

Large Language Models

Recap: Self-Supervised

Skills needed to classify dogs

- Identify different textures
- Identify different colors
- Identify which pixels are parts of an object
- Understand parts of objects which make up the whole
- Understand the context of the object/animal in the image
- Understand lighting conditions
- Understand objects close up/far away
- Have common-sense reasoning skills
- Learn which features are associated with which dog

Recap: Self-Supervised

General visual
modelling skills
(Can be learned with
Self Supervised)

Specific (Supervised)

- Identify different textures
- Identify different colors
- Identify which pixels are parts of an object
- Understand parts of objects which make up the whole
- Understand the context of the object/animal in the image
- Understand lighting conditions
- Understand objects close up/far away
- Have common-sense reasoning skills
- Learn which features are associated with which dog

Self-Supervised for Language

Skills needed to classify book genres

Knowledge of words/letters

Knowledge of grammar

Meanings of words

Understanding context of words

Keeping track of entities over time

Understanding expressions/idioms

Understanding tone

Learning which features are associated with each genre

Self-Supervised for Language

General language
modelling skills
(Can be learned with
Self Supervised)

Specific (Supervised)

Knowledge of words/letters

Knowledge of grammar

Meanings of words

Understanding context of words

Keeping track of entities over time

Understanding expressions/idioms

Understanding tone

Learning which features are associated with each genre

Self-Supervised for Language

General language
modelling skills
(Can be learned with
Self Supervised)

Knowledge of words/letters
Knowledge of grammar
Meanings of words
Understanding context of words
Keeping track of entities over time
Understanding expressions/idioms
Understanding tone

Specific (Supervised) | Learning which features are associated with each genre

Want: model with general understanding of language (Language model!)

LLMs

Building LLMs: Pre-training objectives + architectures

- Encoder only
- Decoder only
- Encoder Decoder

GPT

Gradient-Free Performance Improvement

LLMs

Building LLMs: Pre-training objectives + architectures

- **Encoder only**
- **Decoder only**
- **Encoder Decoder**

GPT

Gradient-Free Performance Improvement

Last time

Pre-training tasks

rotation



[Gidaris et al. 2018](#))

in-painting



[Pathak et al., 2016z](#)

colorization



Source: [Google AI blog post](#)

1. Solving the pretext tasks allow the model to learn good features.
2. We can automatically generate labels for the pretext tasks.

Last time

Pre-training tasks

Common recipe in vision:

1. Take original data
2. Remove/obscure information
3. Ask model to get back to the original instance

colorization



Source: [Google AI blog post](#)

rotation



[Gidaris et al. 2018](#))

in-painting



[Pathak et al., 2016z](#)

Last time

Pre-training tasks

rotation



90° rotation

in-painting



colorization



What to use as pre-training task for language?

Encoder only LLMs

It's cold today! Don't forget to wear a _____.

The _____ is a popular tourist attraction in Seattle.

I missed ___ bus.

I had 3 pencils and lost one so now I have _____ pencils.

Encoder only LLMs

It's cold today! Don't forget to wear a jacket / coat / sweater.

The Space Needle is a popular tourist attraction in Seattle.

I missed the bus.

I had 3 pencils and lost one so now I have 2 / two pencils.

Encoder only LLMs

It's cold today! Don't forget to wear a jacket / coat / sweater.

Common Sense

The Space Needle is a popular tourist attraction in Seattle.

Factual knowledge

I missed the bus.

Grammar

I had 3 pencils and lost one so now I have 2 / two pencils.

Math/ Reasoning

LLMs

Encoder Only:

I love cake

Decoder Only:

I love

Encoder-Decoder:

I love cake

me gusta

LLMs

Encoder Only:

I love cake

Decoder Only:

I love

Encoder-Decoder:

I love cake


me gusta

ELMo (Embeddings from Language Models)


Pre-training task

This wall needs another coat of paint

Predict words based on previous words



Predict words based on following words




Peters et al. Deep contextualized word representations. 2018.

ELMo (Embeddings from Language Models)

Application to downstream tasks

This wall needs another coat of paint




Peters et al. Deep contextualized word representations. 2018.

ELMo (Embeddings from Language Models)

Application to downstream tasks

This wall needs another coat of paint

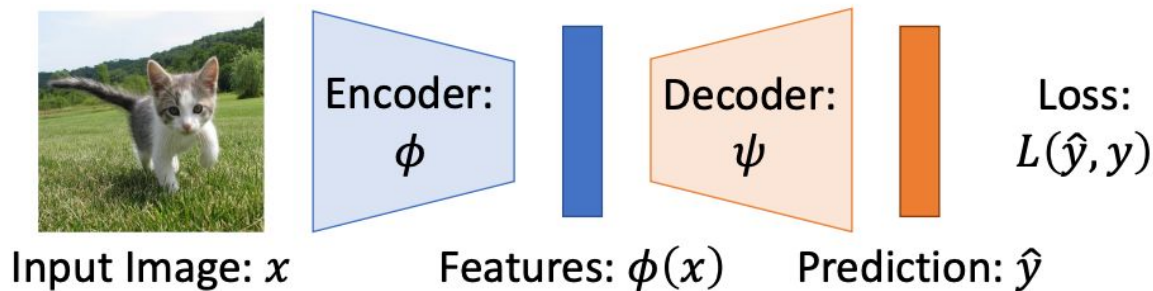


Fine-tune


Peters et al. Deep contextualized word representations. 2018.

Vision: Pre-train and then fine-tune

Step 1: Pretrain a network on a pretext task that doesn't require supervision




Step 2: Transfer encoder to downstream tasks via linear classifiers, KNN, finetuning

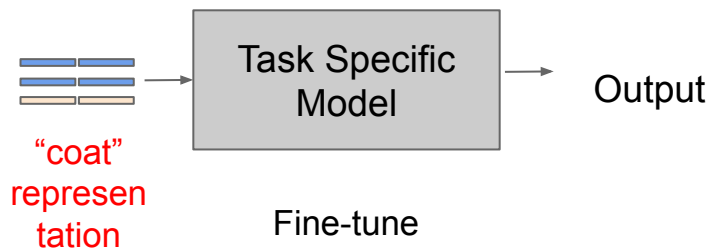


Language: Pre-train and then fine-tune

Step 1: Pretrain a network on a pretext task that doesn't require supervision



Step 2: Transfer encoder to downstream tasks via linear classifiers, KNN, finetuning



Downstream tasks:
Sentiment classification, NLI,
...

Encoder only LLMs

It's cold today! Don't forget to wear a coat.

This wall needs another coat of paint

coat =

$$\begin{bmatrix} 0.286 \\ 0.792 \\ -0.177 \\ -0.107 \\ 0.109 \\ -0.542 \\ 0.349 \\ 0.271 \end{bmatrix}$$

Glove embeddings use the same vector for every instance of a word, no matter the context!

ELMo Results

TASK	PREVIOUS SOTA		OUR BASELINE	ELMo + BASELINE	INCREASE (ABSOLUTE/ RELATIVE)
SQuAD	Liu et al. (2017)	84.4	81.1	85.8	4.7 / 24.9%
SNLI	Chen et al. (2017)	88.6	88.0	88.7 \pm 0.17	0.7 / 5.8%
SRL	He et al. (2017)	81.7	81.4	84.6	3.2 / 17.2%
Coref	Lee et al. (2017)	67.2	67.2	70.4	3.2 / 9.8%
NER	Peters et al. (2017)	91.93 \pm 0.19	90.15	92.22 \pm 0.10	2.06 / 21%
SST-5	McCann et al. (2017)	53.7	51.4	54.7 \pm 0.5	3.3 / 6.8%

Source: [Peters et al. Deep contextualized word representations. 2018.](#)

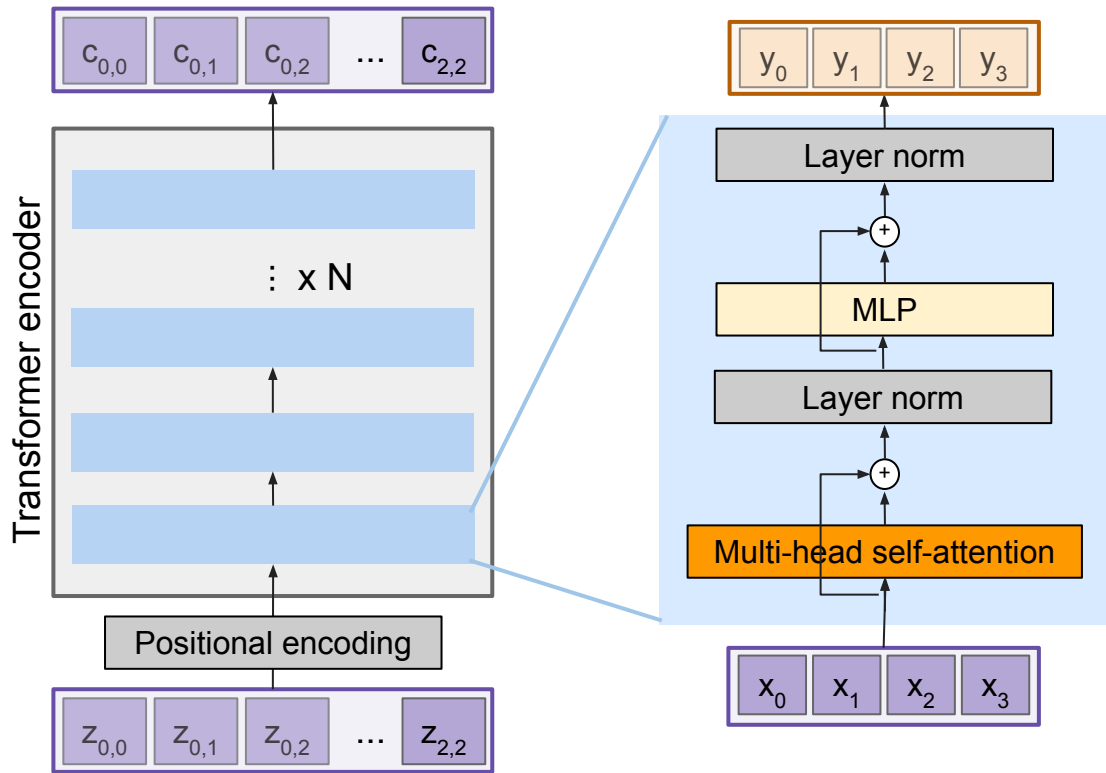
ELMo

Used RNNs... what about transformers?

Let's combine:

- (1) ELMo approach of gathering general knowledge about world
- (2) Powerful multi-purpose architecture that is the transformer

Recall: Transformer encoder block



Transformer Encoder Block:

Inputs: Set of vectors x

Outputs: Set of vectors y

Self-attention is the only interaction between vectors.

Layer norm and MLP operate independently per vector.

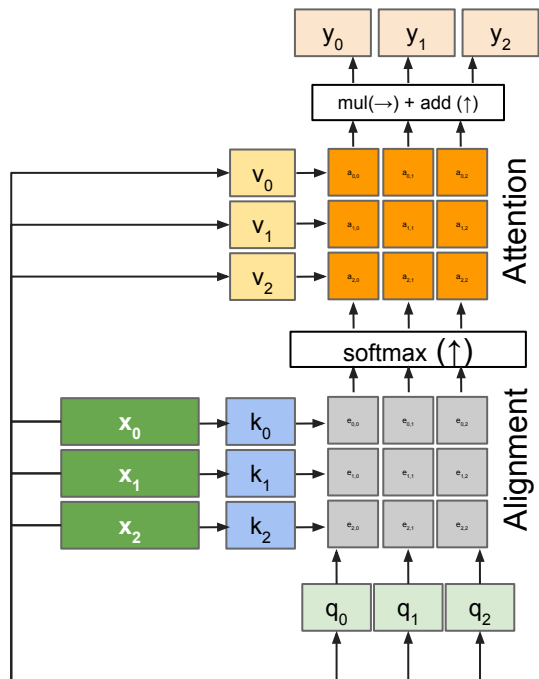
Highly scalable, highly parallelizable, but high memory usage.

Vaswani et al, "Attention is all you need", NeurIPS 2017

Encoder Only: Bert (Bidirectional Encoder Representations from Transformers)

Input: Text sequence

Output: Feature Vector




Outputs:
context vectors: \mathbf{y} (shape: D_v)

Devlin et al. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. 2018.

Encoder Only: Bert

Input: Text sequence

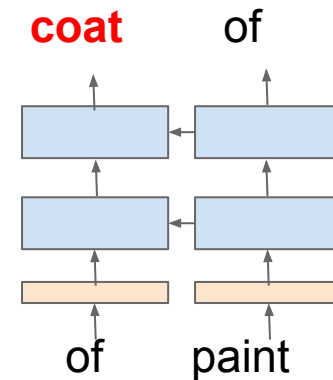
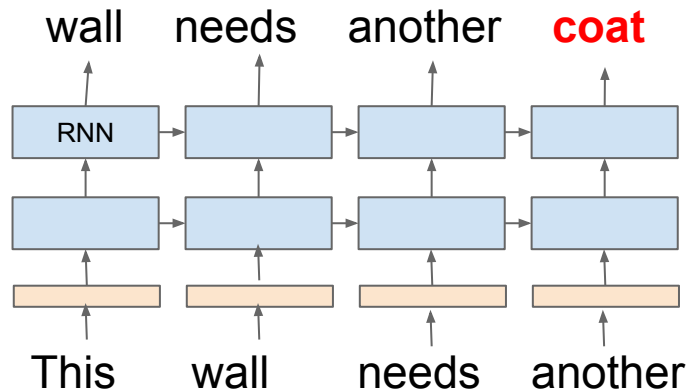
Output: Feature Vector



Outputs:
context vectors: y (shape: D_v)

Devlin et al. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. 2018.

ELMo (Embeddings from Language Models)

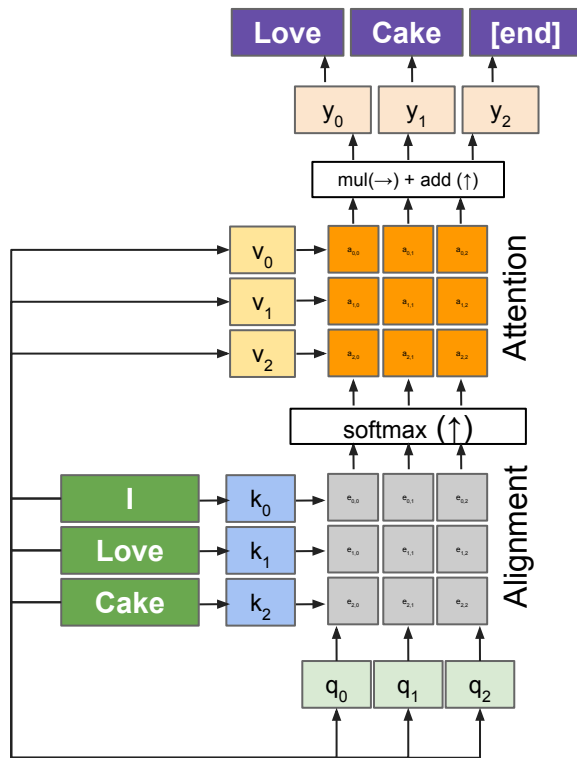


Encoder Only: Bert

Input: Text sequence

Output: Feature Vector

What information do the y vectors contain?



Outputs:
context vectors: y (shape: D_v)

Devlin et al. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. 2018.

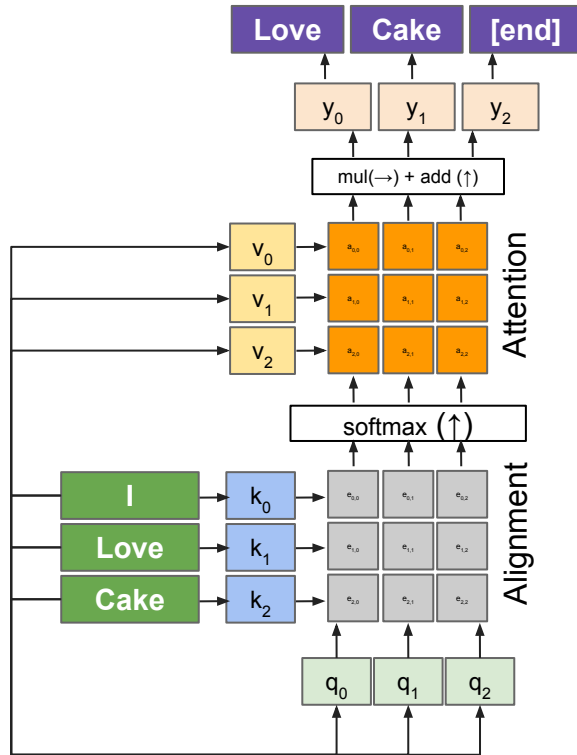
Encoder Only: Bert

Input: Text sequence

Output: Feature Vector

What information do the y vectors contain?

Just copying input



Outputs:
context vectors: y (shape: D_v)

Devlin et al. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. 2018.

Encoder Only: Bert

Input: Text sequence

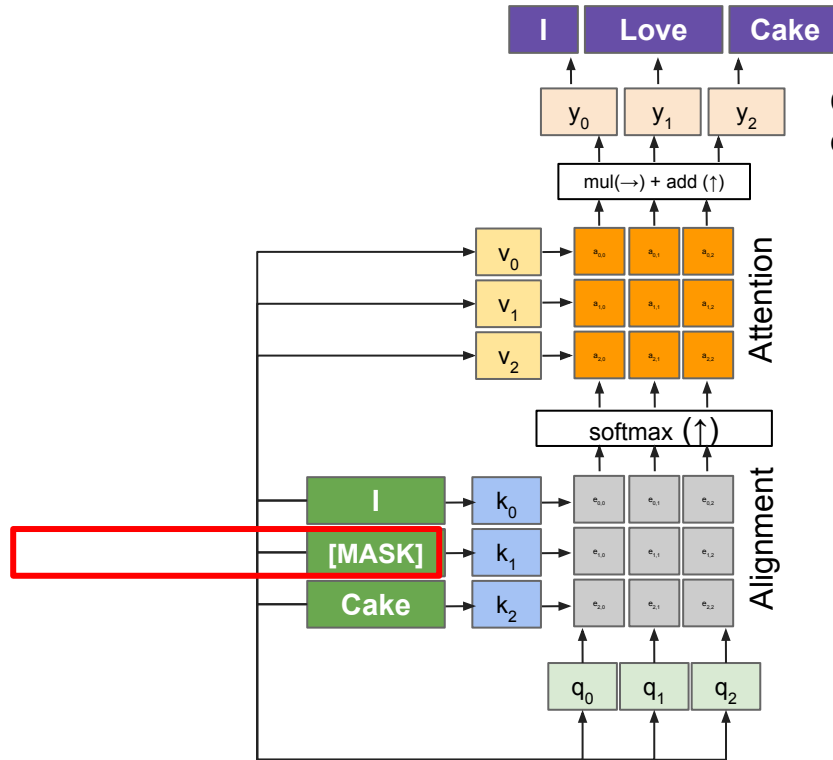
Output: Feature Vector

Randomly select 15% of tokens.

80% - [MASK]

10% - random token

10% - keep same



Outputs:
context vectors: \mathbf{y} (shape: D_v)


Devlin et al. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. 2018.

Encoder Only: Bert

Input: Text sequence

Output: Feature Vector

1 Masked Language Model



Randomly select 15% of tokens.

80% - [MASK]

10% - random token

10% - keep same

Devlin et al. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. 2018.

Encoder Only: Bert

Input: Text sequence

Output: Feature Vector

- 1 Masked Language Model
2. Next Sentence Prediction

Does sentence 2 follow sentence 1?

1 I love Cake [SEP] It's so good [SEP]

Transformer

[CLS] I [MASK] Cake [SEP] It's so good [SEP]

Randomly select 15% of tokens.

80% - [MASK]

10% - random token

10% - keep same

Devlin et al. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. 2018.


Encoder Only: Bert

Input: Text sequence

Output: Feature Vector

- 1 Masked Language Model
2. Next Sentence Prediction

Does sentence 2 follow sentence 1?



Randomly select 15% of tokens.

80% - [MASK]

10% - random token

10% - keep same

Devlin et al. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. 2018.

Encoder Only: Bert

Application to downstream tasks

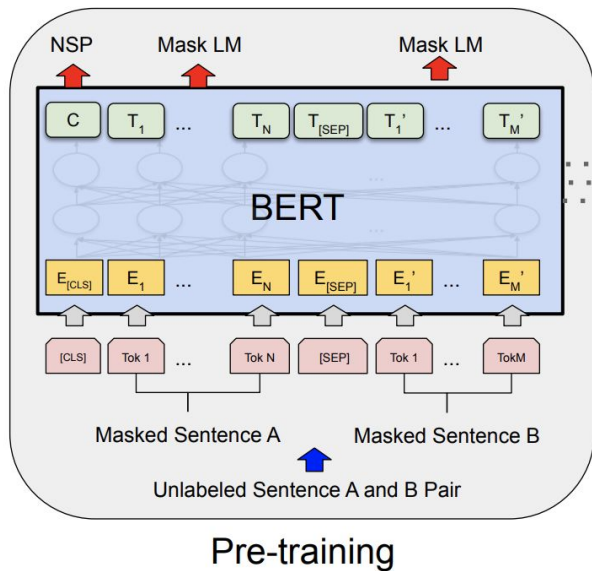


Image Source: [Devlin et al. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. 2018.](#)

Encoder Only: Bert

Application to downstream tasks





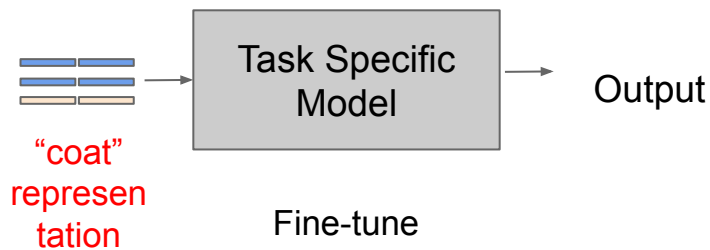
Image Source: [Devlin et al. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. 2018.](#)

ELMO: Two step process

Step 1: Pretrain a network on a pretext task that doesn't require supervision



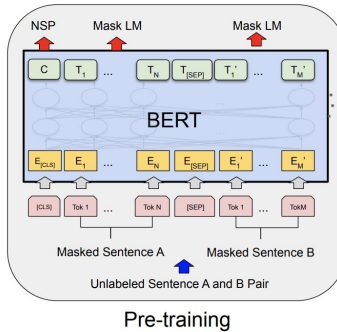
Step 2: Transfer encoder to downstream tasks via linear classifiers, KNN, finetuning



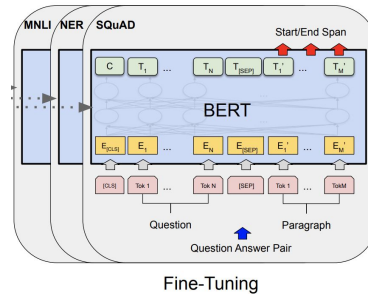
Downstream tasks:
Sentiment classification, NLI,
...

BERT: Two step process

Step 1: Pretrain a network on a pretext task that doesn't require supervision



Step 2: Transfer encoder to downstream tasks via linear classifiers, KNN, finetuning



Downstream tasks:
Sentiment
classification, NLI,
...

Image Source: [Devlin et al. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. 2018.](#)

BERT Results

System	Dev	Test
ESIM+GloVe	51.9	52.7
ESIM+ELMo	59.1	59.2
OpenAI GPT	-	78.0
BERT_{BASE}	81.6	-
BERT_{LARGE}	86.6	86.3
Human (expert) [†]	-	85.0
Human (5 annotations) [†]	-	88.0

Table 4: SWAG Dev and Test accuracies. [†]Human performance is measured with 100 samples, as reported in the SWAG paper.

System	Dev		Test	
	EM	F1	EM	F1
Top Leaderboard Systems (Dec. 10th, 2018)				
Human	-	-	82.3	91.2
#1 Ensemble - nlnet	-	-	86.0	91.7
#2 Ensemble - QANet	-	-	84.5	90.5
Published				
BiDAF+ELMo (Single)	-	85.6	-	85.8
R.M. Reader (Ensemble)	81.2	87.9	82.3	88.5
Ours				
BERT _{BASE} (Single)	80.8	88.5	-	-
BERT _{LARGE} (Single)	84.1	90.9	-	-
BERT _{LARGE} (Ensemble)	85.8	91.8	-	-
BERT _{LARGE} (Sgl.+TriviaQA)	84.2	91.1	85.1	91.8
BERT _{LARGE} (Ens.+TriviaQA)	86.2	92.2	87.4	93.2

Table 2: SQuAD 1.1 results. The BERT ensemble is 7x systems which use different pre-training checkpoints and fine-tuning seeds.

Image Source: [Devlin et al. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. 2018.](#)

LLMs

Encoder Only:

I love cake

ELMO: Bi-directional next word prediction,
BERT: Masked language objective, Next Sentence Prediction

Decoder Only:

I love

Encoder-Decoder:

I love cake

me gusta

LLMs

Encoder Only:

I love cake

ELMO: Bi-directional next word prediction,
BERT: Masked language objective, Next Sentence Prediction

Decoder Only:

I love

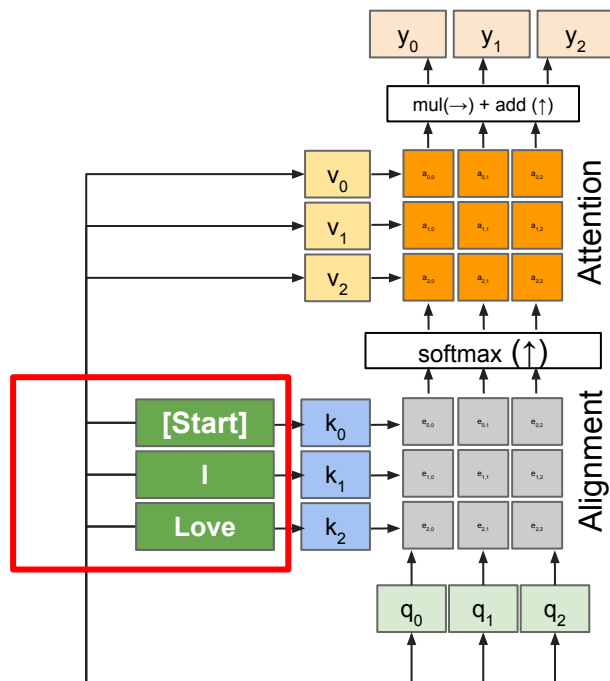
Encoder-Decoder:

I love cake

me gusta

Decoder Only: GPT

Input: Text sequence

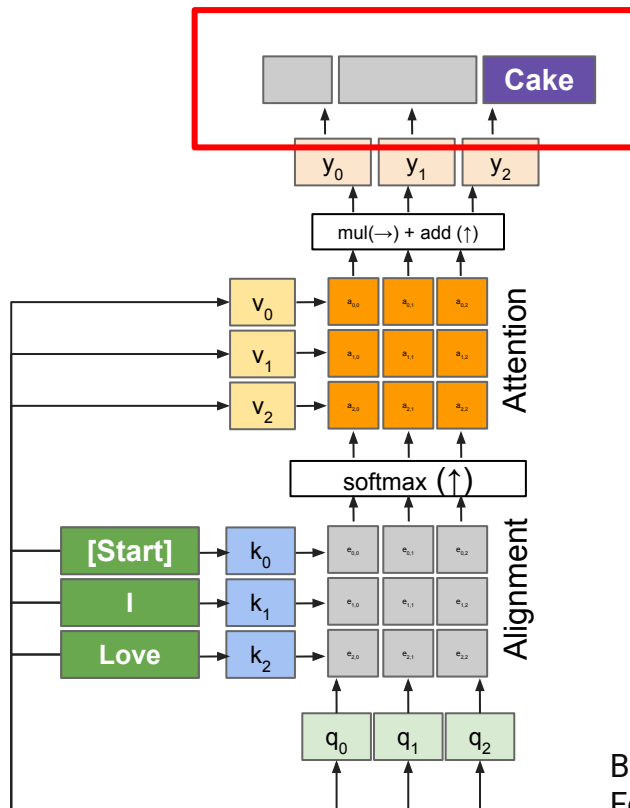


Brown et al. Language Models are Few-Shot Learners. 2020.

Decoder Only: GPT

Input: Text sequence

Output: Completed text sequence




Brown et al. Language Models are Few-Shot Learners. 2020.

Decoder Only: GPT

Input: Text sequence

Output: Completed text sequence

Cons: Need to process entire sentence in order to get loss from one word - not very much signal for the amount of processing



Brown et al. Language Models are Few-Shot Learners. 2020.


Decoder Only: GPT

Input: Text sequence

Output: Completed text sequence

Cons: Need to process entire sentence in order to get loss from one word - not very much signal for the amount of processing

Solution: predict each word given previous words so far



Brown et al. Language Models are Few-Shot Learners. 2020.

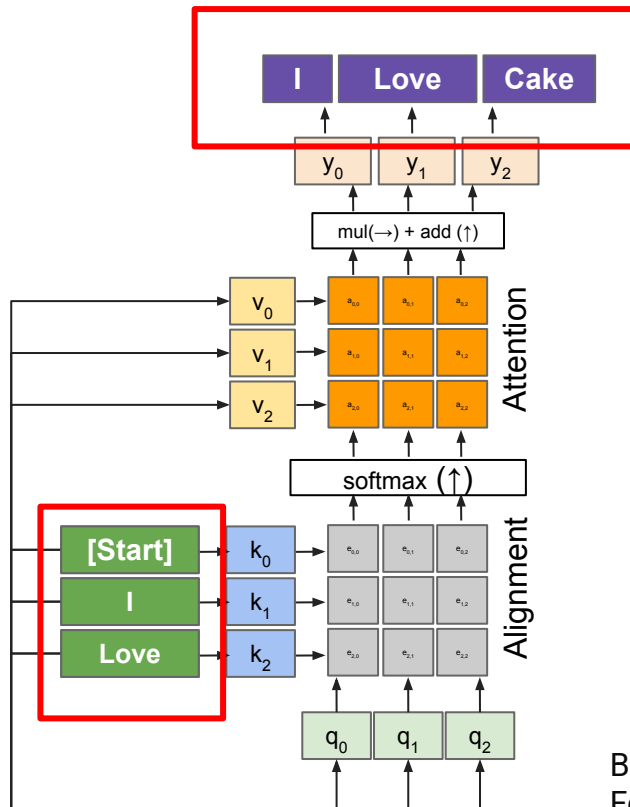
Decoder Only: GPT

Input: Text sequence

Output: Completed text sequence

Cons: Need to process entire sentence in order to get loss from one word - not very much signal for the amount of processing

Solution: predict each word given previous words so far




Brown et al. Language Models are Few-Shot Learners. 2020.

Decoder Only: GPT

Input: Text sequence

Output: Completed text sequence

What's wrong with this?



Brown et al. Language Models are Few-Shot Learners. 2020.


Decoder Only: GPT

Input: Text sequence

Output: Completed text sequence

What's wrong with this?

It can see the answer!



Brown et al. Language Models are Few-Shot Learners. 2020.

Decoder Only: GPT


Input: Text sequence

Output: Completed text sequence

What's wrong with this?

It can see the answer!

Solution: zero out values from future words



Brown et al. Language Models are Few-Shot Learners. 2020.

Decoder Only: GPT


Input: Text sequence

Output: Completed text sequence

What's wrong with this?

It can see the answer!

Solution: zero out values from future words




Brown et al. Language Models are Few-Shot Learners. 2020.

Decoder Only: GPT

Input: Text sequence

Output: Completed text sequence

To pre-train: predict next words from previous words for large text corpus



Brown et al. Language Models are Few-Shot Learners. 2020.

LLMs

Encoder Only:



ELMO: Bi-directional next word prediction,
BERT: Masked language objective, Next Sentence Prediction

Decoder Only:



GPT: next token prediction (autoregressive)

Encoder-Decoder:



LLMs

Encoder Only:



ELMO: Bi-directional next word prediction,
BERT: Masked language objective, Next Sentence Prediction

Decoder Only:




GPT: next token prediction (autoregressive)

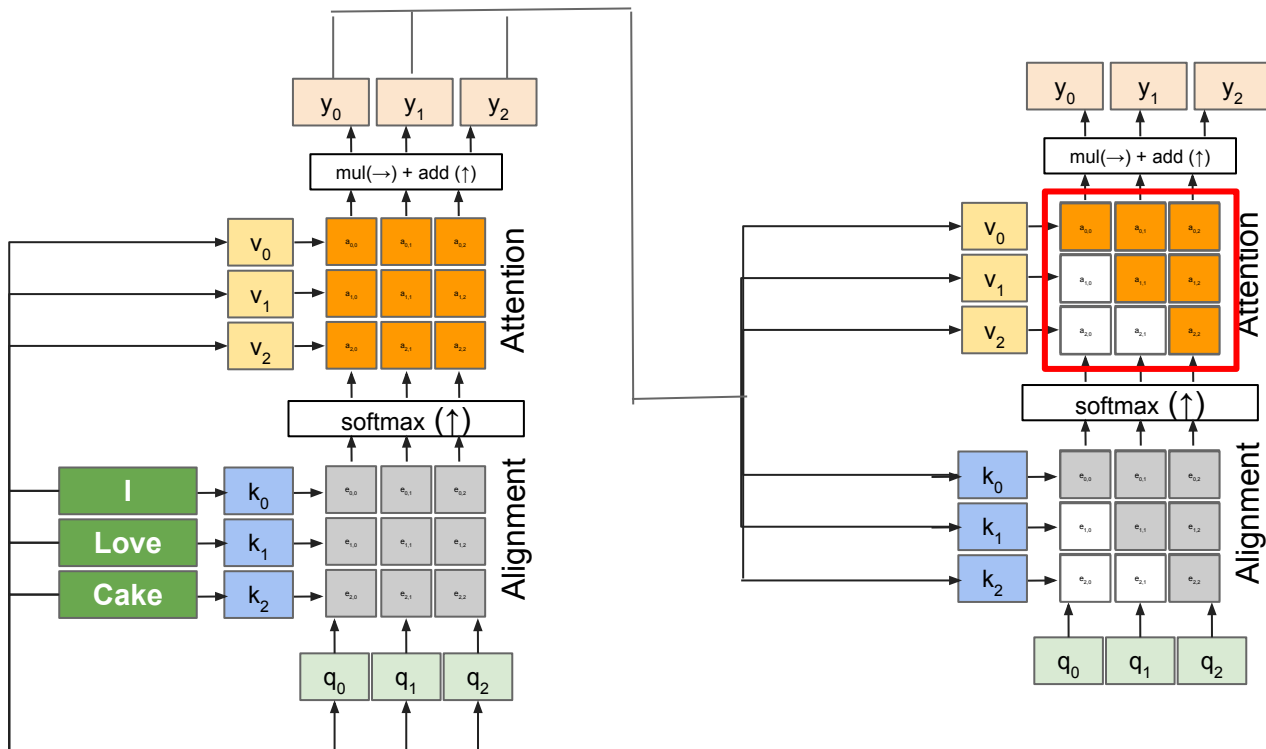
Encoder-Decoder:



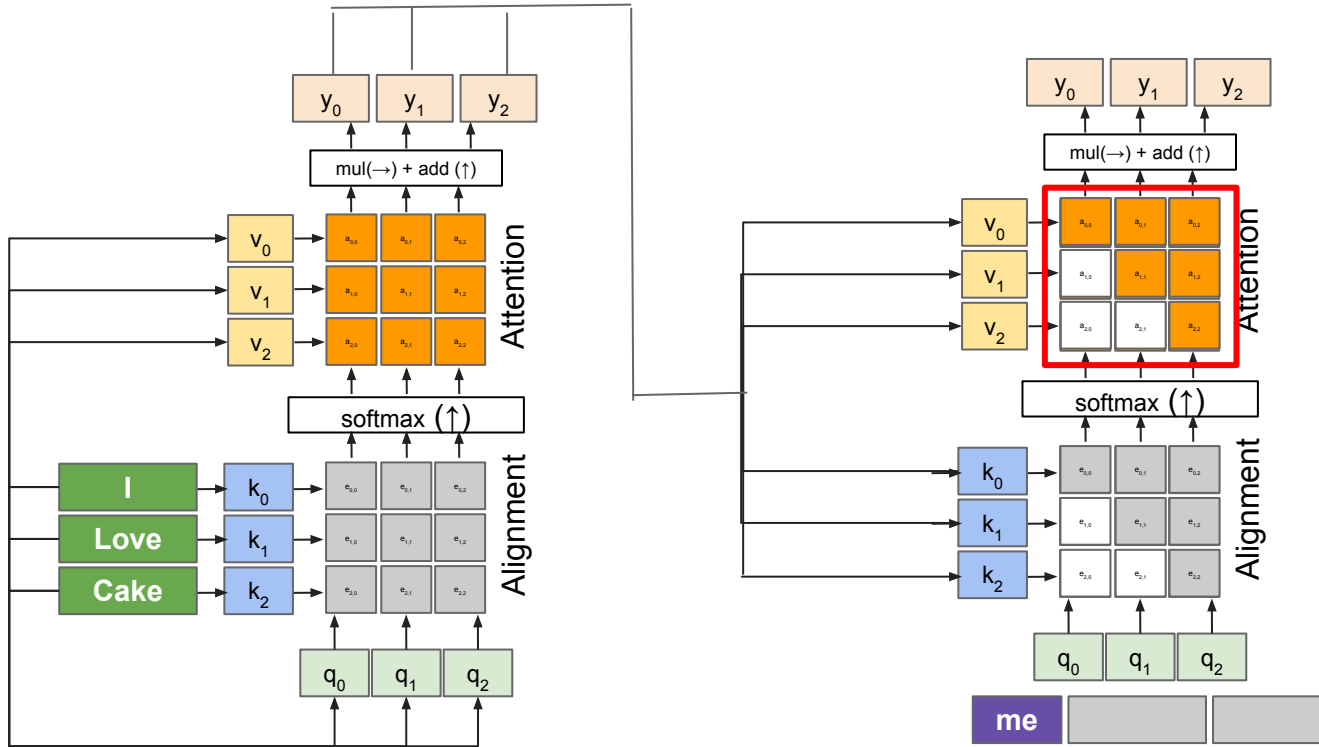
Encoder-Decoder: Generate text based on previously generated text and the meaning of a separate sequence



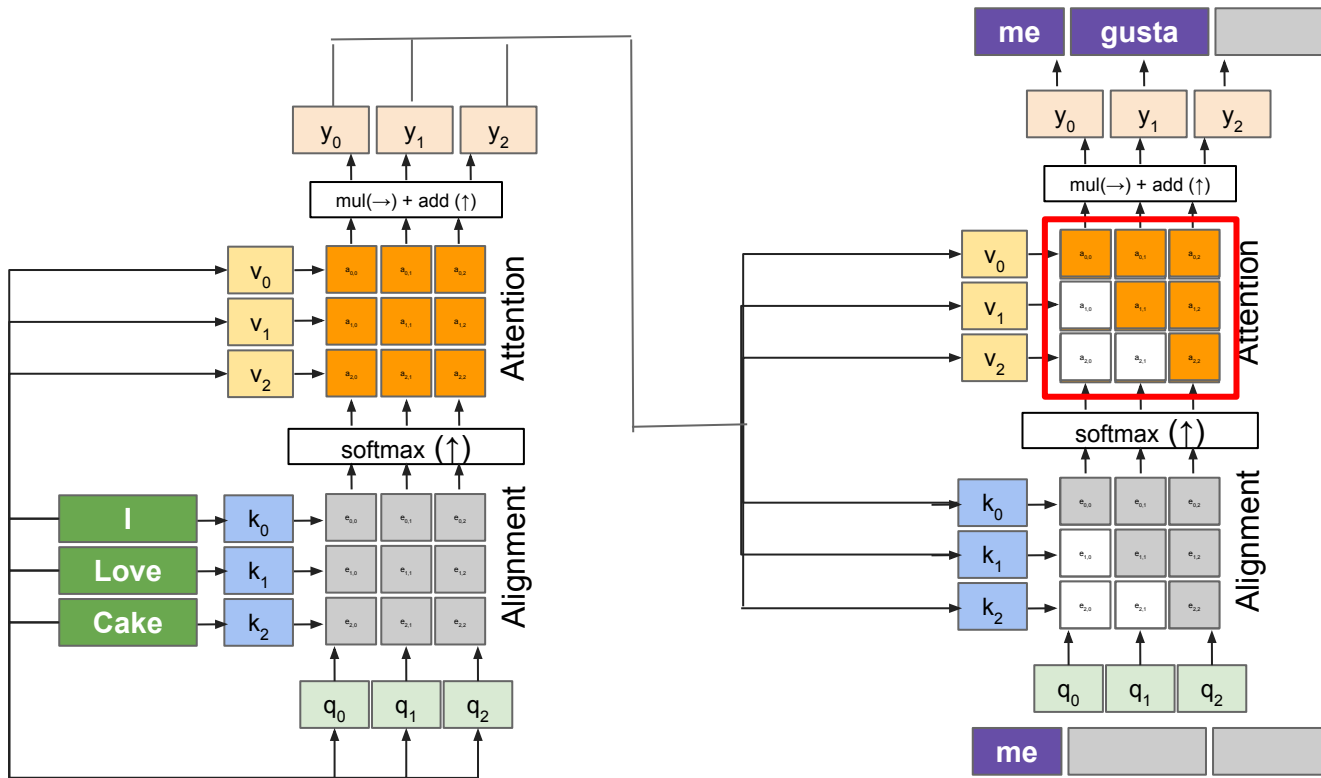
Encoder-Decoder: Generate text based on previously generated text and the meaning of a separate sequence




Encoder-Decoder: Generate text based on previously generated text and the meaning of a separate sequence



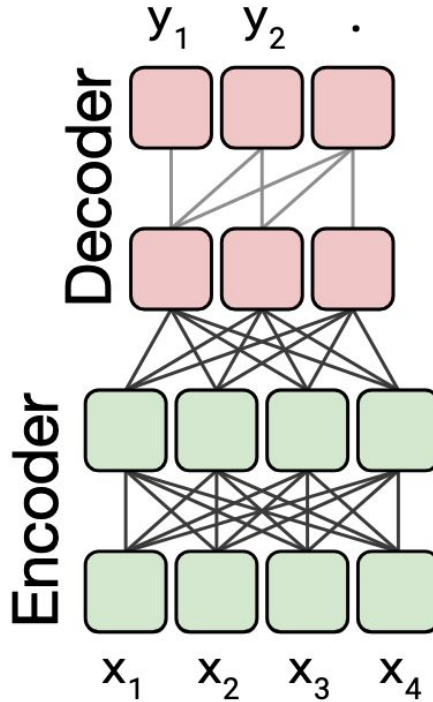
Encoder-Decoder: Generate text based on previously generated text and the meaning of a separate sequence



Encoder-Decoder: Generate text based on previously generated text and the meaning of a separate sequence



Encoder-Decoder: Generate text based on previously generated text and the meaning of a separate sequence



[Raffel et al., 2019](#)

Encoder-Decoder: Generate text based on previously generated text and the meaning of a separate sequence

Objective	Inputs	Targets
Prefix language modeling	Thank you for inviting	me to your party last week .
BERT-style Devlin et al. (2018)	Thank you <M> <M> me to your party apple week .	<i>(original text)</i>
Deshuffling	party me for your to . last fun you inviting week Thank	<i>(original text)</i>
MASS-style Song et al. (2019)	Thank you <M> <M> me to your party <M> week .	<i>(original text)</i>
I.i.d. noise, replace spans	Thank you <X> me to your party <Y> week .	<X> for inviting <Y> last <Z>
I.i.d. noise, drop tokens	Thank you me to your party week .	for inviting last
Random spans	Thank you <X> to <Y> week .	<X> for inviting me <Y> your party last <Z>

Image Source: Raffel et al. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. 2019

Encoder-Decoder: Generate text based on previously generated text and the meaning of a separate sequence

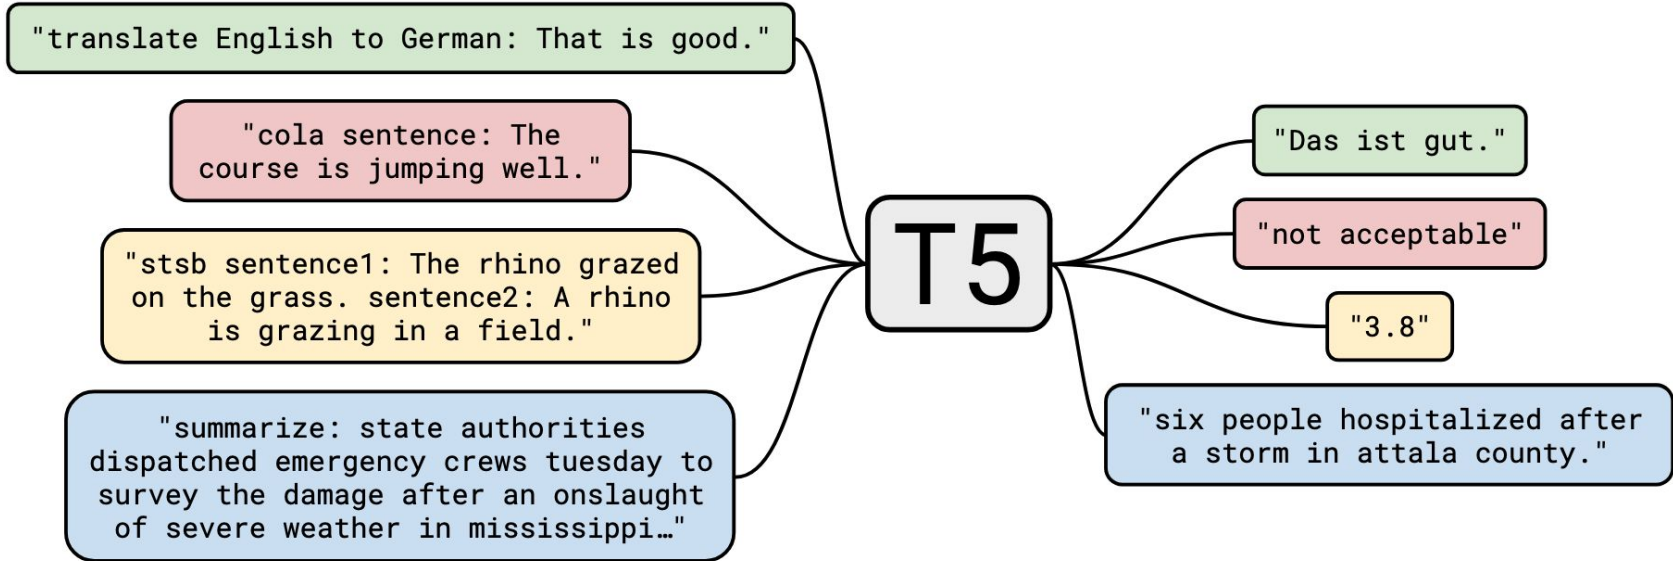


Image Source: Raffel et al. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. 2019

LLMs

Encoder Only:



ELMO: Bi-directional next word prediction,
BERT: Masked language objective, Next Sentence Prediction

Decoder Only:



GPT: text token prediction

Encoder-Decoder:



T5: Masked language objective



LLMs

Building LLMs: Pre-training objectives + architectures

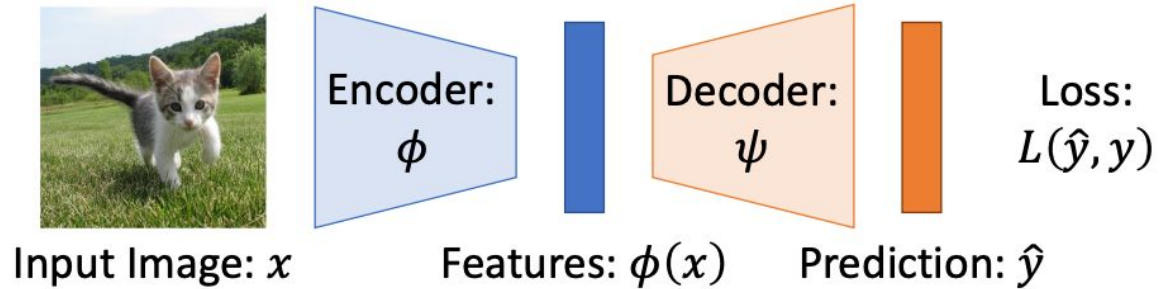
- Encoder only
- Decoder only
- Encoder Decoder

GPT


Gradient-Free Performance Improvement

Using pre-trained models out of the box

Step 1: Pretrain a network on a pretext task that doesn't require supervision




Step 2: Transfer encoder to downstream tasks via linear classifiers, KNN, finetuning



Using pre-trained models out of the box

Step 1: Pretrain a network on a pretext task that doesn't require supervision

reference frame




Source: [Vondrick et al., 2018](#)


Step 2: Transfer encoder to downstream tasks via linear classifiers, KNN, finetuning



Input Image: x



Features: $\phi(x)$





Downstream tasks:
Image classification,
object detection,
semantic segmentation

Using pre-trained models out of the box


Step 1: Pretrain a network on a pretext task that doesn't require supervision

reference frame



Source: [Vondrick et al., 2018](#)

Step 2: Use the model out of the box in a creative way!



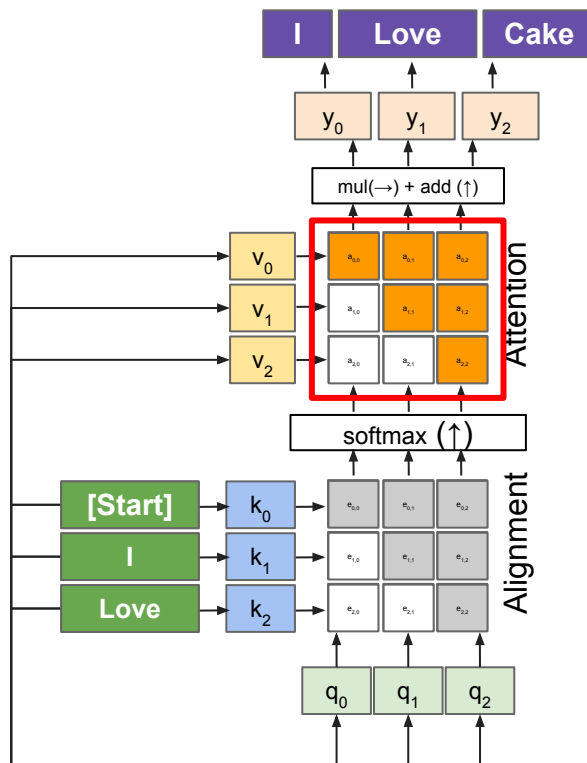
Source: [Google AI blog post](#)

Decoder Only: GPT

Input: Text sequence

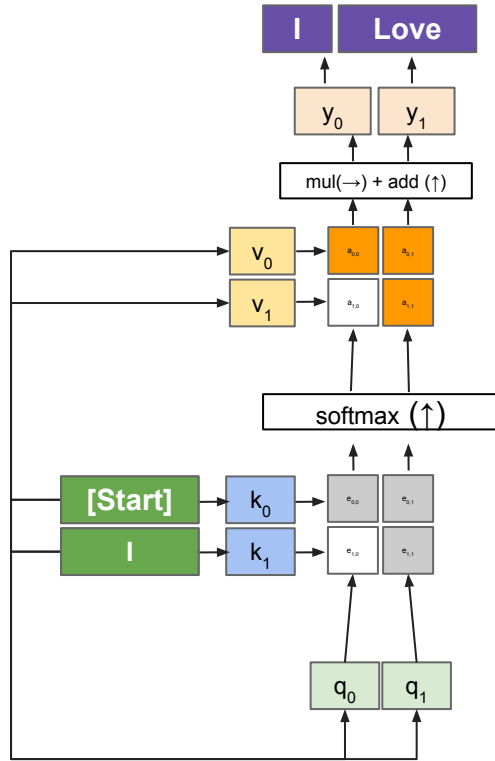
Output: Completed text sequence

To pre-train: predict next words from previous words for large text corpus

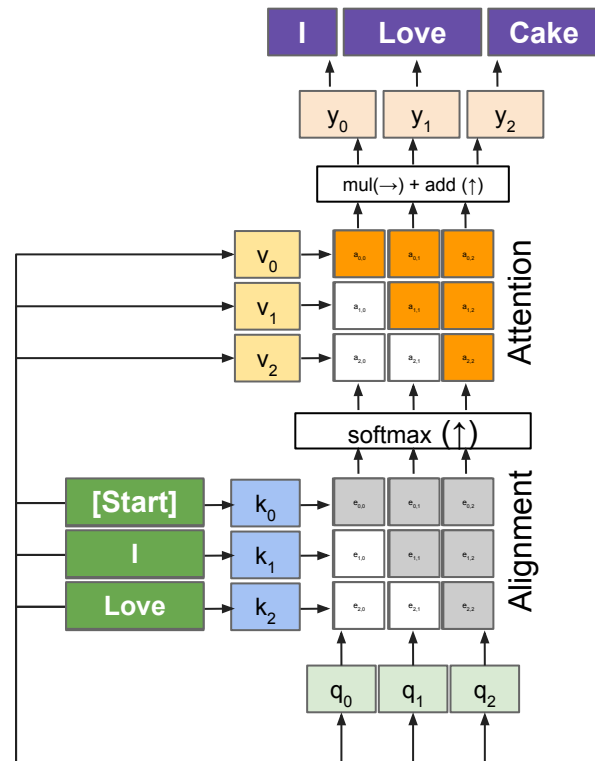
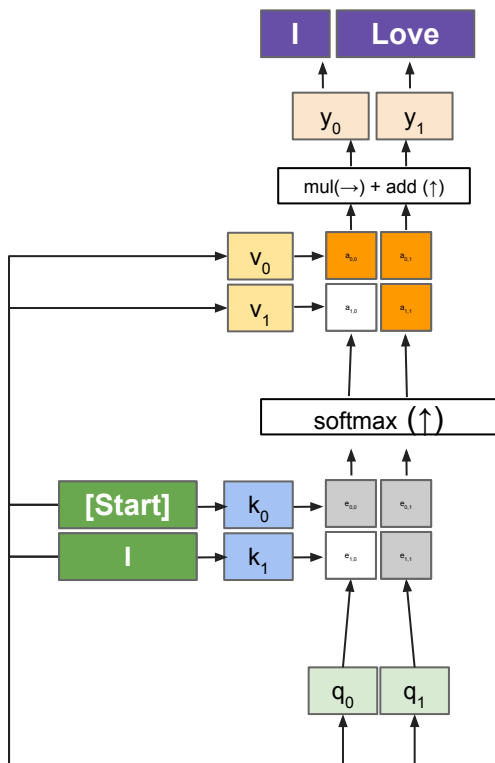


Decoder Only: Inference

Decoder Only: Inference




Decoder Only: Inference




Using pre-trained models out of the box

Step 1: Pretrain a network on a pretext task that doesn't require supervision




Step 2: Transfer encoder to downstream tasks via linear classifiers, KNN, finetuning




Using pre-trained models out of the box


Step 1: Pretrain a network on a pretext task that doesn't require supervision



Step 2: Use the model out of the box in a creative way!



Zero-shot



Fine-tuning is effective (but annoying)

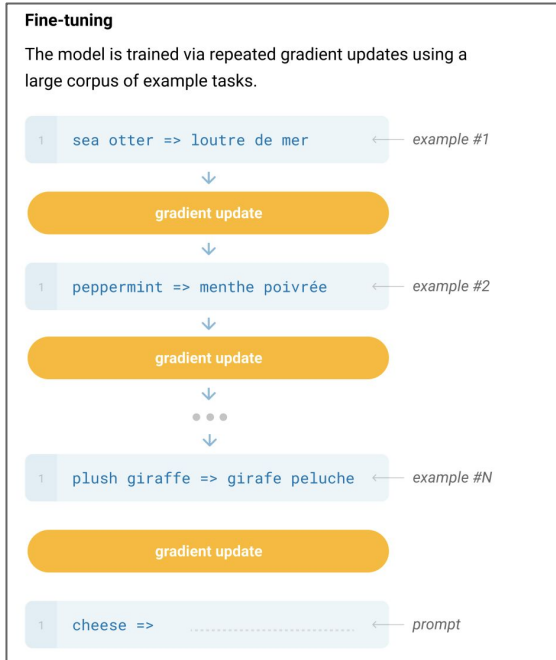
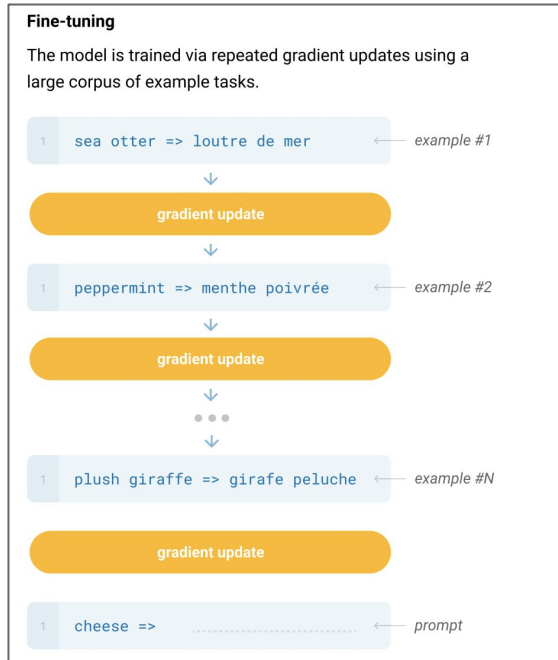


Image Source: [Language Models are Few-Shot Learners, Brown et al](#)

GPT-3 Results



Language Models are Few-Shot Learners

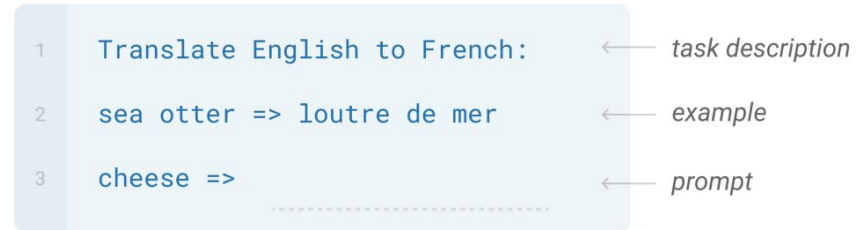


Image Source: [Language Models are Few-Shot Learners, Brown et al](#)

In-Context Learning

context

Translate English to French:
sea otter => loutre de mer
cheese =>

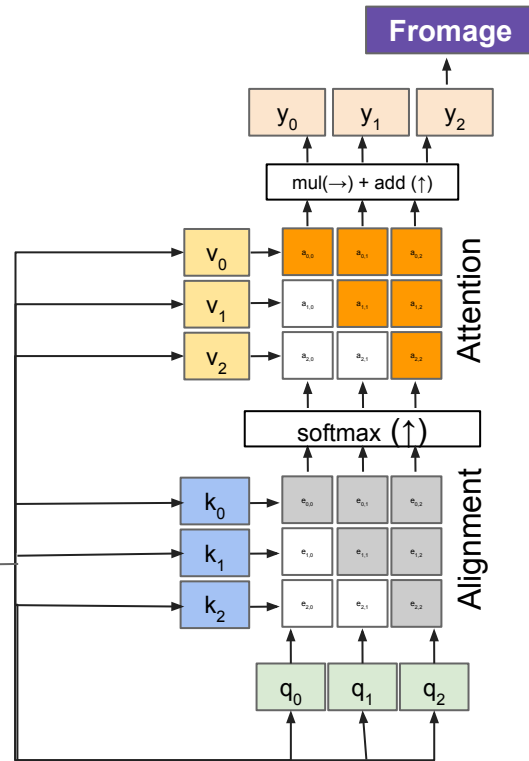


Image Source: [Language Models are Few-Shot Learners, Brown et al](#)

In-Context Learning

Zero-shot

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

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

One-shot

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

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

Few-shot

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

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

No training. No gradients.

Image Source: [Language Models are Few-Shot Learners, Brown et al](#)

Effect of In-Context Learning

s.u!c/c!e.s s i/o/n = succession





Image Source: [Language Models are Few-Shot Learners, Brown et al](#)

Effect of In-Context Learning

s.u!c/c!e.s s i/o/n = succession



Prompt becomes unnecessary

Image Source: [Language Models are Few-Shot Learners, Brown et al](#)

Effect of In-Context Learning

Example Question:

audacious is to boldness as

- (a) sanctimonious is to hypocrisy,
- (b) anonymous is to identity,
- (c) remorseful is to misdeed,
- (d) deleterious is to result,
- (e) impressionable is to temptation

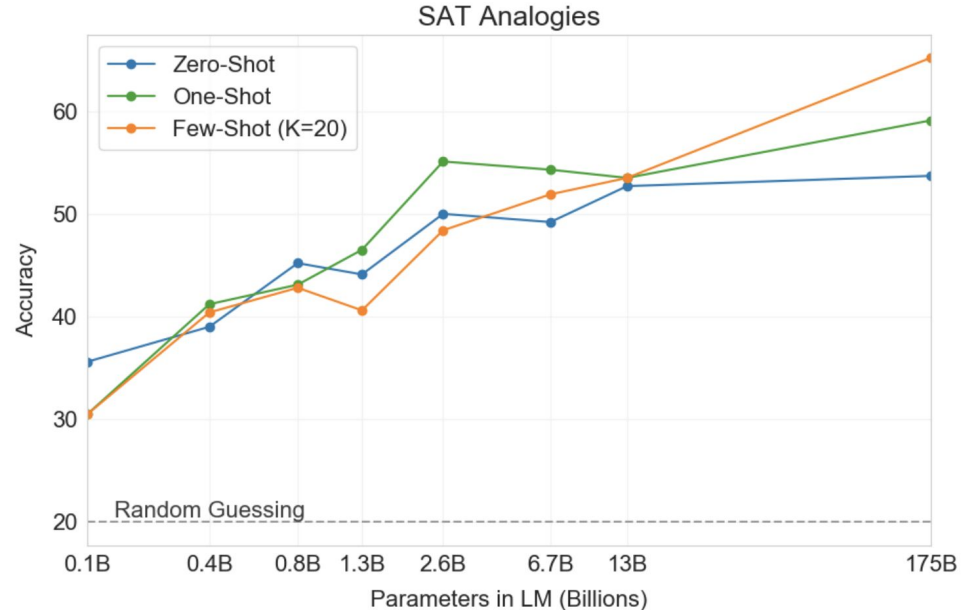


Image Source: [Language Models are Few-Shot Learners, Brown et al](#)

GPT Results




Figure 3.3: On TriviaQA GPT3’s performance grows smoothly with model size, suggesting that language models continue to absorb knowledge as their capacity increases. One-shot and few-shot performance make significant gains over zero-shot behavior, matching and exceeding the performance of the SOTA fine-tuned open-domain model, RAG [LPP+20]

Image Source: [Language Models are Few-Shot Learners, Brown et al](#)

GPT Results




Figure 3.6: GPT-3 results on PIQA in the zero-shot, one-shot, and few-shot settings. The largest model achieves a score on the development set in all three conditions that exceeds the best recorded score on the task.

Image Source: [Language Models are Few-Shot Learners, Brown et al](#)

GPT Results

	SuperGLUE Average	BoolQ Accuracy	CB Accuracy	CB F1	COPA Accuracy	RTE Accuracy
Fine-tuned SOTA	89.0	91.0	96.9	93.9	94.8	92.5
Fine-tuned BERT-Large	69.0	77.4	83.6	75.7	70.6	71.7
GPT-3 Few-Shot	71.8	76.4	75.6	52.0	92.0	69.0

	WiC Accuracy	WSC Accuracy	MultiRC Accuracy	MultiRC F1a	ReCoRD Accuracy	ReCoRD F1
Fine-tuned SOTA	76.1	93.8	62.3	88.2	92.5	93.3
Fine-tuned BERT-Large	69.6	64.6	24.1	70.0	71.3	72.0
GPT-3 Few-Shot	49.4	80.1	30.5	75.4	90.2	91.1

Table 3.8: Performance of GPT-3 on SuperGLUE compared to fine-tuned baselines and SOTA. All results are reported on the test set. GPT-3 few-shot is given a total of 32 examples within the context of each task and performs no gradient updates.

Image Source: [Language Models are Few-Shot Learners, Brown et al](#)

GPT Results


	SuperGLUE Average	BoolQ Accuracy	CB Accuracy	CB F1	COPA Accuracy	RTE Accuracy
Fine-tuned SOTA	89.0	91.0	96.9	93.9	94.8	92.5
Fine-tuned BERT-Large	69.0	77.4	83.6	75.7	70.6	71.7
GPT-3 Few-Shot	71.8	76.4	75.6	52.0	92.0	69.0

	WiC Accuracy	WSC Accuracy	MultiRC Accuracy	MultiRC F1a	ReCoRD Accuracy	ReCoRD F1
Fine-tuned SOTA	76.1	93.8	62.3	88.2	92.5	93.3
Fine-tuned BERT-Large	69.6	64.6	24.1	70.0	71.3	72.0
GPT-3 Few-Shot	49.4	80.1	30.5	75.4	90.2	91.1

Table 3.8: Performance of GPT-3 on SuperGLUE compared to fine-tuned baselines and SOTA. All results are reported on the test set. GPT-3 few-shot is given a total of 32 examples within the context of each task and performs no gradient updates.

All GPT was trained to do was predict the next token, how is it so good???


Scale!
(in params)



BERT (Large)
340 million params

GPT - 3
175 billion params

Scale! (In data)



BERT

3.3 Billion tokens¹

- **All of english wikipedia**
- **11,000 Books**

GPT - 3

~300 billion tokens

- **Common Crawl (Much of the internet)**

¹<https://aclanthology.org/W19-4828.pdf>

Loss vs Model and Dataset Size




Image source: [Kaplan et al. Scaling Laws for Neural Language Models. 2020.](#)

The Bitter Lesson

Rich Sutton

March 13, 2019

The biggest lesson that can be read from 70 years of AI research is that general methods that leverage computation are ultimately the most effective, and by a large margin. The ultimate reason for this is Moore's law, or rather its generalization of continued exponentially falling cost per unit of computation. Most AI research has been conducted as if the computation available to the agent were constant (in which case leveraging human knowledge would be one of the only ways to improve performance) but, over a slightly longer time than a typical research project, massively more computation inevitably becomes available. Seeking an improvement that makes a difference in the shorter term, researchers seek to leverage their human knowledge of the domain, but the only thing that matters in the long run is the leveraging of computation. These two need not run counter to each other, but in practice they tend to. Time spent on one is time not spent on the other. There are psychological commitments to investment in one approach or the other. And the human-knowledge approach tends to complicate methods in ways that make them less suited to taking advantage of general methods leveraging computation. There were many examples of AI researchers' belated learning of this bitter lesson, and it is instructive to review some of the most prominent.

<http://www.incompleteideas.net/Incldeas/BitterLesson.html>

LLMs

Building LLMs: Pre-training objectives + architectures

- Encoder only
- Decoder only
- Encoder Decoder

GPT

Gradient-Free Performance Improvement

Improving accuracy without fine-tuning

In-context learning is better than zero-shot with no additional training

Can we do better than in context learning with no additional training?

Chain of thought prompting

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: John takes care of 10 dogs. Each dog takes .5 hours a day to walk and take care of their business. How many hours a week does he spend taking care of dogs?

A:

GPT (or similar)

The answer is 50. ❌

	GSM8K
Finetuned GPT-3 175B	33%
Finetuned GPT-3 175B + verifier (prior SOTA)	55%
9–12 year olds (Cobbe et al., 2021)	60%
PaLM 540B: standard prompting	17.9%

Image Source: [Wei et al. Chain of Thought Prompting Elicits Reasoning in Large Language Models. 2022.](#)

Chain of thought prompting

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: John takes care of 10 dogs. Each dog takes .5 hours a day to walk and take care of their business. How many hours a week does he spend taking care of dogs?

A:

GPT (or similar)

	GSM8K
Finetuned GPT-3 175B	33%
Finetuned GPT-3 175B + verifier (prior SOTA)	55%
9–12 year olds (Cobbe et al., 2021)	60%
PaLM 540B: standard prompting	17.9%

Image Source: [Wei et al. Chain of Thought Prompting Elicits Reasoning in Large Language Models. 2022.](#)

Chain of thought prompting

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: John takes care of 10 dogs. Each dog takes .5 hours a day to walk and take care of their business. How many hours a week does he spend taking care of dogs?

A:

GPT (or similar)

John takes care of 10 dogs. Each dog takes .5 hours a day to walk and take care of their business. So that is $10 \times .5 = 5$ hours a day. $5 \text{ hours a day} \times 7 \text{ days a week} = 35 \text{ hours a week}$. The answer is 35 hours a week. ✓

	GSM8K
Finetuned GPT-3 175B	33%
Finetuned GPT-3 175B + verifier (prior SOTA) 9–12 year olds (Cobbe et al., 2021)	55% 60%
PaLM 540B: standard prompting	17.9%
PaLM 540B: chain of thought prompting	58.1%

Image Source: [Wei et al. Chain of Thought Prompting Elicits Reasoning in Large Language Models. 2022.](#)

Chain of thought for Math Problems

Table 20: Few-shot exemplars for full chain of thought prompt for math word problems. This set of exemplars was used for all math word problem datasets except AQuA.

PROMPT FOR MATH WORD PROBLEMS

Q: There are 15 trees in the grove. Grove workers will plant trees in the grove today. After they are done, there will be 21 trees. How many trees did the grove workers plant today?

A: There are 15 trees originally. Then there were 21 trees after some more were planted. So there must have been $21 - 15 = 6$. The answer is 6.

Q: If there are 3 cars in the parking lot and 2 more cars arrive, how many cars are in the parking lot?

A: There are originally 3 cars. 2 more cars arrive. $3 + 2 = 5$. The answer is 5.

Q: Leah had 32 chocolates and her sister had 42. If they ate 35, how many pieces do they have left in total?

A: Originally, Leah had 32 chocolates. Her sister had 42. So in total they had $32 + 42 = 74$. After eating 35, they had $74 - 35 = 39$. The answer is 39.

Q: Jason had 20 lollipops. He gave Denny some lollipops. Now Jason has 12 lollipops. How many lollipops did Jason give to Denny?

A: Jason started with 20 lollipops. Then he had 12 after giving some to Denny. So he gave Denny $20 - 12 = 8$. The answer is 8.

Image Source: [Wei et al. Chain of Thought Prompting Elicits Reasoning in Large Language Models. 2022.](#)

Chain of thought for Math Problems

Table 8: Example of correct chains of thought produced by the model for the GSM8K dataset.

QUESTION: James decides to run 3 sprints 3 times a week. He runs 60 meters each sprint. How many total meters does he run a week?

MODEL ANSWER (CORRECT): James decides to run 3 sprints 3 times a week. He runs 60 meters each sprint. So he runs 60 meters x 3 sprints x 3 times a week. That is 60 meters x 9. The answer is 540. ✓

QUESTION: Brandon's iPhone is four times as old as Ben's iPhone. Ben's iPhone is two times older than Suzy's iPhone. If Suzy's iPhone is 1 year old, how old is Brandon's iPhone?

MODEL ANSWER (CORRECT): Brandon's iPhone is 4 times as old as Ben's iPhone. Ben's iPhone is 2 times older than Suzy's iPhone. So Brandon's iPhone is $4 \times 2 = 8$ times older than Suzy's iPhone. Suzy's iPhone is 1 year old. So Brandon's iPhone is $8 \times 1 = 8$ years old. The answer is 8. ✓

Image Source: [Wei et al. Chain of Thought Prompting Elicits Reasoning in Large Language Models. 2022.](#)

Chain of thought for Symbolic Reasoning

PROMPT FOR LAST LETTER CONCATENATION

Q: Take the last letters of the words in "Elon Musk" and concatenate them.

A: The last letter of "Elon" is "n". The last letter of "Musk" is "k". Concatenating them is "nk". The answer is nk.

Q: Take the last letters of the words in "Larry Page" and concatenate them.

A: The last letter of "Larry" is "y". The last letter of "Page" is "e". Concatenating them is "ye". The answer is ye.

Q: Take the last letters of the words in "Sergey Brin" and concatenate them.

A: The last letter of "Sergey" is "y". The last letter of "Brin" is "n". Concatenating them is "yn". The answer is yn.

Q: Take the last letters of the words in "Bill Gates" and concatenate them.

A: The last letter of "Bill" is "l". The last letter of "Gates" is "s". Concatenating them is "ls". The answer is ls.

Image Source: [Wei et al. Chain of Thought Prompting Elicits Reasoning in Large Language Models. 2022.](#)

Chain of thought for Symbolic Reasoning

Table 13: Examples of correct and incorrect chains of thought produced by LaMDA 137B on the letter concatenation task.

QUESTION: Take the last letters of the words in “Waldo Schmidt” and concatenate them.

MODEL ANSWER (CORRECT): The last letter of “Waldo” is “o”. The last letter of “Schmidt” is “t”. Concatenating them is “ot”. So the answer is ot. ✓

Image Source: [Wei et al. Chain of Thought Prompting Elicits Reasoning in Large Language Models. 2022.](#)

Chain of thought for Physical Reasoning

PROMPT FOR COIN FLIP

Q: A coin is heads up. Ka flips the coin. Sherrie flips the coin. Is the coin still heads up?

A: The coin was flipped by Ka and Sherrie. So the coin was flipped 2 times, which is an even number. The coin started heads up, so after an even number of flips, it will still be heads up. So the answer is yes.

Q: A coin is heads up. Jamey flips the coin. Teressa flips the coin. Is the coin still heads up?

A: The coin was flipped by Jamey and Teressa. So the coin was flipped 2 times, which is an even number. The coin started heads up, so after an even number of flips, it will still be heads up. So the answer is yes.

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.

Image Source: [Wei et al. Chain of Thought Prompting Elicits Reasoning in Large Language Models. 2022.](#)

Chain of thought for Physical Reasoning

QUESTION: A coin is heads up. Andree flips the coin. Audrie does not flip the coin. Is the coin still heads up?

MODEL ANSWER (CORRECT): The coin was flipped by Andree. 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.



Image Source: [Wei et al. Chain of Thought Prompting Elicits Reasoning in Large Language Models. 2022.](#)


Chain of thought results

	GSM8K	SVAMP	ASDiv	MAWPS
Standard prompting	6.5 \pm 0.4	29.5 \pm 0.6	40.1 \pm 0.6	43.2 \pm 0.9
Chain of thought prompting	14.3 \pm 0.4	36.7 \pm 0.4	46.6 \pm 0.7	57.9 \pm 1.5

	Commonsense			Symbolic	
	Date	Sports	SayCan	Concat	Coin
Standard prompting	21.5 \pm 0.6	59.5 \pm 3.0	80.8 \pm 1.8	5.8 \pm 0.6	49.0 \pm 2.1
Chain of thought prompting	26.8 \pm 2.1	85.8 \pm 1.8	91.7 \pm 1.4	77.5 \pm 3.8	99.6 \pm 0.3

Image Source: [Wei et al. Chain of Thought Prompting Elicits Reasoning in Large Language Models. 2022.](#)

In-context CoT traces need not be accurate (!)



[Min et al., 2022](#)

Explicit CoT is usually not required today

Models are **fine-tuned** on CoT traces, so they often do CoT on their own

However, few-shot examples are still useful

b/c CoT elicits a *behavior*, while few-shot provides useful information

Think step by step

(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.* ✗

(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* ✗

(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.* ✓

(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.* ✓

Image Source: [Chowdhery et al. Large Language Models are Zero-Shot Reasoners. 2022.](#)

Think step by step

	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

Image Source: [Chowdhery et al. Large Language Models are Zero-Shot Reasoners. 2022.](#)

Self Consistency

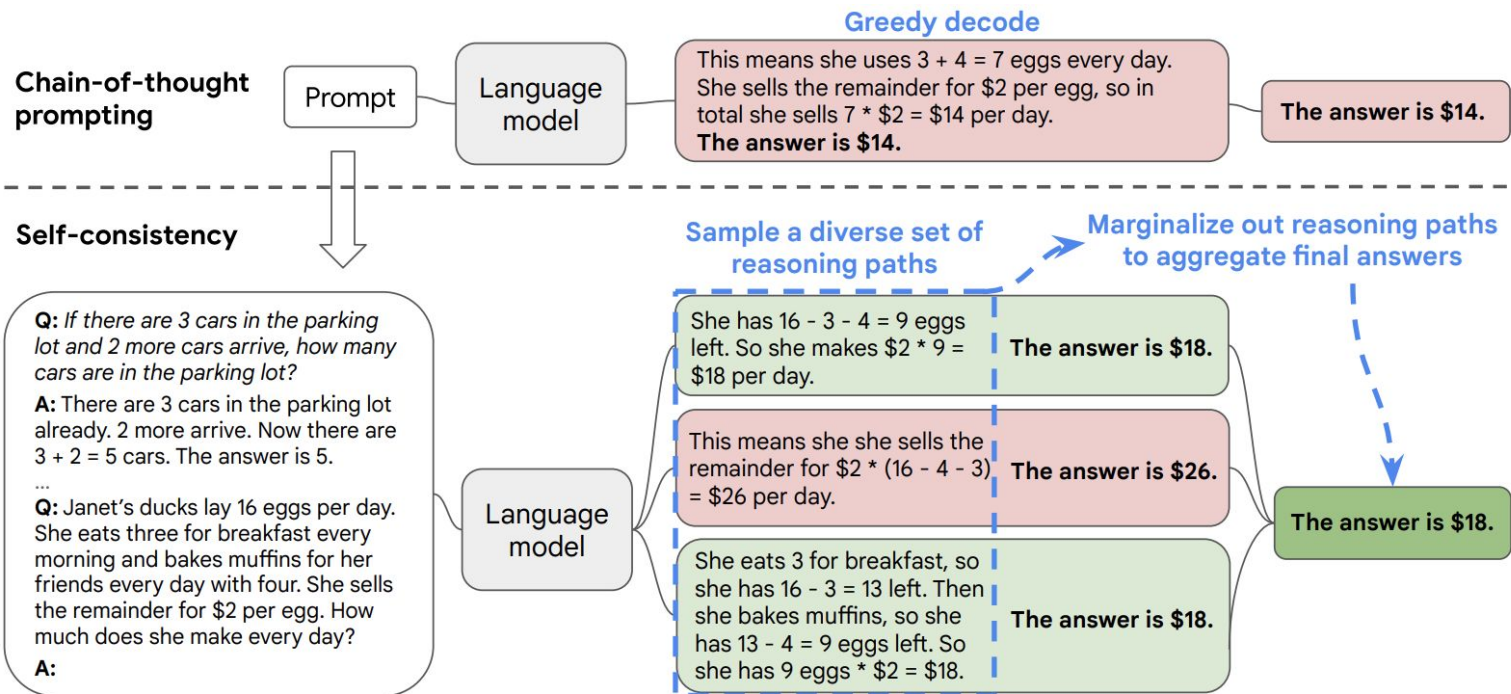


Image Source: [Xie et al. Self-Consistency Improves Chain of Thought Reasoning in Language Models. 2022.](#)

Self Consistency

	GSM8K	MultiArith	AQuA	SVAMP	CSQA	ARC-c
Greedy decode	56.5	94.7	35.8	79.0	79.0	85.2
Weighted avg (unnormalized)	56.3 \pm 0.0	90.5 \pm 0.0	35.8 \pm 0.0	73.0 \pm 0.0	74.8 \pm 0.0	82.3 \pm 0.0
Weighted avg (normalized)	22.1 \pm 0.0	59.7 \pm 0.0	15.7 \pm 0.0	40.5 \pm 0.0	52.1 \pm 0.0	51.7 \pm 0.0
Weighted sum (unnormalized)	59.9 \pm 0.0	92.2 \pm 0.0	38.2 \pm 0.0	76.2 \pm 0.0	76.2 \pm 0.0	83.5 \pm 0.0
Weighted sum (normalized)	74.1 \pm 0.0	99.3 \pm 0.0	48.0 \pm 0.0	86.8 \pm 0.0	80.7 \pm 0.0	88.7 \pm 0.0
Unweighted sum (majority vote)	74.4 \pm 0.1	99.3 \pm 0.0	48.3 \pm 0.5	86.6 \pm 0.1	80.7 \pm 0.1	88.7 \pm 0.1

Table 1: Accuracy comparison of different answer aggregation strategies on PaLM-540B.

Image Source: [Xie et al. Self-Consistency Improves Chain of Thought Reasoning in Language Models. 2022.](#)

Summary Slide

Encoder Only: Capture the meaning of an entire sequence

I love cake

ELMO: Bi-directional next word prediction,
BERT: Masked language objective, Next Sentence Prediction

Decoder Only: Generate text based on previously generated text

I love

GPT: next token prediction (autoregressive)

Encoder-Decoder: Generate text based on previously generated text and the meaning of a separate sequence **T5**: Masked language objective

I love cake

me gusta

Summary

Zero-shot

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

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

One-shot

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

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

Few-shot

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

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

Image Source: [Language Models are Few-Shot Learners, Brown et al](#)

Summary




Image source: [Kaplan et al. Scaling Laws for Neural Language Models. 2020.](#)

Summary

Inference Only Performance Improvement

- Chain-of-thought
- Think step by step
- Self consistency

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.

Image Source: [Wei et al. Chain of Thought Prompting Elicits Reasoning in Large Language Models. 2022.](#)



<https://twitter.com/alirahimi0>