

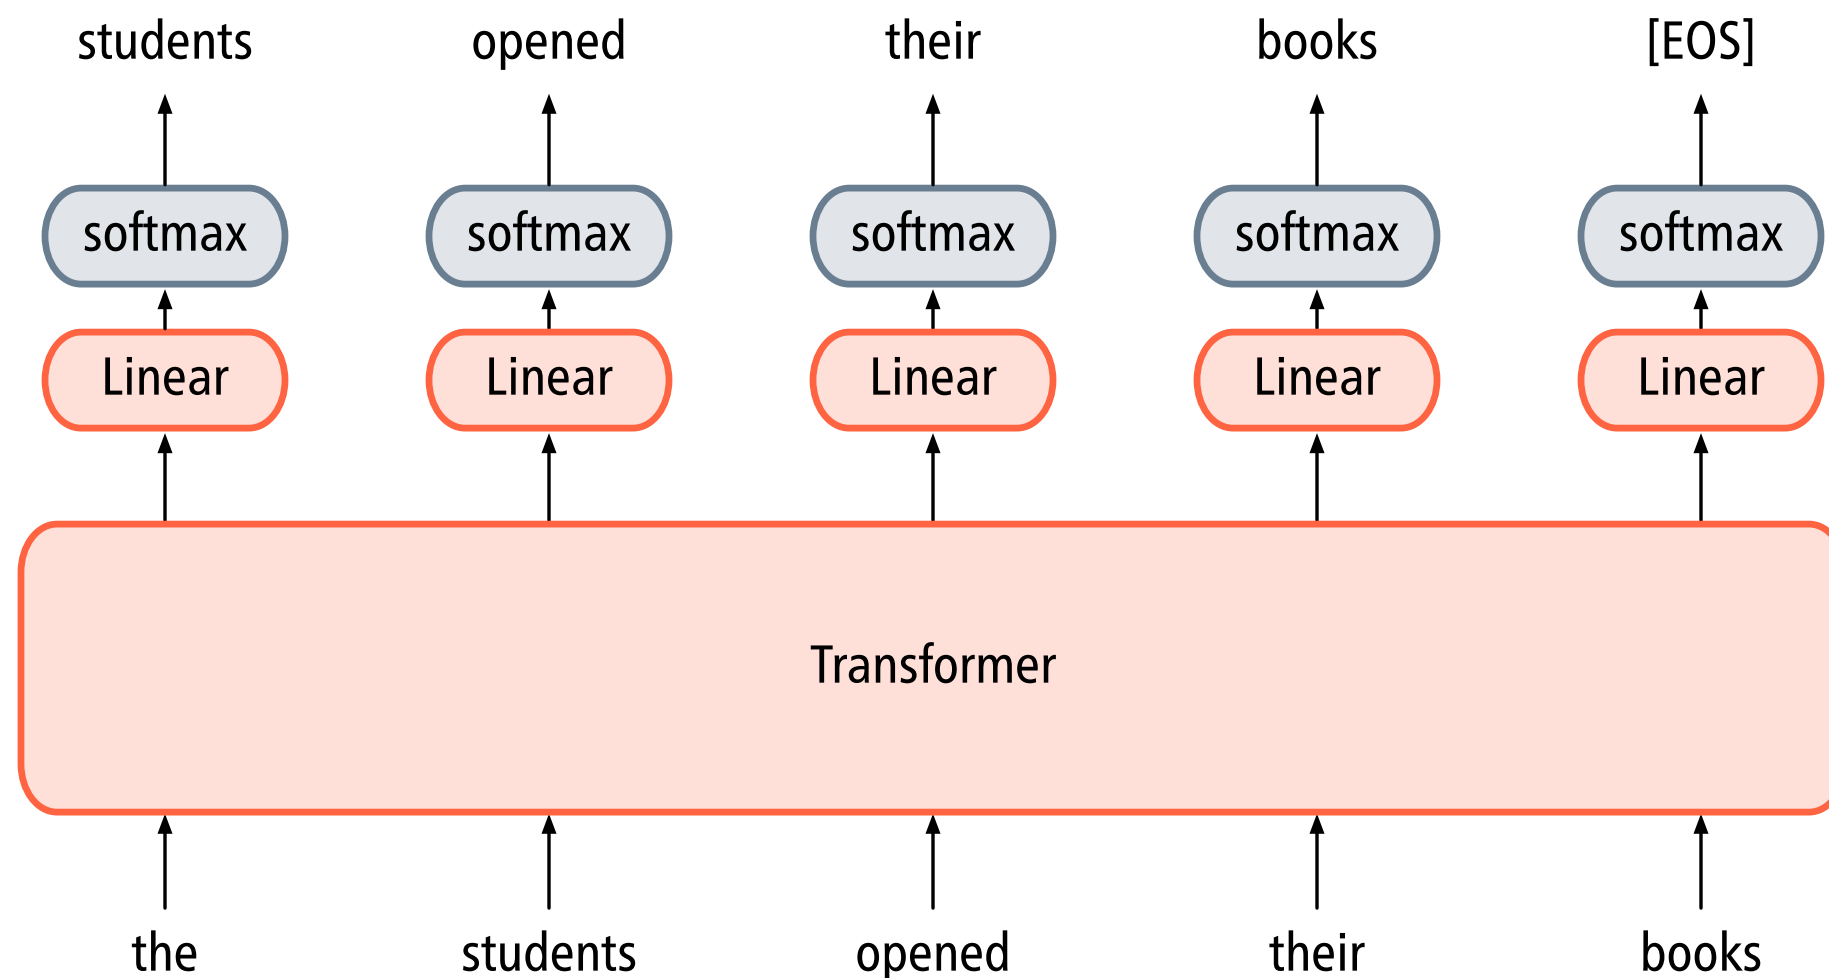
Natural Language Processing

Emergent abilities of LLMs

Marco Kuhlmann

Department of Computer and Information Science


Pretraining and finetuning



Step 1: Pretrain on language modelling

Large quantities of text, general facts about language

Pretraining and finetuning



Step 2: Finetune on specific tasks

Small quantities of labelled data, task-specific knowledge

Model growth


	GPT-1	GPT-2	GPT-3	GPT-4
Number of dimensions	768	1,600	12,288	?
Number of layers	12	48	96	120
Trainable parameters	0.117 B	1.542 B	175 B	1,800 B
Training data size	4 GB	40 GB	570 GB	?

[Radford et al. \(2018\)](#), [Radford et al. \(2019\)](#), [Brown et al. \(2020\)](#)

GPT-1: Effective pretraining

Language modelling is an effective pretraining method for a broad range of tasks in natural language understanding.

[Radford et al. \(2018\)](#)



Effective pretraining

Natural Language Inference (NLI)

Premise: A man inspects the uniform of a figure in some East Asian country.

Hypothesis: The man is sleeping.

Label: contradiction

Question answering

Question: The first postage stamp was made ...

Candidate answers: A. in England, B. in America, C. by Alice, D. in 1910


Method	MNLI +	MNLI -	QNLI	RTE	SNLI
Previous state-of-the-art	80.6	80.1	82.3	61.7	89.3
GPT-1 (Radford et al., 2018)	82.1	81.4	88.1	56.0	89.9

GPT-2: Emergent zero-shot learning

zero-shot learning

the ability of a machine learning model to solve tasks out-of-the-box, with no examples and no gradient updates

[Radford et al. \(2019\)](#)



Zero-shot learning

Radford et al. (2019)

Sequence prediction

Question: Who took the first steps on the moon in 1969?

Answer: **Neil Armstrong**

model output



Question: Which Stanford University alumna co-founded Coursera?

Answer: **Daphne Koller**

Sequence modelling

The trophy doesn't fit into the brown suitcase because it is too large.

it = ... **$p(\text{trophy}) > p(\text{suitcase})$**

The trophy doesn't fit into the brown suitcase because it is too small.

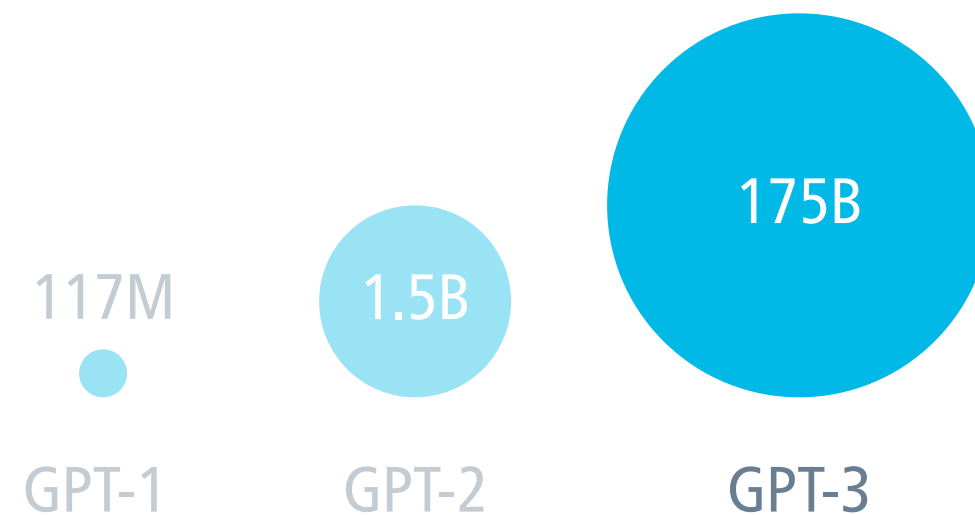
it = ... **$p(\text{trophy}) < p(\text{suitcase})$**

GPT-3: Emergent in-context learning

in-context learning

the ability of a machine learning model to learn tasks from a few examples, with no gradient updates

[Brown et al. \(2020\)](#)



In-context learning

[Brown et al. \(2020\)](#); examples from Jesse Mu

Word unscrambling

gaot => goat

sakne => snake

brid => bird

fsih => fish

dcuk => duck

cmihp => **chimp**

in-context learning



Machine translation

thanks => merci

hello => bonjour

mint => menthe

wall => mur

otter => loutre

bread => **pain**

in-context learning



Chain-of-thought prompting

Wei et al. (2022)

Standard prompting

Q: Roger has 5 tennis balls. He buys 2 more cans of tennis balls. Each can has 3 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?

A: **The answer is 11.**

Chain-of-thought prompting

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

A: Roger started with 5 balls. 2 cans of 3 balls each is 6 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?

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 had $3 + 6 = 9$. The answer is 9.**

Zero-shot chain-of-thought prompting

Kojima et al. (2022)

Chain-of-thought prompting

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

A: Roger started with 5 balls. 2 cans of 3 balls each is 6 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?

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 had $3 + 6 = 9$. The answer is 9.

Zero-shot chain-of-thought prompting

Q: Roger has 5 tennis balls. He buys 2 more cans of tennis balls. Each can has 3 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?

A: Let's think step by step. **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 had $3 + 6 = 9$. The answer is 9.**

Prompt engineering

designed by a LM



Prompt	Accuracy
Let's work this out in a step by step way to be sure we have the right answer.	82.0
Let's think step by step.	78.7
First, ...	77.3
Let's think about this logically.	74.5
Let's solve this problem by splitting it into steps	72.2
Let's be realistic and think step by step.	70.8
Let's think like a detective step by step.	70.3
(Zero-shot)	17.7