



GPT-3

Hung-yi Lee 李宏毅

GPT-3

Language Models are Few-Shot Learners

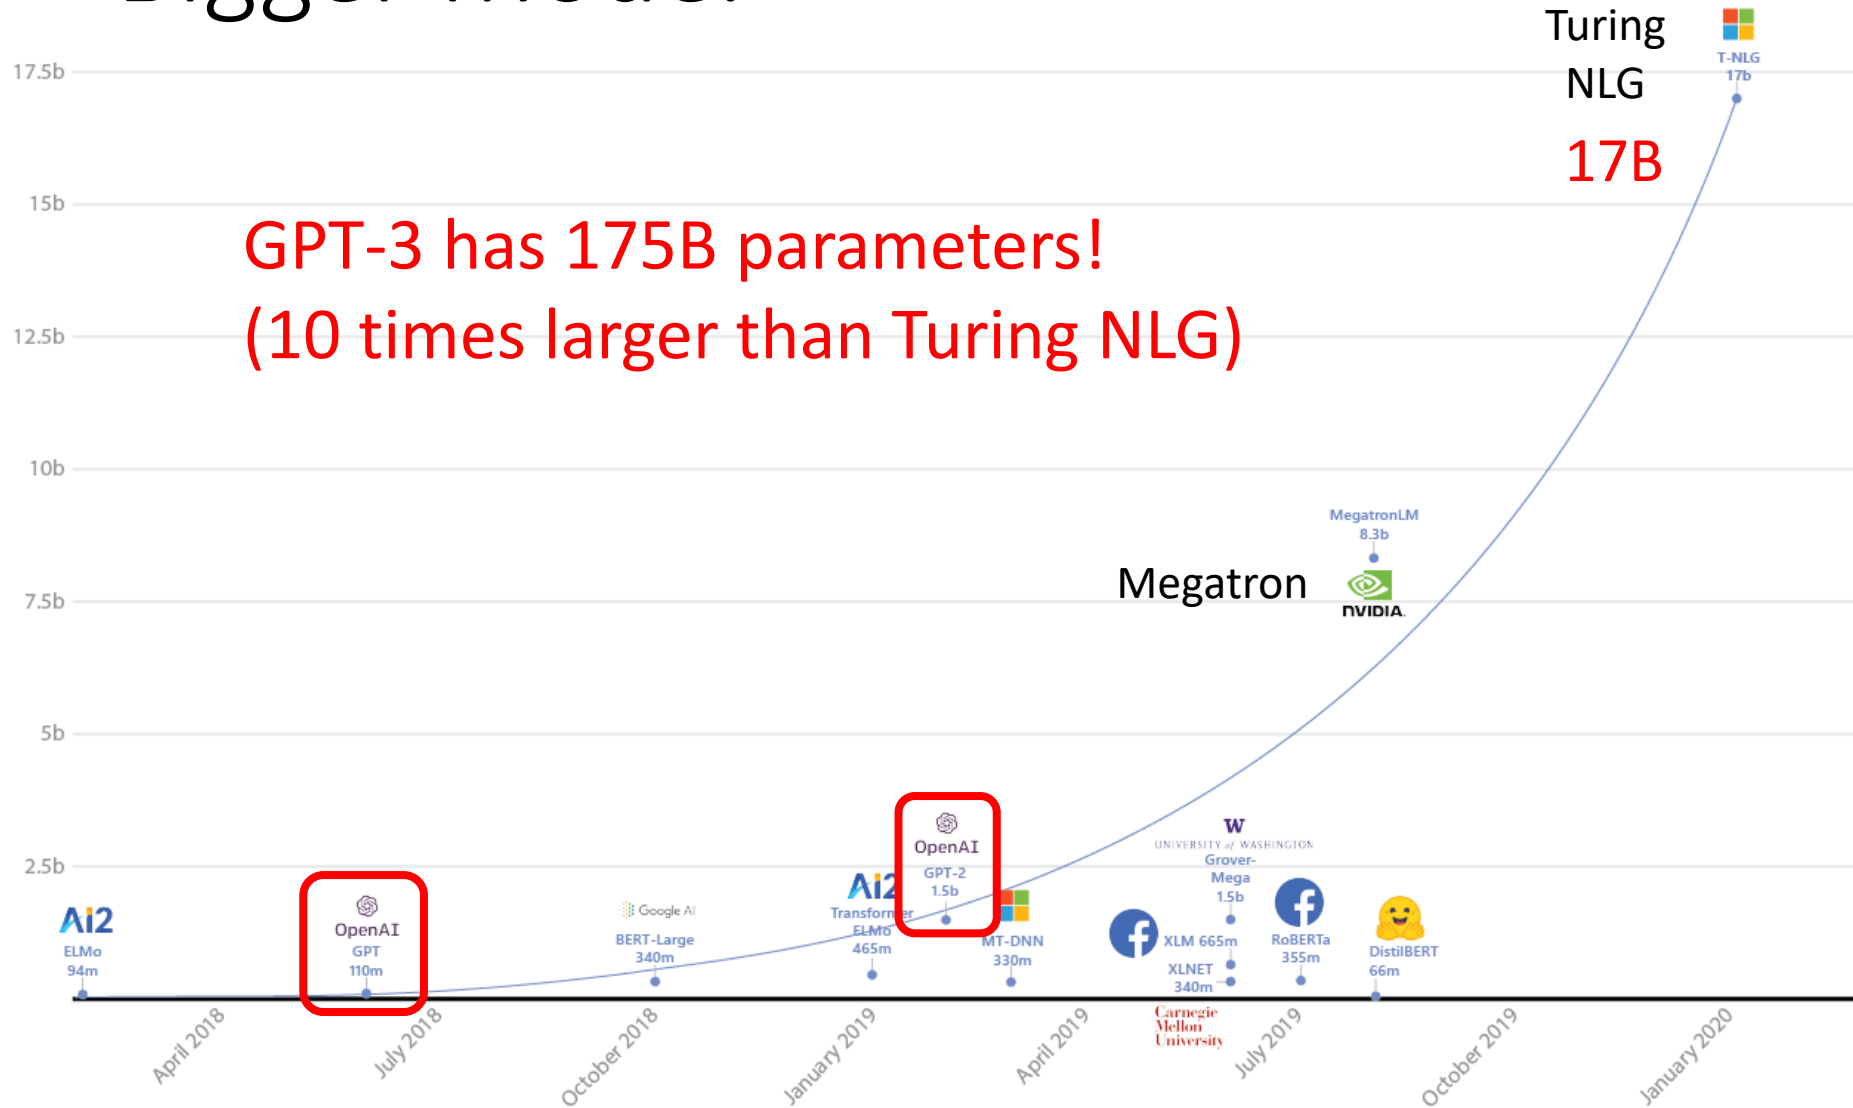
Tom B. Brown*	Benjamin Mann*	Nick Ryder*	Melanie Subbiah*	
Jared Kaplan†	Prfulla Dhariwal	Arvind Neelakantan	Pranav Shyam	Girish Sastry
Amanda Askell	Sandhini Agarwal	Ariel Herbert-Voss	Gretchen Krueger	Tom Henighan
Rewon Child	Aditya Ramesh	Daniel M. Ziegler	Jeffrey Wu	Clemens Winter
Christopher Hesse	Mark Chen	Eric Sigler	Mateusz Litwin	Scott Gray
Benjamin Chess		Jack Clark	Christopher Berner	
Sam McCandlish	Alec Radford	Ilya Sutskever	Dario Amodei	

OpenAI



Bigger Model

**GPT-3 has 175B parameters!
(10 times larger than Turing NLG)**





假設 ELMO 的參數量是長 30 公分的尺
GPT-3 的參數量大約是 ELMO 的 2000 倍
那麼 GPT-3 比台北 101 還高



此處著實

GPT-3 是來自於暗黑大陸的模型



Elliot Turner
@eturner303

Reading the OpenAI GPT-3 paper. Impressive performance on many few-shot language tasks. The cost to train this 175 billion parameter language model appears to be staggering: Nearly \$12 million

自然语言处理

炫富

核武器

深度学习 (Deep Learning)

OpenAI

如何评价1700亿参数的GPT-3?

<https://www.zhihu.com/q>



minimaxir commented 26 days ago • edited

...I'm not sure if it's even possible for the 175B model to be distributed in a reasonable manner.

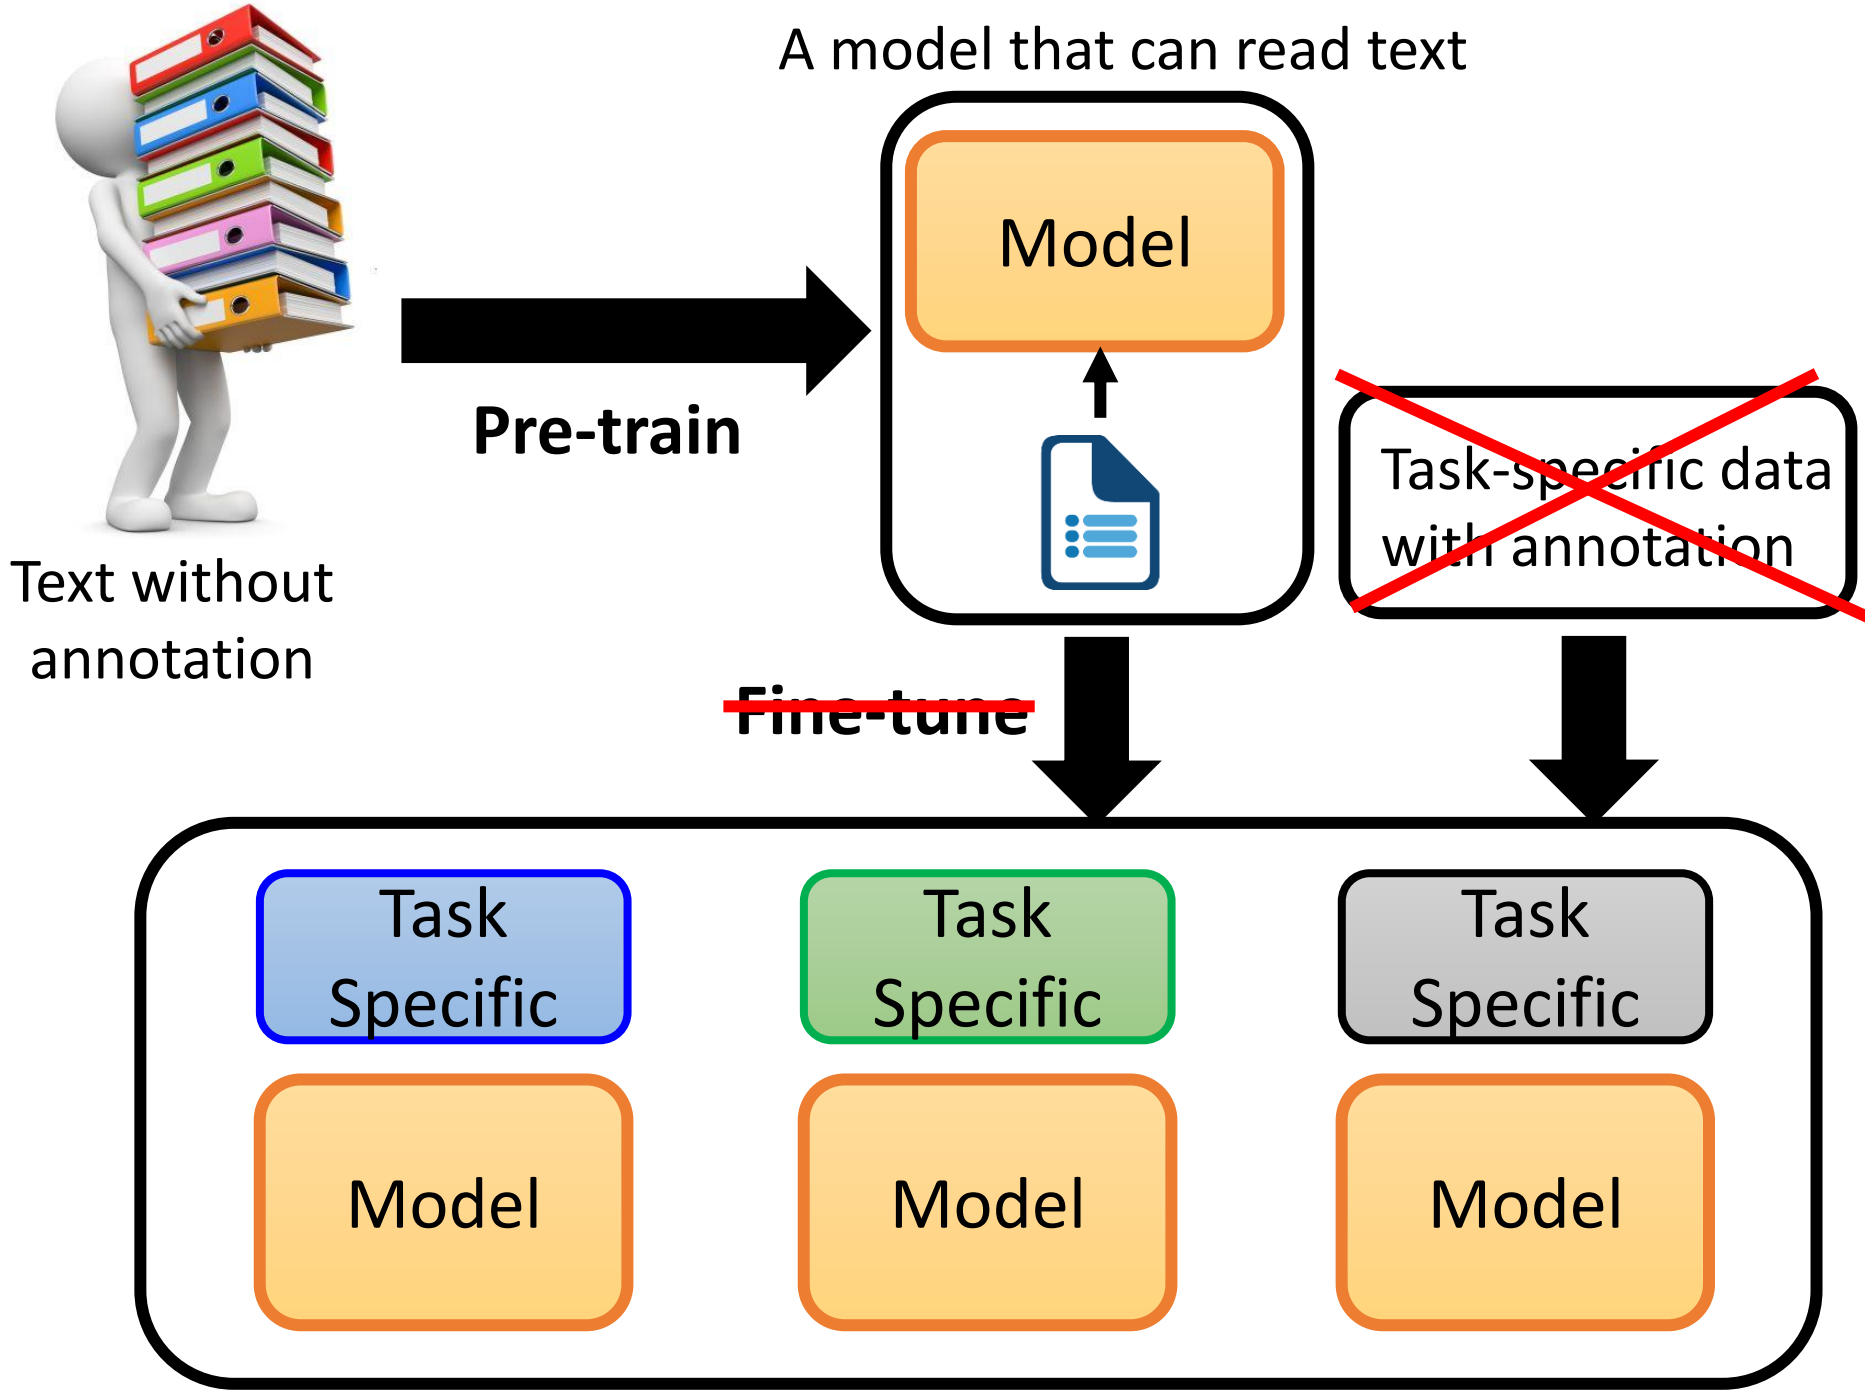
The size of the 1.5B GPT-2 model was about 6GB on disk, which would imply that the 175B model is *at least* 700GB!



Sebastian Gehrmann
@sebgehr

The ELMo paper? 15 pages.
BERT? 16 pages.
GPT-2? 24 pages.
T5? 53 pages.
GPT-3?? 72 pages! arxiv.org/pdf/2005.14165...

github.com/openai/gpt-3/issues/1



GPT 系列的野望

第一部份：詞彙和結構

本部份共 15 題，每題含一個空格。請就試題冊上 A、B、C、D 四個選項中選出最適合題意的字或詞，標示在答案紙上。

題型說明

例：

It's eight o'clock now. Sue _____ in her bedroom.

- A. study
- B. studies
- C. studied
- D. is studying

正確答案為 D，請在答案紙上塗黑作答。

少數範例

“In-context” Learning

Few-shot Learning

(no gradient descent)

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

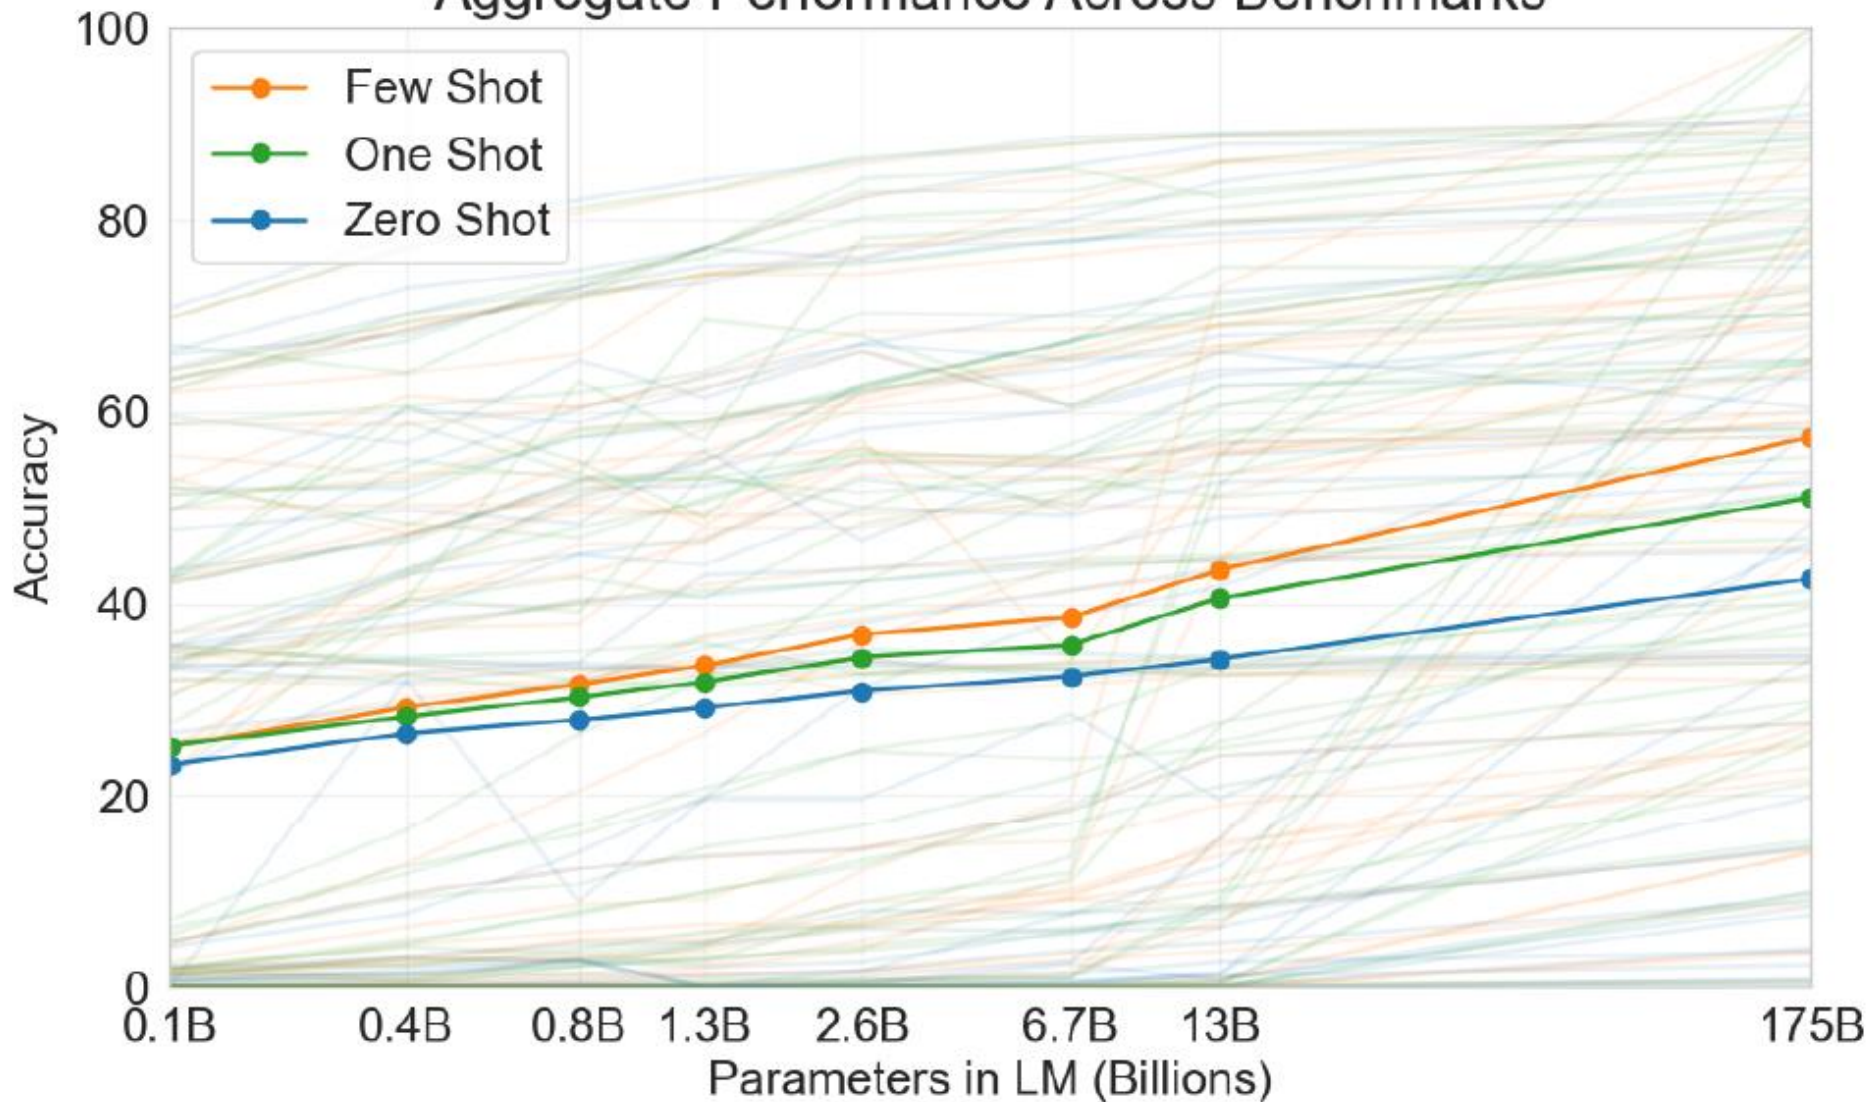
One-shot Learning

1	Translate English to French:	← task description
2	sea otter => loutre de mer	← example
3	cheese =>	← prompt

Zero-shot Learning

1	Translate English to French:	← task description
2	cheese =>	← prompt

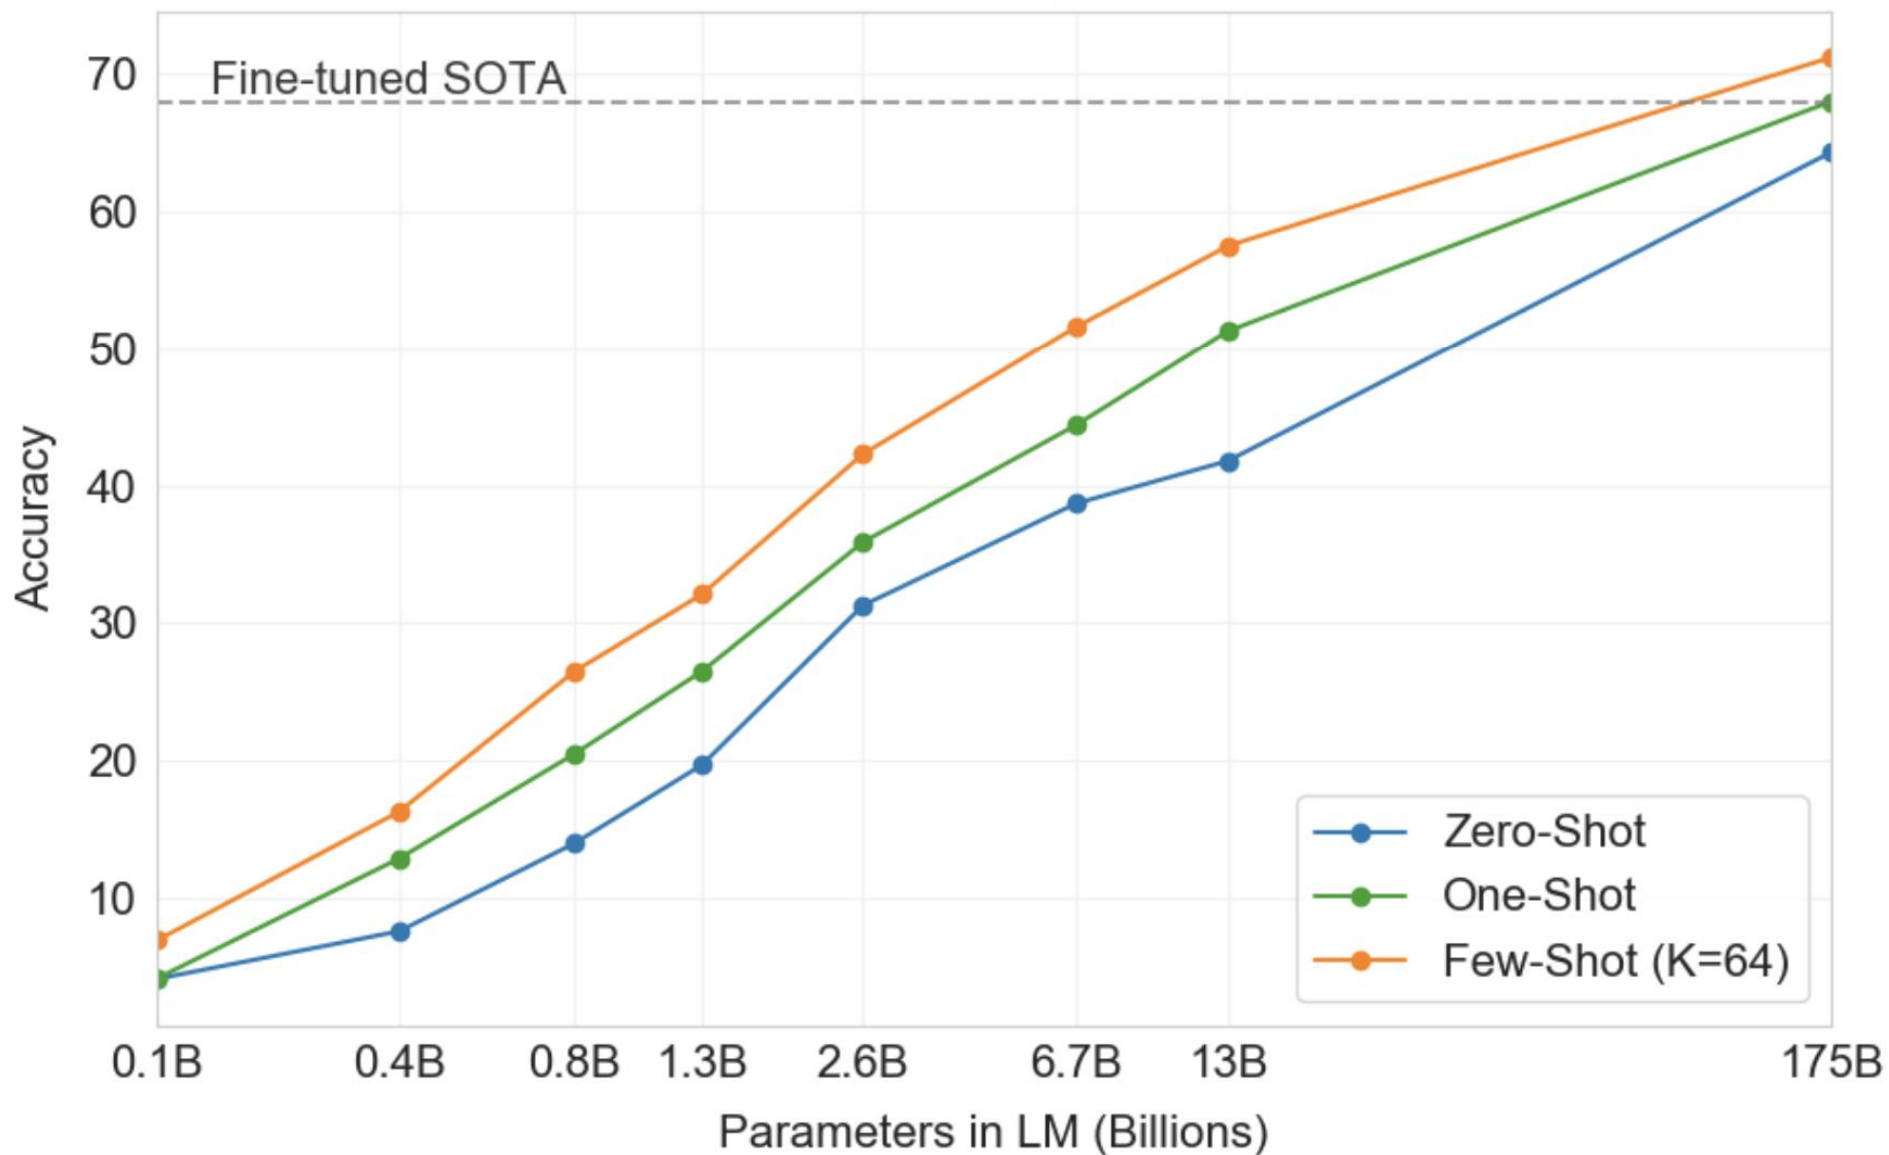
Aggregate Performance Across Benchmarks



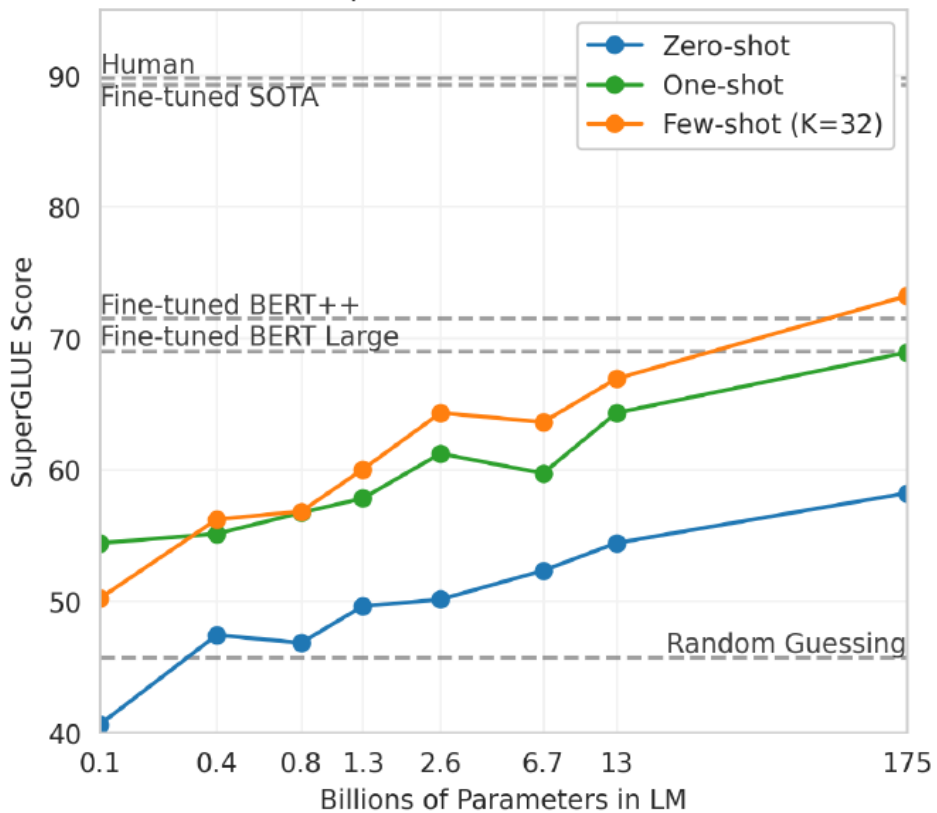
Average of **42** tasks

Closed Book QA

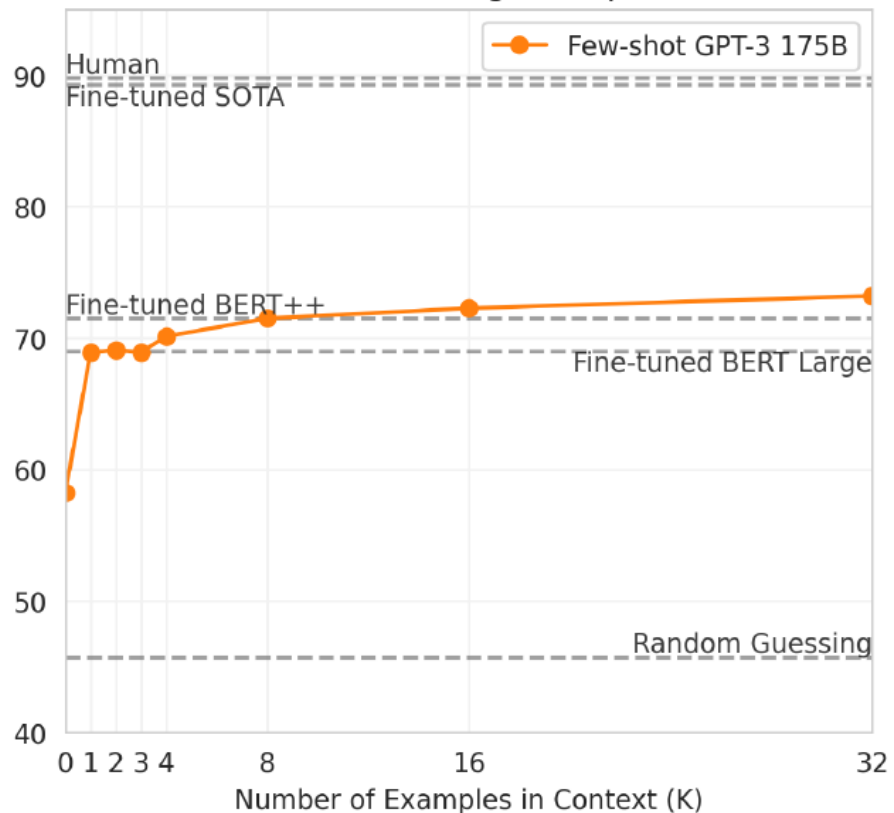
TriviaQA




SuperGLUE Performance



In-Context Learning on SuperGLUE



Human ability to detect model generated news articles



A "whatpu" is a small, furry animal native to Tanzania. An example of a sentence that uses the word whatpu is:

We were traveling in Africa and we saw these very cute whatpus.

To do a "farduddle" means to jump up and down really fast. An example of a sentence that uses the word farduddle is:

One day when I was playing tag with my little sister, she got really excited and she started doing these crazy farduddles.

A "yalubalu" is a type of vegetable that looks like a big pumpkin. An example of a sentence that uses the word yalubalu is:

I was on a trip to Africa and I tried this yalubalu vegetable that was grown in a garden there. It was delicious.

A "Burringo" is a car with very fast acceleration. An example of a sentence that uses the word Burringo is:

In our garage we have a Burringo that my father drives to work every day.

A "Gigamuru" is a type of Japanese musical instrument. An example of a sentence that uses the word Gigamuru is:

I have a Gigamuru that my uncle gave me as a gift. I love to play it at home.


To "screeg" something is to swing a sword at it. An example of a sentence that uses the word screeg is:

We screeghed at each other for several minutes and then we went outside and ate ice cream.

Figure 3.16: Representative GPT-3 completions for the few-shot task of using a new word in a sentence. Boldface is GPT-3's completions, plain text is human prompts. In the first example both the prompt and the completion are provided by a human; this then serves as conditioning for subsequent examples where GPT-3 receives successive additional prompts and provides the completions. Nothing task-specific is provided to GPT-3 other than the conditioning shown here.

Context → Please unscramble the letters into a word, and write that word:
r e ! c . i p r o . c a / l =

Target Completion → reciprocal

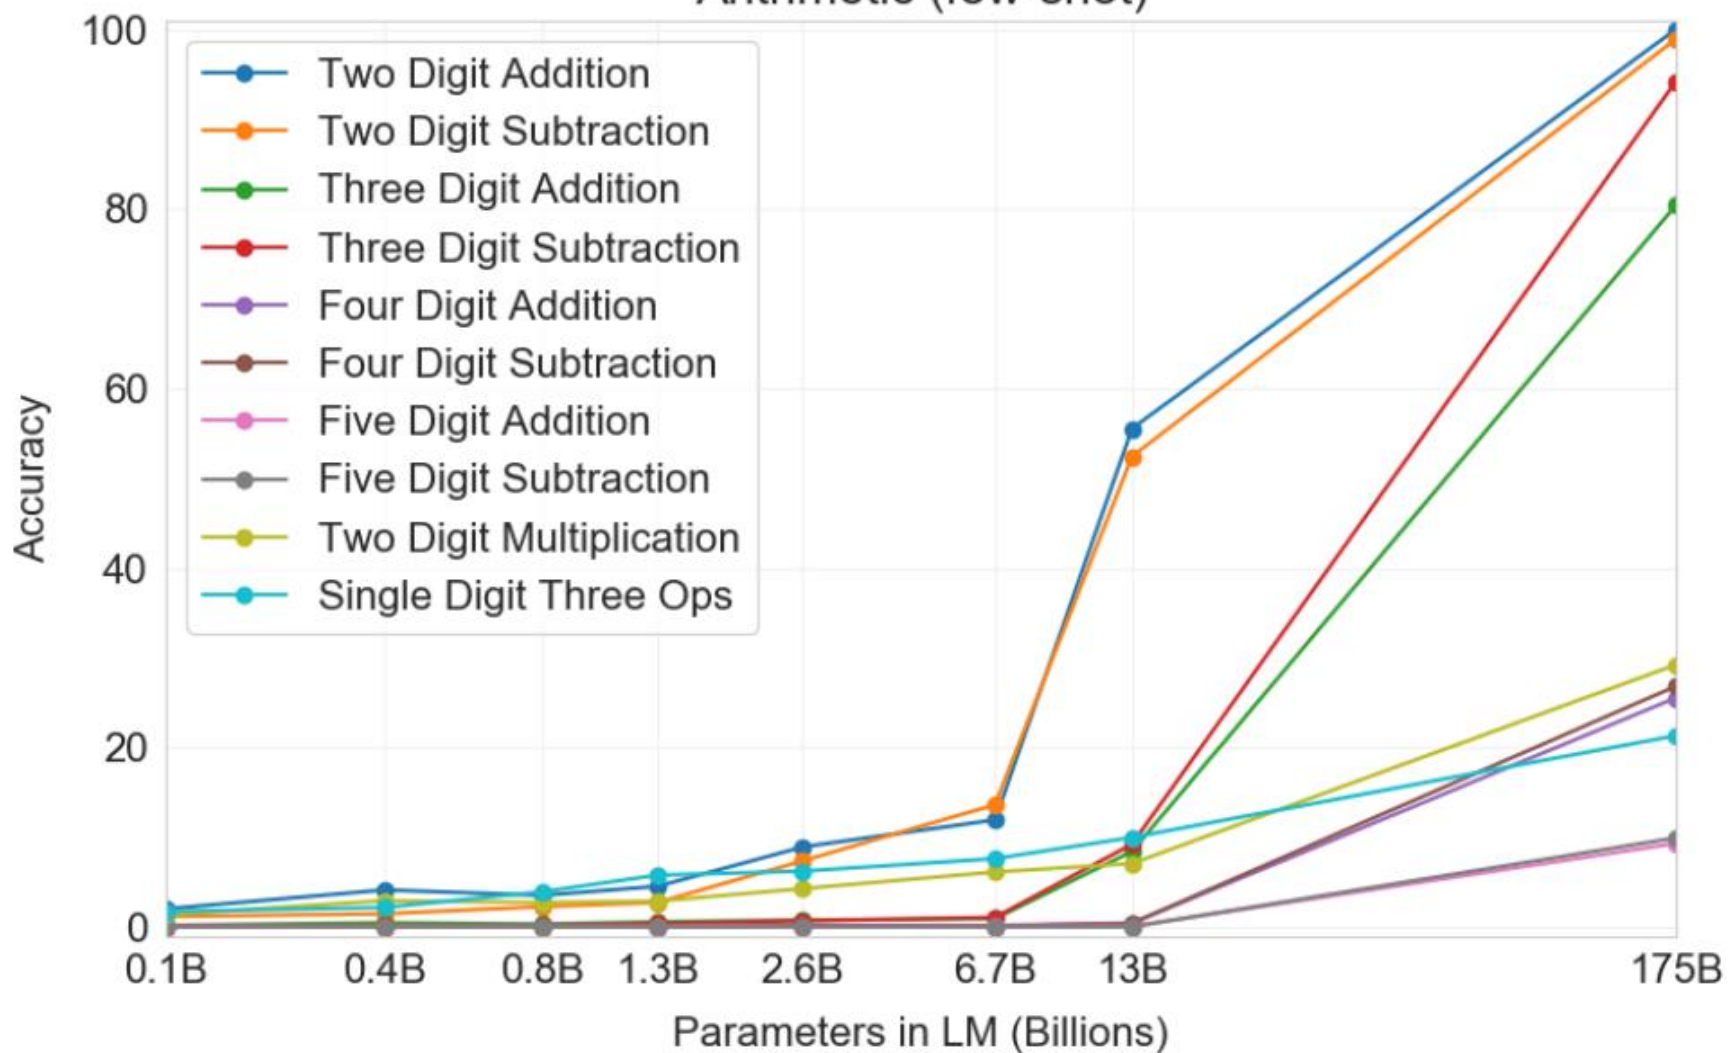


Context → Q: What is 17 minus 14?

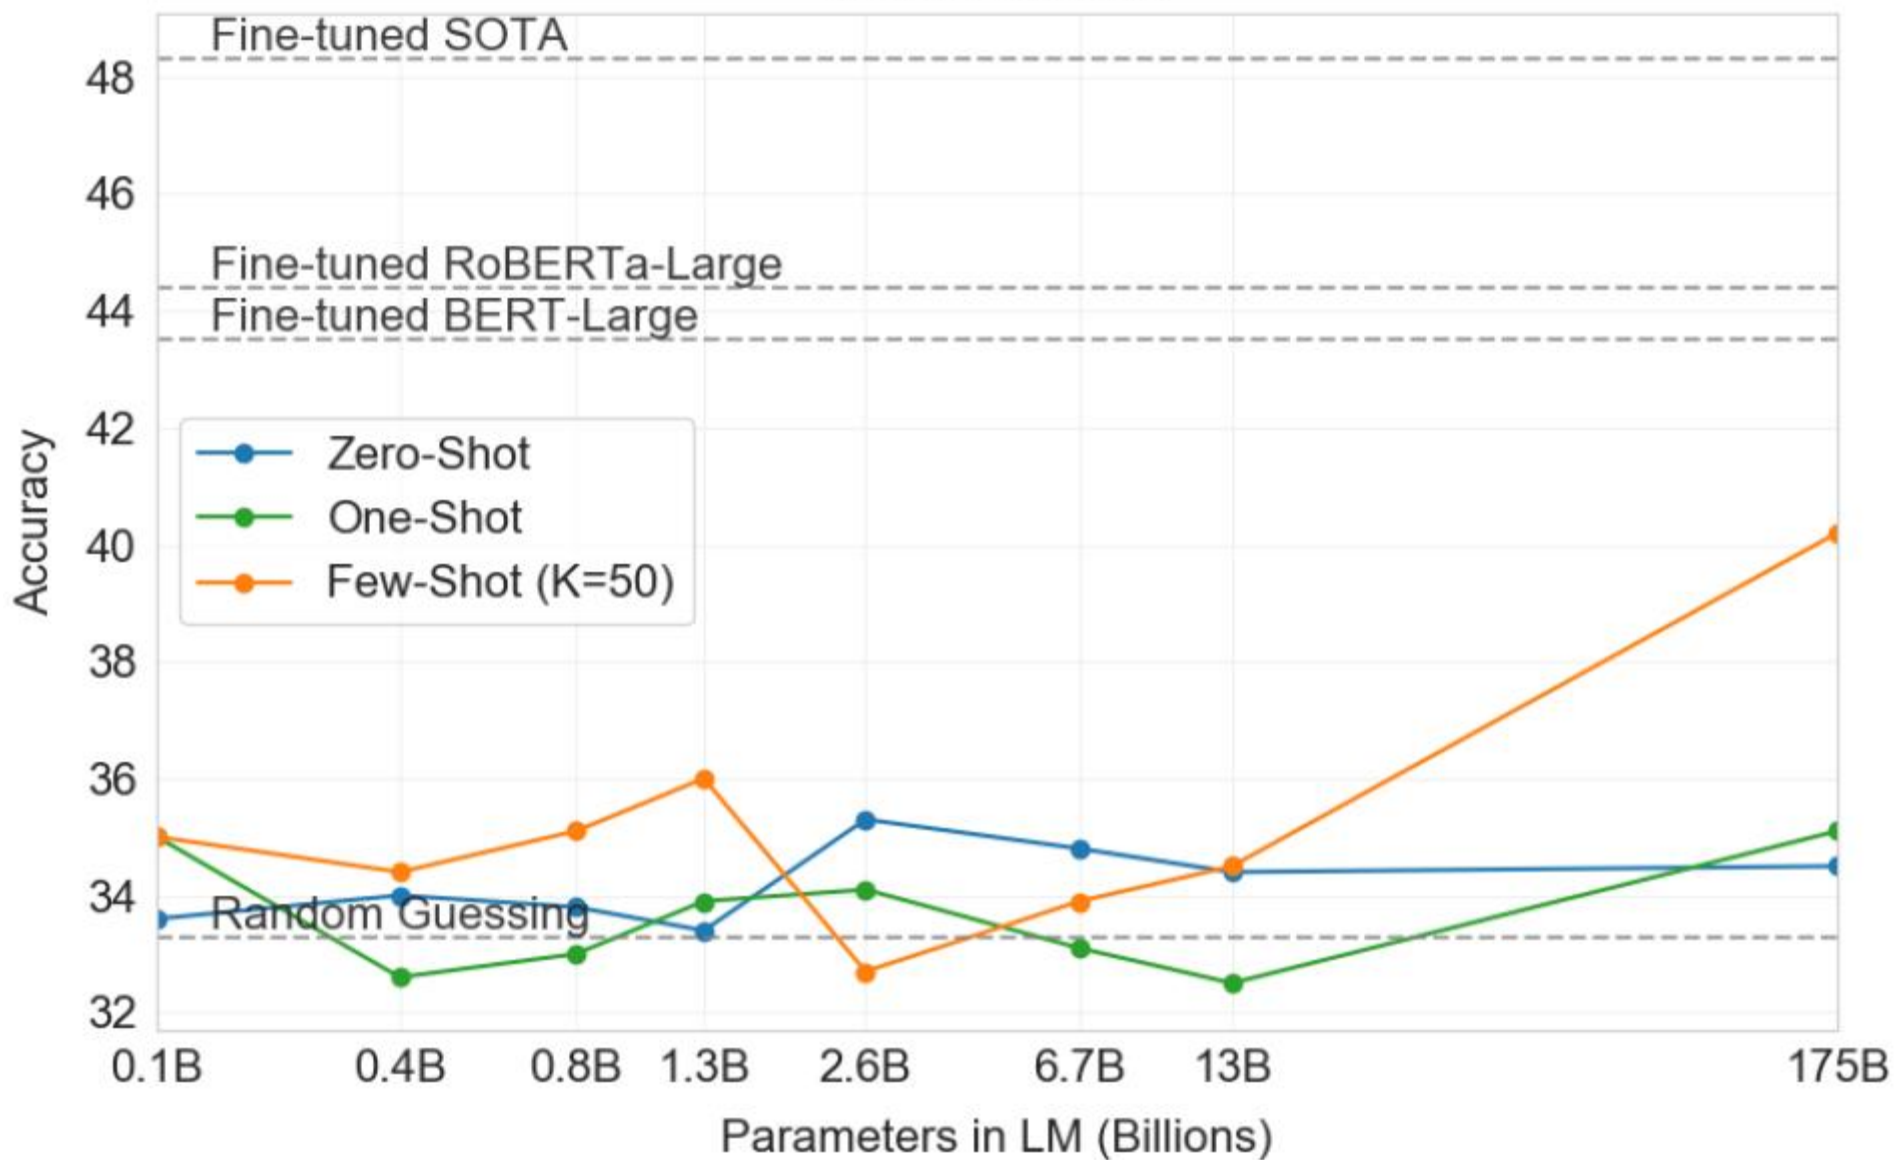
A:

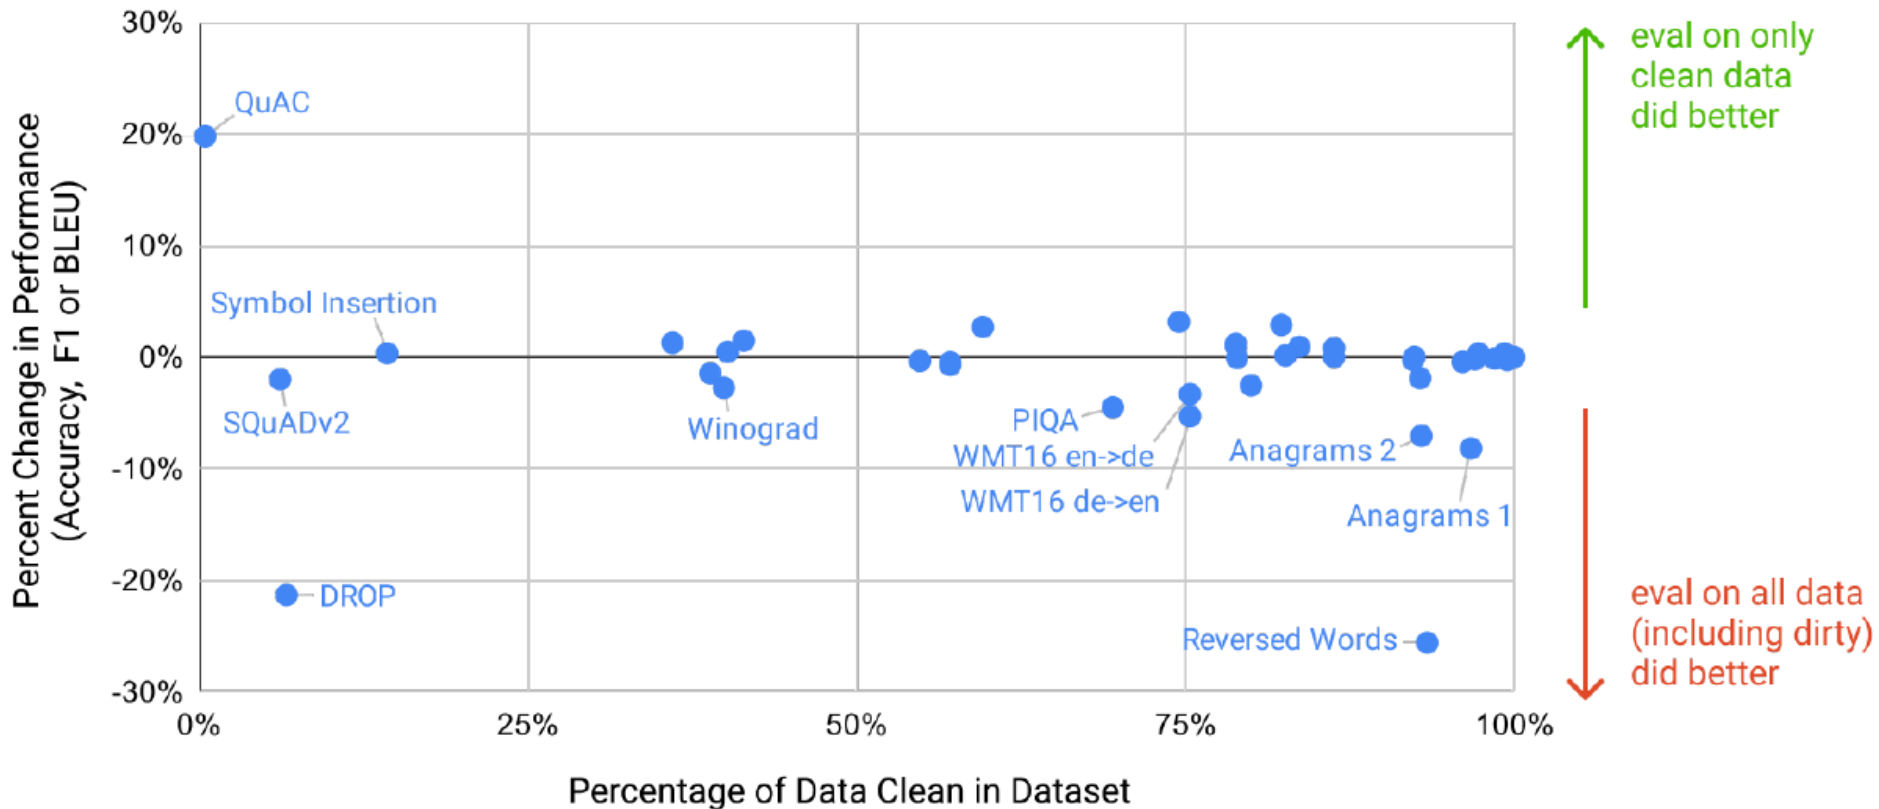
Target Completion → 3

Arithmetic (few-shot)



ANLI Round3





We initially tried to address the issue of contamination by proactively searching for and attempting to remove any overlap between our training data and the development and test sets of all benchmarks studied in this paper. Unfortunately, a bug resulted in only partial removal of all detected overlaps from the training data. Due to the cost of training, it wasn't feasible to retrain the model. To address this, we investigate in detail how the remaining detected overlap impacts results.

Turing A

<http://rowanz>

關於感情問題



我一律建議分手

所有關於感情的問題



一律建議分手



我一律建議分手

wanted
azy. Fo
er neat
didn't
etely c



感情問題我一律建議分手

情的問題
建議分手

Not sure what you

Evaluating Machines by their Real-World Language Use

Rowan Zellers[♦] Ari Holtzman[♦] Elizabeth Clark[♦]
Lianhui Qin[♦] Ali Farhadi[♦] Yejin Choi[♥]

[♦]Paul G. Allen School of Computer Science & Engineering, University of Washington
[♥]Allen Institute for Artificial Intelligence
rowanzellers.com/advice



I have to do a dissection for my high school class, but I'm distressed by dead animals. Last time we dissected an animal in class, I had a panic attack. I asked my teacher for another assignment, but she refused. I don't want to play a 'victim' card, but I don't know what to do. Help!



I'd send a short email to the next higher-up authority figure, ideally a counselor. Be forthright; it's the best approach when self-advocating as a student.




Thanks!

Helpful



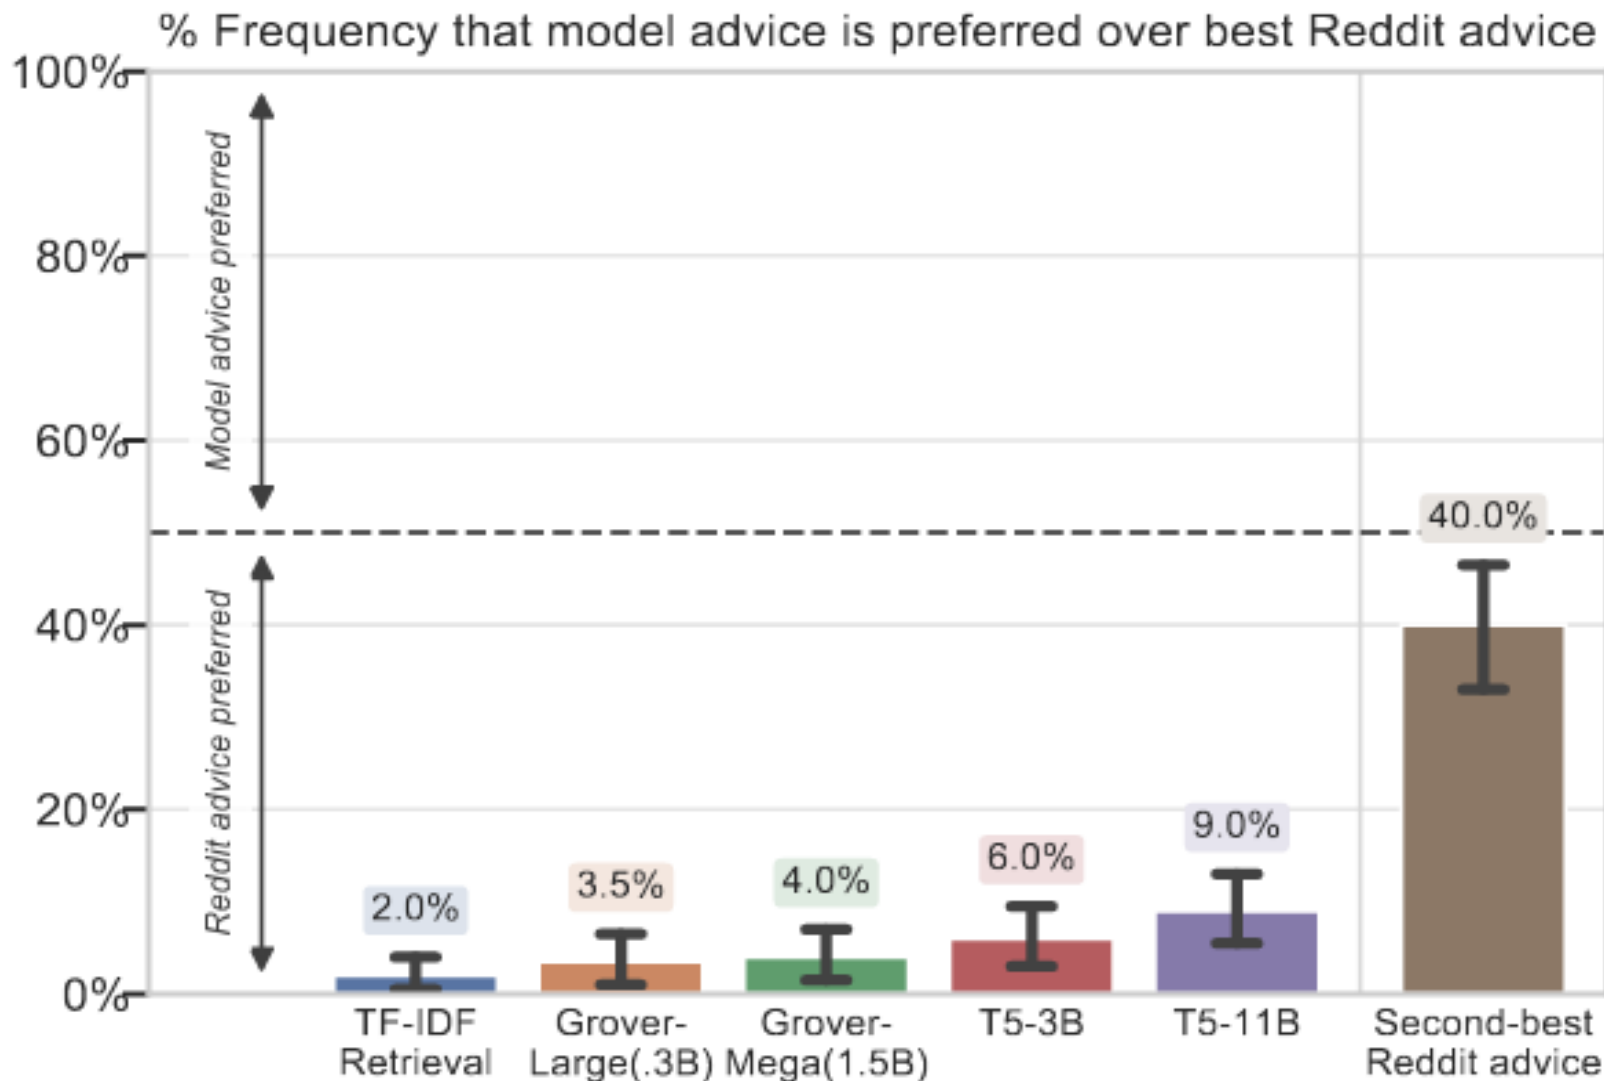
Go to your teacher and say "I'm asking you to do a project that requires me to see dead animals. This is a dealbreaker." If she doesn't concede, tell your principal about your trauma.




??????

Not helpful




Turing Advice Challenge





raster order

My Favorite Ones



Source of image: <https://openai.com/blog/image-gpt/>