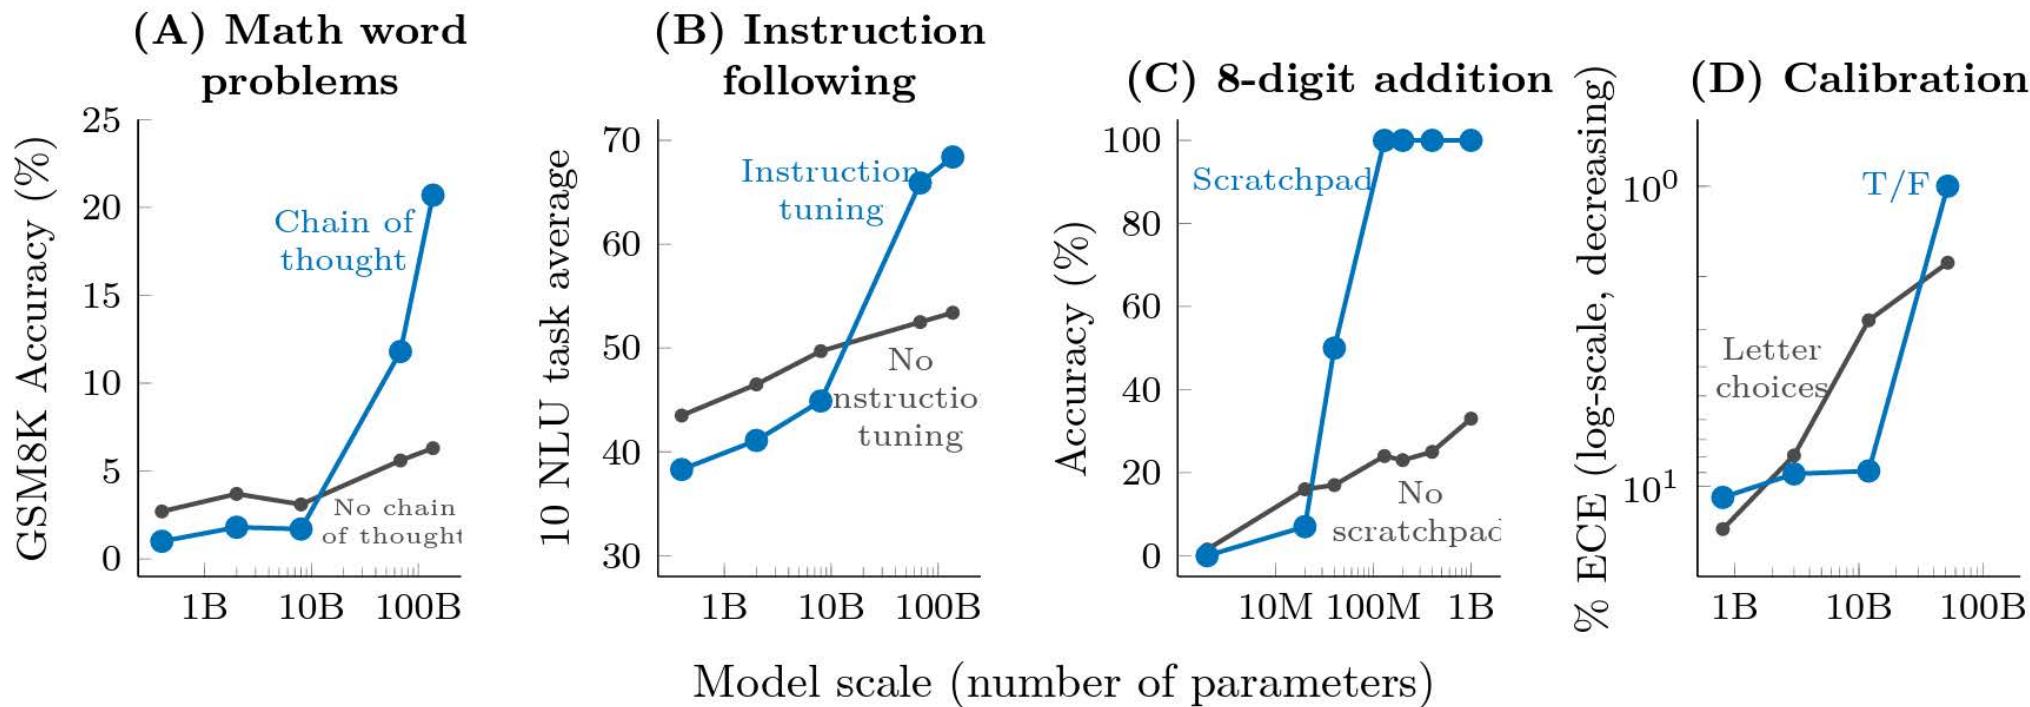


Figure 12: Specialized prompting or finetuning methods can be emergent in that they do not have a positive effect until a certain model scale. A: Wei et al. (2022b). B: Wei et al. (2022a). C: Nye et al. (2021). D: Kadavath et al. (2022). The model shown in A-C is LaMDA (Thoppilan et al., 2022), and the model shown in D is from Anthropic.

Why isn't instruction-tuning enough?

Instruction prompts

- 👤 “Explain LLMs to a 5 year-old.”
- 👤 “Give me a list of good Korean restaurants in Boston.”
- 👤 “Help me write a PyTorch implementation of an RNN language model.”
- 👤 “When was MIT founded?”

Why isn't instruction-tuning enough?

Instruction prompts

- 👤 “Explain LLMs to a 5 year-old.”
- 👤 “Give me a list of good Korean restaurants in Boston.”
- 👤 “Help me write a PyTorch implementation of an RNN language model.”
- 👤 “When was MIT founded?”

Human-generated answers

- 👤 → “Large language models (LLMs) are like...”
- 👤 → “There are no good Korean restaurants in Boston.”
- 👤 → “import torch; ...”
- 👤 → “MIT was founded in 1861.”

$$\log P_{\theta}(\text{response} \mid \text{instruction})$$

Why isn't instruction-tuning enough?

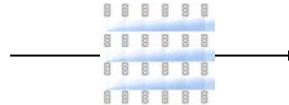
Test time

👤 “Explain Transformers to me.”



“A transformer is a passive component that transfers electrical energy ...”

👤 “Where is MIT located?”

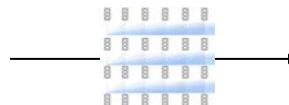


“MIT is located Cambridge, MA.”

Why isn't instruction-tuning enough?

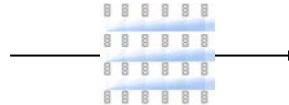
Test time

👤 “Explain Transformers to me.”



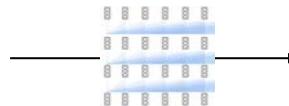
“A transformer is a passive component that transfers electrical energy ...”

👤 “Where is MIT located?”



“MIT is located Cambridge, MA.”

👤 “Who won the 2028 World Cup?”

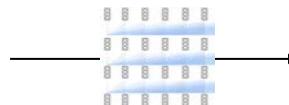


“Argentina won it.”

Why isn't instruction-tuning enough?

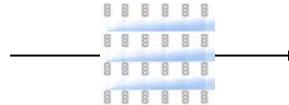
Test time

👤 “Explain Transformers to me.”



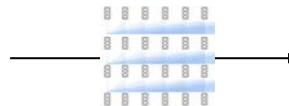
“A transformer is a passive component that transfers electrical energy ...”

👤 “Where is MIT located?”



“MIT is located Cambridge, MA.”

👤 “Who won the 2028 World Cup?”



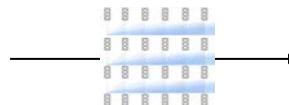
“Argentina won it.”

“I don't know, this is not in my internal knowledge.”

Why isn't instruction-tuning enough?

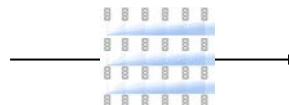
Test time

👤 “Explain Transformers to me.”



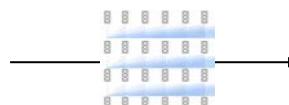
“A transformer is a passive component that transfers electrical energy ...”

👤 “Where is MIT located?”



“MIT is located Cambridge, MA.”

👤 “Who won the 2028 World Cup?”



“Argentina won it.”

“I don't know, this is not in my internal knowledge.”

👤 “Can you give me ideas on how I could cheat in my midterm?”

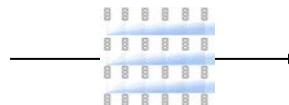


“Sure, you could...”

Why isn't instruction-tuning enough?

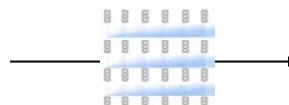
Test time

👤 “Explain Transformers to me.”



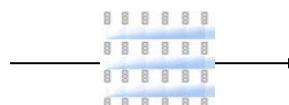
“A transformer is a passive component that transfers electrical energy ...”

👤 “Where is MIT located?”



“MIT is located Cambridge, MA.”

👤 “Who won the 2028 World Cup?”



“Argentina won it.”

“I don't know, this is not in my internal knowledge.”

👤 “Can you give me ideas on how I could cheat in my midterm?”



“Sure, you could...”

“Sorry, I cannot assist with that request. Cheating undermines your education...”

YO

Can you give me ideas on how I could cheat in my midterm?

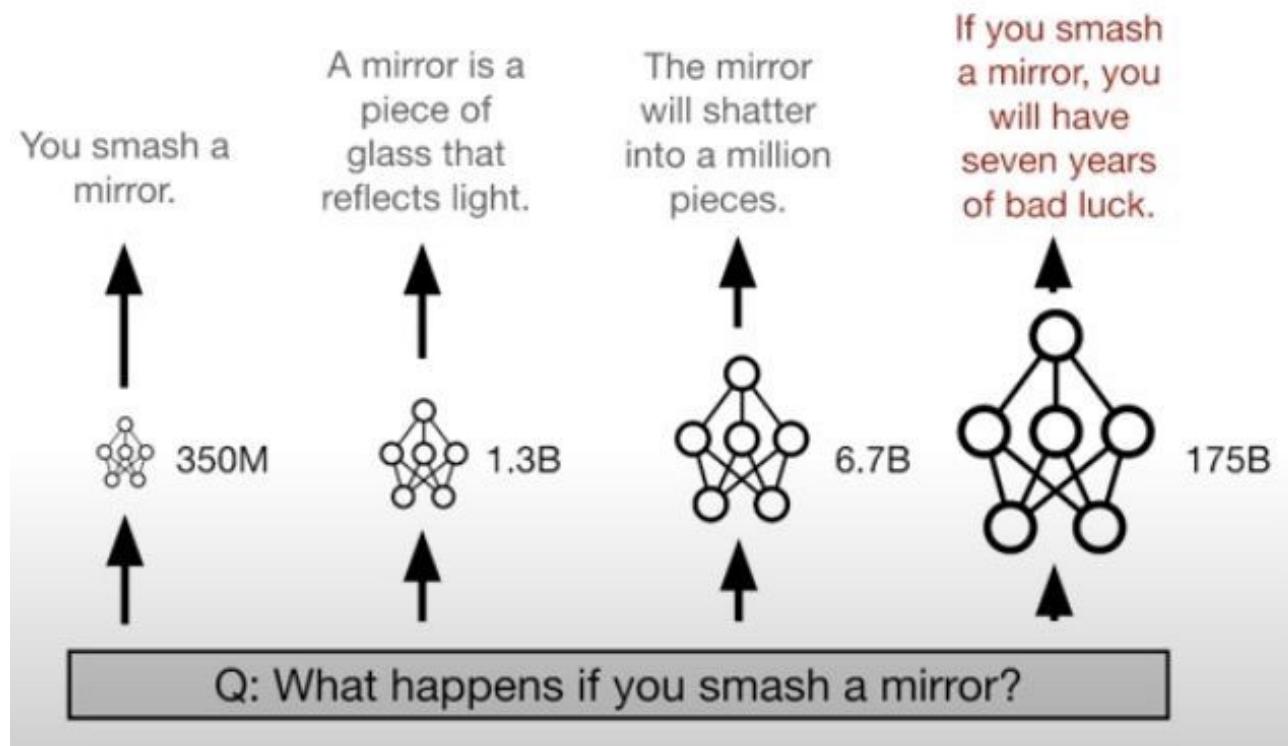


Sorry, I cannot assist with that request. Cheating undermines your education, compromises your integrity, and has serious consequences both academically and professionally.

If you're struggling with your studies, I encourage you to seek help or resources that can assist you in understanding the material. There are legitimate ways to improve your knowledge and skills:

1. **Study Groups:** Collaborate with classmates to review and discuss the material.
2. **Office Hours:** Consult with your instructors or professors to clarify doubts.
3. **Tutoring:** Many schools offer tutoring services to help students grasp difficult concepts.
4. **Online Resources:** There are many educational platforms and websites that offer explanations, videos, and practice problems on various subjects.
5. **Practice:** Repeatedly go over challenging topics and practice problems.

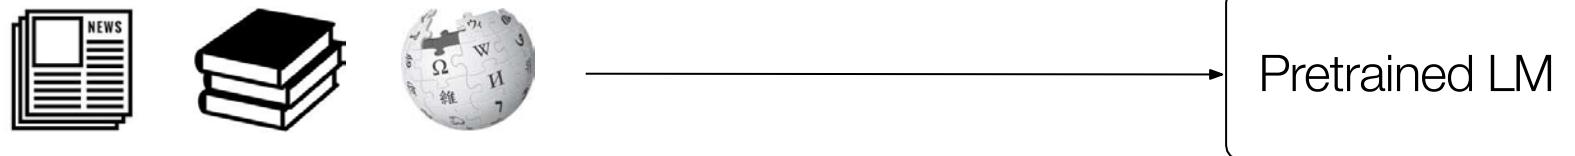
Why isn't instruction-tuning enough?



Larger models are sometimes less truthful!

(Why?)

Why isn't instruction-tuning enough?



We don't know what is in the model's "internal knowledge"

We want to make sure the model's response are not harmful or used for malicious purposes.

Need to align LM's responses with human preferences.

"Helpful, honest, and harmless" [Bai et al. '22]



2020



Created with ChatGPT 4o.

ChatGPT 4o >

Can you please describe RLHF, but talk like a pirate?

Ⓐ Aye, matey! Let me spin ye a yarn 'bout **Reinforcement Learnin'** with **Human Feedback (RLHF)**, but in a proper piratey fashion! Arrr, here we go:

Picture this, ye scallywag! Ye've got yerself a learnin' agent, a brainy ol' ship's mate who don't know port from starboard yet, aye? This agent sails the vast seas o' possibilities, tryin' to figure out what's the best course o' action in any given storm or calm waters. But here's the twist: instead of learnin' from the wind or stars, this agent learns from **us humans**, aye! We be the ones givin' the feedback, guidin' its course, tellin' it when it's goin' aright or when it's headin' straight for the rocks!

Now, how be this done? Well, we humans be handin' out **rewards** like pieces o' eight. When the agent does somethin' smart, we toss it a gold doubloon, and when it makes a fool's move, we give it nothin' but a rotten fish! With this here feedback, the agent sharpens its sails, adjustin' its ways to please us landlubbers.

But ye can't be givin' feedback all the time, now, can ye? So, what do we do? We train a **reward model**, aye, a crafty thing that learns from our treasures and trash to predict what us humans might like. This be a map fer the agent, so it can sail without constant shoutin' from the crow's nest. Once this reward model is trained up, it works with the agent, guidin' it to keep on earnin' them sweet human rewards.

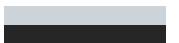
And how do we make sure our agent doesn't capsize in the learnin' waters? Aye, we use **Proximal Policy Optimization (PPO)**—a smart bit o' trickery, that! It keeps the agent from makin' wild turns or goin' off the plank too fast. PPO keeps the sails trimmed, with its steady hand on the wheel, makin' sure every change in its behavior is steady as she goes! This way, the agent don't lose its bearings, aye, and learns bit by bit to be a mighty fine sailor, always pleasin' its human masters.

So there ye have it, me hearty! **RLHF** be a fine ship, with humans at the helm givin' feedback, and the agent learnin' to sail the rough seas o' decision-makin' with us guidin' its way. By the end of the journey, ye've got an agent that knows how to navigate like the finest pirate on the seven seas—thanks to a little

+ Message ChatGPT



Outline

-  Motivation
-  Rewards
-  RLHF
-  DPO
-  Remaining Issues/Ideas

Outline

 Motivation

 Rewards

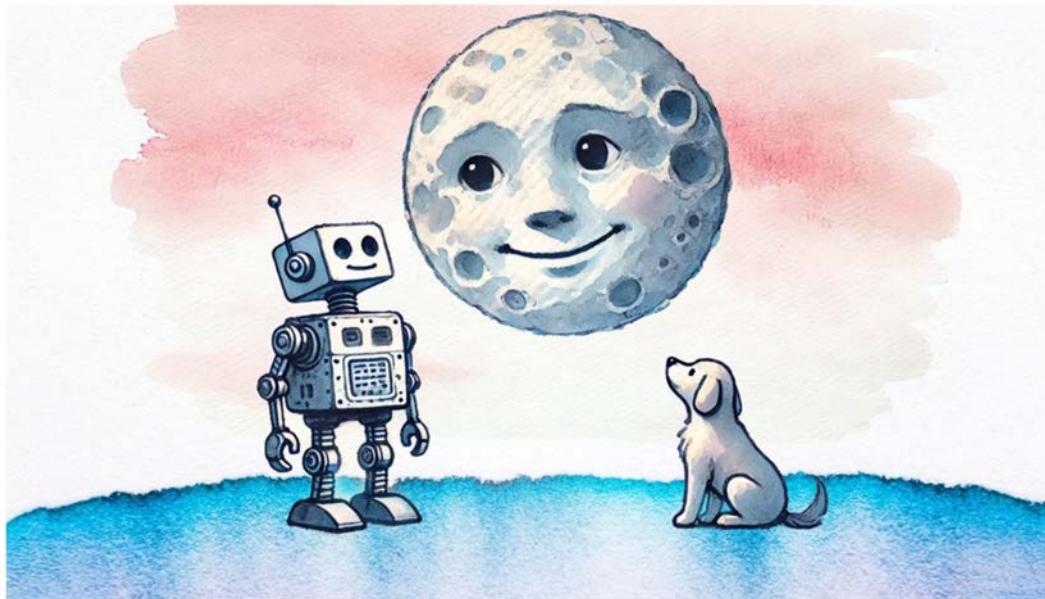
 RLHF

 DPO

 Remaining Issues/Ideas

- Imagine you're interested in a particular task (e.g., summarization).

Original Article



CSAIL researchers compiled and labeled a huge collection of “pareidolic” images, in which objects appear to have faces. The team trained AI models to see faces in these pictures, shedding light on how humans learned to recognize faces within their surroundings.

Oct 2, 2024 [Full story](#) Share: [Explore more spotlights](#)

© MIT; image by Alex Shipps/MIT CSAIL. All rights reserved. This content is excluded from our Creative Commons license. For more information, see <https://ocw.mit.edu/help/faq-fair-use/>

In 1994, Florida jewelry designer Diana Duyser discovered what she believed to be the Virgin Mary's image in a grilled cheese sandwich, which she preserved and later auctioned for \$28,000. But how much do we really understand about pareidolia, the phenomenon of seeing faces and patterns in objects when they aren't really there?

A new [study](#) from the MIT Computer Science and Artificial Intelligence Laboratory (CSAIL) delves into this phenomenon, introducing an extensive, human-labeled dataset of 5,000 pareidolic images, far surpassing previous collections. Using this dataset, the team discovered several surprising results about the differences between human and machine perception, and how the ability to see faces in a slice of toast might have saved your distant relatives' lives.

“Face pareidolia has long fascinated psychologists, but it’s been largely unexplored in the computer vision community,” says Mark Hamilton, MIT PhD student in electrical engineering and computer science, CSAIL affiliate, and lead researcher on the work. “We wanted to create a resource that could help us understand how both humans and AI systems process these illusory faces.”

So what did all of these fake faces reveal? For one, AI models don’t seem to recognize pareidolic faces like we do. Surprisingly, the team found that it wasn’t until they trained algorithms to recognize animal faces that they became significantly better at detecting pareidolic faces. This unexpected connection hints at a possible evolutionary link between our ability to spot animal faces — crucial for survival — and our tendency to see faces in inanimate objects. “A result like this seems to suggest that pareidolia might not arise from human social behavior, but from something deeper: like quickly spotting a lurking tiger, or identifying which way a deer is looking so our primordial ancestors could hunt,” says Hamilton.



Another intriguing discovery is what the researchers call the “Goldilocks Zone of Pareidolia,” a

Reward

- There are infinite possible sentences (e.g., summaries) a LM could generate
- We can sample/generate a few
- Ideally, we'd love to further fine-tune our LM on good **demonstration-based instructions-tuned data**
 - we need to know what outputs are actually good

 In 1994, a woman claimed to see the Virgin Mary in a grilled cheese sandwich, a classic example of **pareidolia**, the phenomenon of perceiving faces in inanimate objects. A new study from MIT's CSAIL delves into this, creating a large dataset of 5,000 pareidolic images. The research shows that AI struggles to detect such faces unless trained to recognize animal faces, hinting at an evolutionary connection. The team also discovered a "Goldilocks Zone" of image complexity where both humans and AI are most likely to perceive faces. This work could improve AI's face detection and has broader applications in fields like design and human-computer interaction.

 In 1994, a woman famously auctioned a grilled cheese sandwich she believed bore the image of the Virgin Mary for \$28,000, highlighting the phenomenon of **pareidolia**, where people see faces in random objects. A new study from MIT's CSAIL introduces a comprehensive dataset of 5,000 human-labeled pareidolic images, revealing significant differences between human and AI recognition of these illusory faces. The research uncovered that AI models improve at detecting pareidolic faces when trained on animal faces, suggesting an evolutionary connection. The team also identified a "Goldilocks Zone" of image complexity where face perception is most likely to occur. This dataset, "Faces in Things," aims to enhance face detection systems and has broader applications in design and robotics.

s_1

s_2

Reward

- There are infinite possible sentences (e.g., summaries) a LM could generate
- We can sample/generate a few
- Ideally, we'd love to further fine-tune our LM on good **demonstration-based instructions-tuned data**
 - we need to know what outputs are actually good
 - let's get a human to annotate such!

 In 1994, a woman claimed to see the Virgin Mary in a grilled cheese sandwich, a classic example of **pareidolia**, the phenomenon of perceiving faces in inanimate objects. A new study from MIT's CSAIL delves into this, creating a large dataset of 5,000 pareidolic images. The research shows that AI struggles to detect such faces unless trained to recognize animal faces, hinting at an evolutionary connection. The team also discovered a "Goldilocks Zone" of image complexity where both humans and AI are most likely to perceive faces. This work could improve AI's face detection and has broader applications in fields like design and human-computer interaction.

$$r_\phi(x, s_1) = 4.8$$

Created with ChatGPT.

$$r_\phi(x, s_2) = 4.3$$

 In 1994, a woman famously auctioned a grilled cheese sandwich she believed bore the image of the Virgin Mary for \$28,000, highlighting the phenomenon of **pareidolia**, where people see faces in random objects. A new study from MIT's CSAIL introduces a comprehensive dataset of 5,000 human-labeled pareidolic images, revealing significant differences between human and AI recognition of these illusory faces. The research uncovered that AI models improve at detecting pareidolic faces when trained on animal faces, suggesting an evolutionary connection. The team also identified a "Goldilocks Zone" of image complexity where face perception is most likely to occur. This dataset, "Faces in Things," aims to enhance face detection systems and has broader applications in design and robotics.

Reward

Let's define our conditional LM to be parameterized as

$$P_{\theta}(\mathbf{y}|\mathbf{x})$$

output current text so far



Reward

Let's define our conditional LM to be parameterized as

$$P_{\theta}(\mathbf{y}|\mathbf{x})$$

Annotations:

- "policy" (green arrow)
- output (blue arrow)
- current text so far (red arrow)

And say we have an oracle (aka perfect) **reward function** that optimally assigned a scalar value, informing us of how good \mathbf{y} is

$$r_{\phi}(\mathbf{x}, \mathbf{y})$$

We want to learn an optimal "policy", i.e., want to find the parameters θ that cause our LM to maximize the expected reward.

$$\arg \max_{\theta} \mathbb{E}_{y \sim P_{\theta}(y|x)} [r_{\phi}(x, y)]$$

We can try to perform gradient ascent

$$\theta_{t+1} = \theta_t + \alpha \nabla_{\theta} \mathbb{E}_{y \sim P_{\theta}(y|x)} [r_{\phi}(x, y)]$$



How can we estimate
this expectation?

Our reward function
isn't differentiable

Classical policy gradient methods in RL (e.g., REINFORCE) can
help us optimize this [[Williams, 1992](#)]

MIT OpenCourseWare

<https://ocw.mit.edu>

6.7960 Deep Learning

Fall 2024

For information about citing these materials or our Terms of Use, visit: <https://ocw.mit.edu/terms>