

Ralf Möller, Sylvia Melzer

1d-CNNs LSTMs ELMo Transformers BERT GPT

Acknowledgements

- Some slides are based on
 - Machine Learning in NLP (Spring 2020)
 - <http://courses.engr.illinois.edu/cs546/>
 - Julia Hockenmaier <http://juliahr.cs.illinois.edu>
 - RNs, LSTMs, ELMo, Transformers
 - Machine Learning (Spring 2020)
 - http://speech.ee.ntu.edu.tw/~tlkagk/courses_ML20.html
 - 李宏毅 (Hung-yi Lee) <http://speech.ee.ntu.edu.tw/~tlkagk/>
 - ELMo, BERT: [http://speech.ee.ntu.edu.tw/~tlkagk/courses/ML_2019/Lecture/BERT%20\(v3\).pdf](http://speech.ee.ntu.edu.tw/~tlkagk/courses/ML_2019/Lecture/BERT%20(v3).pdf)
- Respective sources are indicated in the gray line at the bottom
- Slides have been modified (All errors are mine)

Convolution

Input image



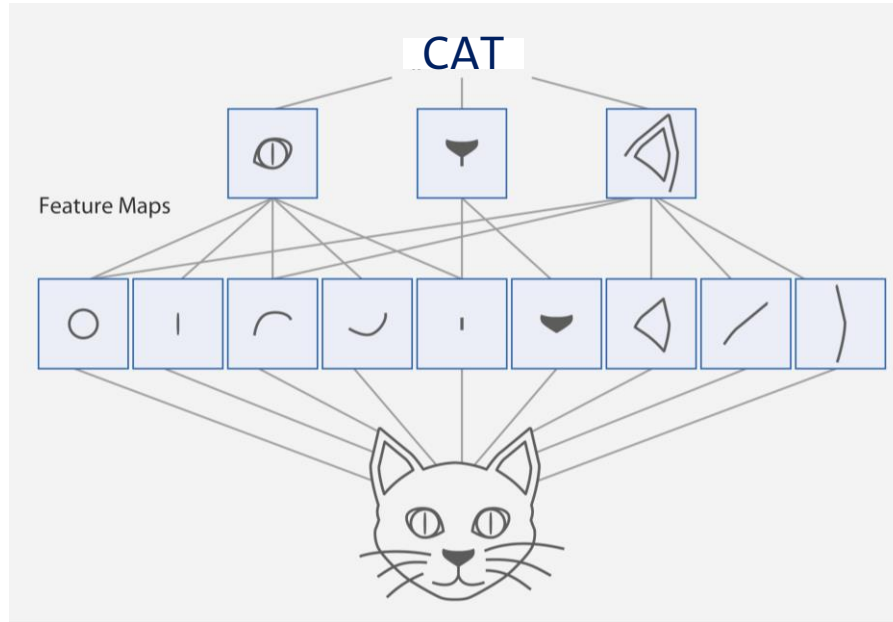
Convolution
Kernel

$$\begin{bmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{bmatrix}$$

Feature map




Feature Maps



<https://www.heise.de/ratgeber/Neuronale-Netz-einfach-erklart-6343697.html?seite=6>

Convolutional Networks (CNs)



Main ConvNet idea for text:

Compute vectors for n-grams and group them afterwards

Example: “this takes too long” compute vectors for:

This takes, takes too, too long, this takes too, takes too long, this takes too long

ConvNets (CNs)

Feature Map

6	4	8	5
5	4	5	8
3	6	7	7
7	9	7	2

Max-Pooling

Main ConvNet idea for text:

Compute vectors for n-grams and group them afterwards

1d-CNNs for text

- Text is a (variable-length) sequence of words (word vectors)
- We can use a 1d-CNN to slide a window of n tokens across:
 - filter size $n = 3$, stride = 1, no padding

The	quick	brown	fox	jumps	over	the	lazy	dog
The	quick	brown	fox	jumps	over	the	lazy	dog
The	quick	brown	fox	jumps	over	the	lazy	dog
The	quick	brown	fox	jumps	over	the	lazy	dog
The	quick	brown	fox	jumps	over	the	lazy	dog
The	quick	brown	fox	jumps	over	the	lazy	dog

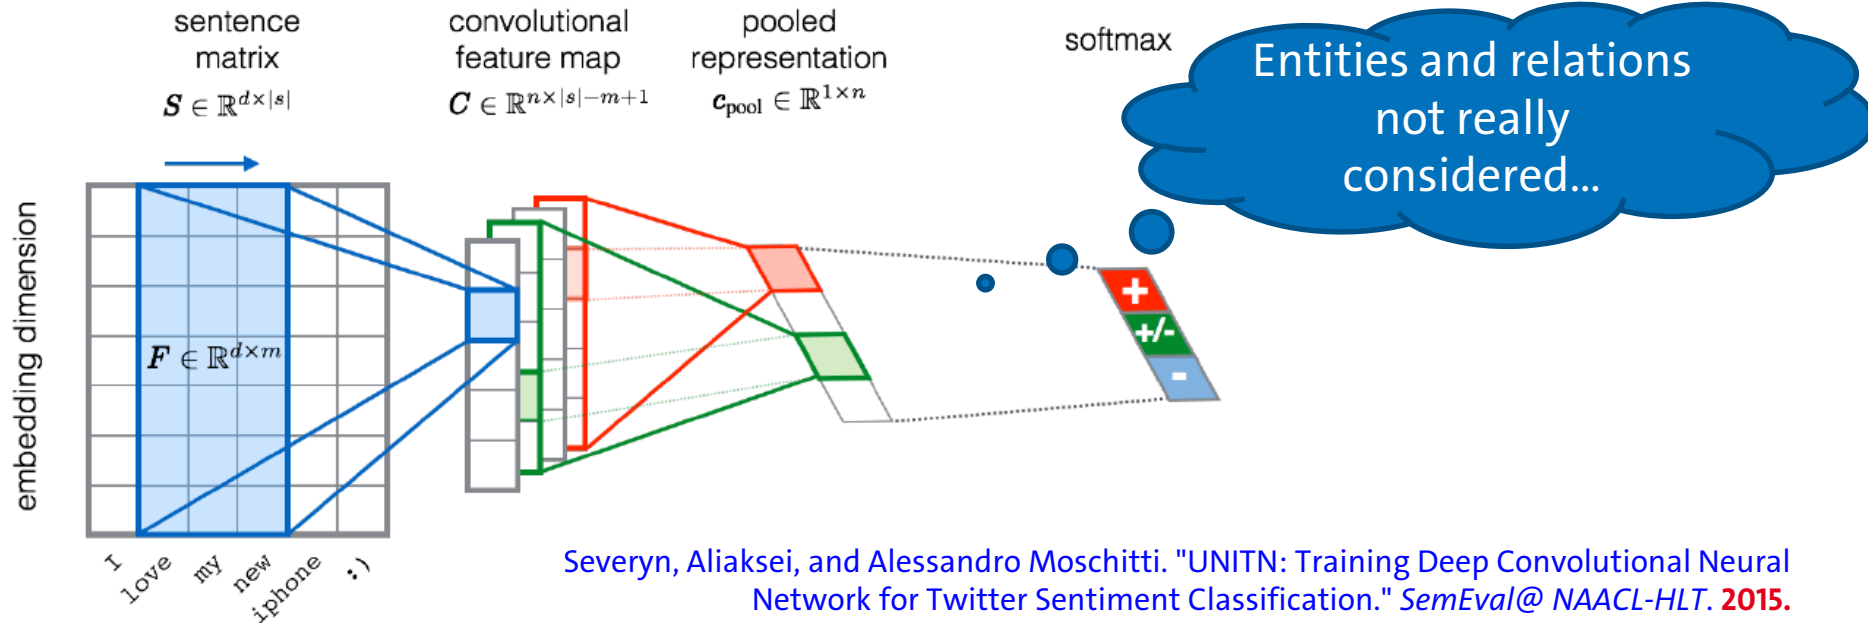
1d-CNNs for text

- filter size $n = 2$, stride = 2, no padding

The	quick	brown	fox	jumps	over	the	lazy	dog
The	quick	brown	fox	jumps	over	the	lazy	dog
The	quick	brown	fox	jumps	over	the	lazy	dog
The	quick	brown	fox	jumps	over	the	lazy	dog

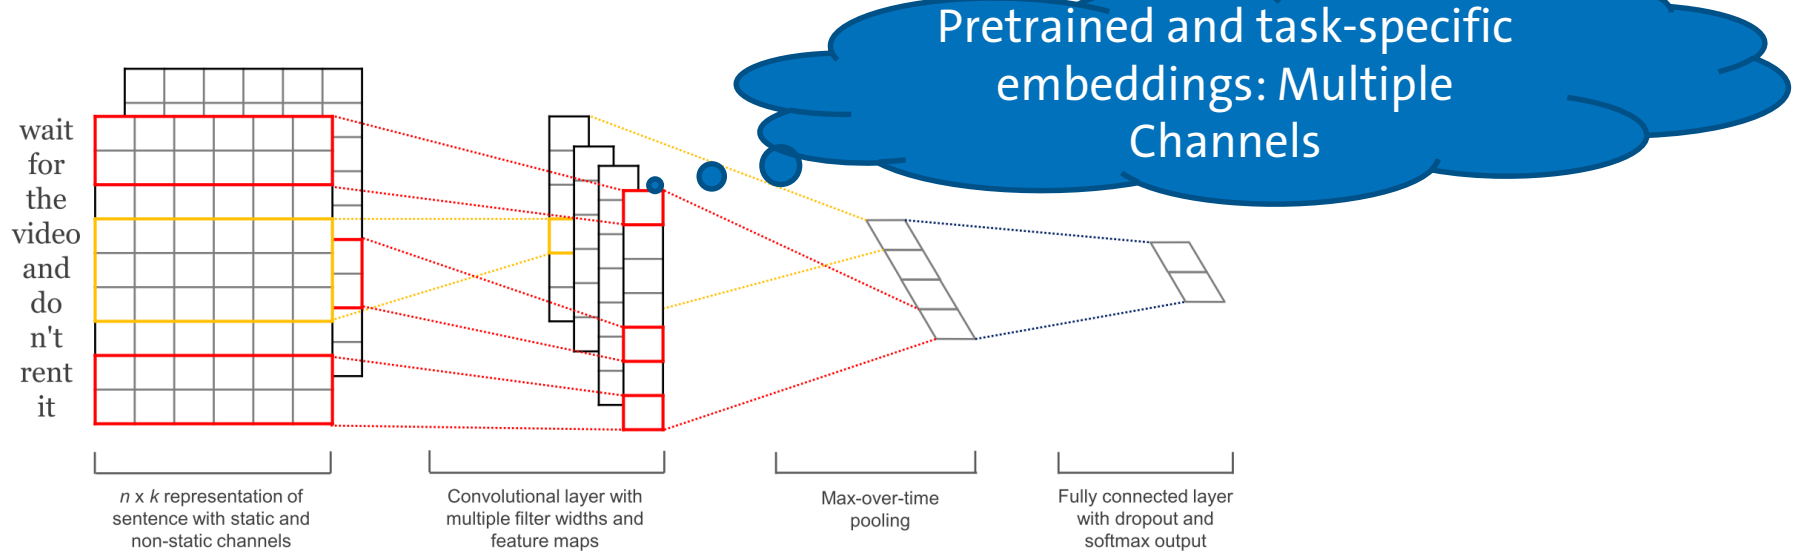
- CNNs (w/ ReLU and maxpool) can be used for classifying (parts of) the text

CNNs for sentiment analysis



Severyn, Aliaksei, and Alessandro Moschitti. "UNITN: Training Deep Convolutional Neural Network for Twitter Sentiment Classification." *SemEval@ NAACL-HLT*. 2015.

CNNs for sentence/text classification




sliding over 3, 4 or 5 words at a time

Kim, Y. "Convolutional Neural Networks for Sentence Classification", EMNLP (2014)


Recurrent Networks (RN) – Or: Copying the Pattern

- Basic computational network copied per time slice
- Input: previous hidden state, output: next hidden state

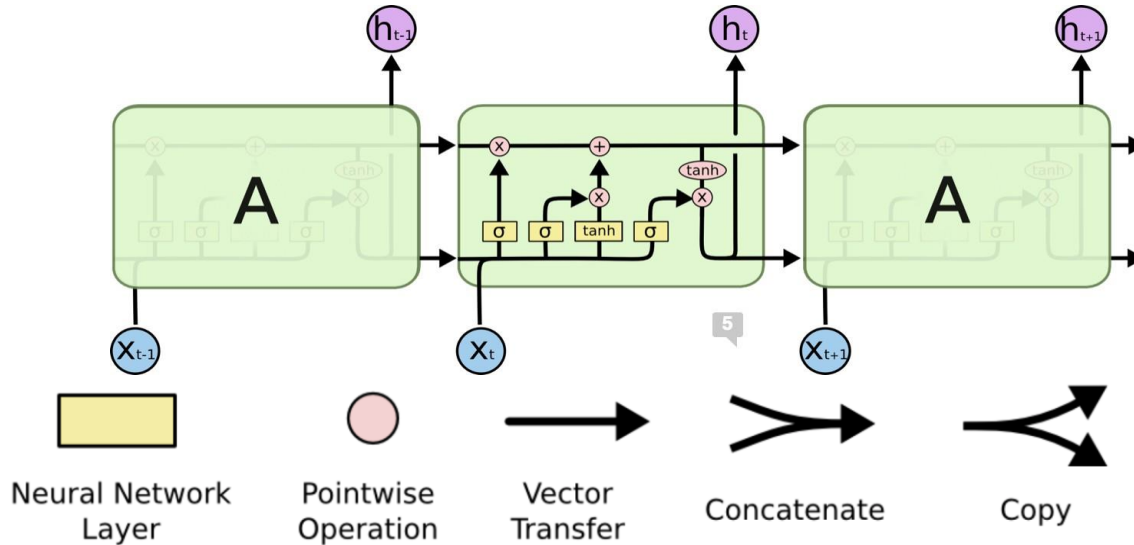
output
hidden
input



Computing the Hidden State



Long Short Term Memory Networks (LSTMs)



Sepp Hochreiter, Jürgen Schmidhuber:
Long Short-Term Memory.
In: *Neural Computation*. 9, 1997, S. 1735,
[doi:10.1162/neco.1997.9.8.1735](https://doi.org/10.1162/neco.1997.9.8.1735)


<https://colah.github.io/posts/2015-08-Understanding-LSTMs/>

Repetitive Variants: LSTMs, GRUs


- **Long Short Term Memory** networks (LSTMs) are RNs with a more complex architecture to combine the last hidden state with the current input.
- **Gated Recurrent Units** (GRUs) are a simplification of LSTMs
- Both contain “**gates**” to control how much of the input or past hidden state to forget or remember

Repetitive Variants: LSTMs, GRUs


RN



LSTM

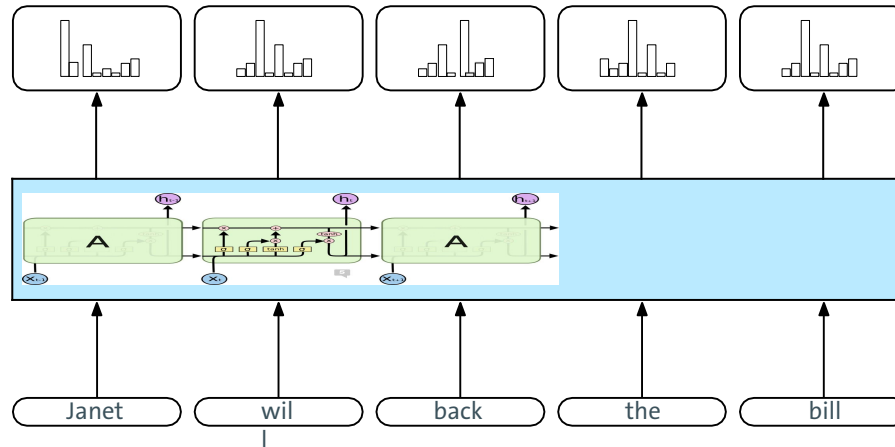


GRU




Basic RNs for Sequence Labeling

- Each time step has a distribution over output classes

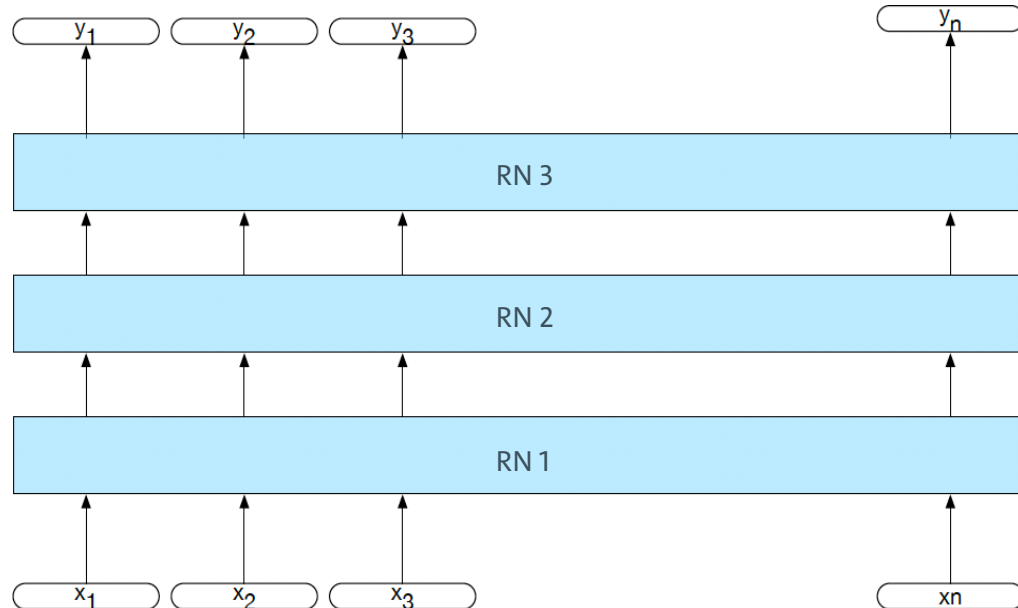


RNs for Sequence Classification




Stacked RNs

We can create an RN that has “vertical” depth (at each time step) by stacking multiple RNs:



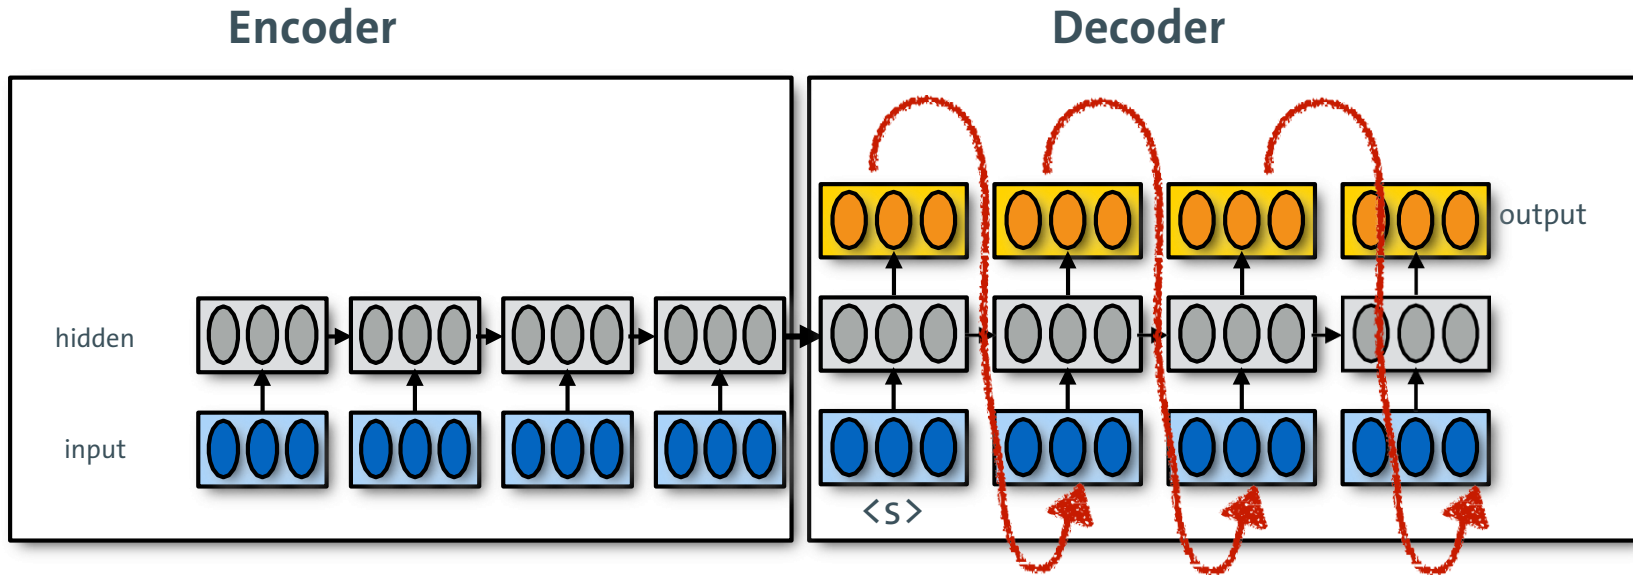
Bidirectional RNs




Encoder-Decoder (seq2seq) model

- Task: Read an input sequence and return an output sequence
 - Machine translation: translate source into target language
 - Dialog system/chatbot: generate a response
- Reading the input sequence: RN Encoder
- Generating the output sequence: RN Decoder

Encoder-Decoder (seq2seq) model



A More General View of seq2seq



Dzmitry Bahdanau, Kyunghyun Cho, Yoshua Bengio:
Neural Machine Translation by Jointly Learning to Align and Translate. ICLR **2015**

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS'17). Curran Associates Inc., Red Hook, NY, USA, 6000–6010. **2017**

Parallelize
with
GPUs

Embeddings from Language Models

Replace static embeddings (lexicon lookup) with **context-dependent embeddings** (produced by a deep language model)

=> Each token's representation is a function of the entire input sentence, computed by a deep **(multi-layer) bidirectional language model**


=> Return for each token a **(task-dependent) linear combination of its representation across layers.**

=> Different layers capture different information

Embeddings from **L**anguage **M**odel (ELMO)

- RN-based language models (trained from lots of sentences)

e.g., given “潮水 退了 就 知道 誰 沒穿 褲子”




潮水 退了 就 知道 誰 沒穿 褲子 = When the tide goes out, you know who's not wearing pants.

<https://arxiv.org/abs/1802.05365>

Each layer in deep LSTM can generate a latent representation.


Which one should we use???

ELMO



ELMO

High computational effort, word2vec to the rescue?

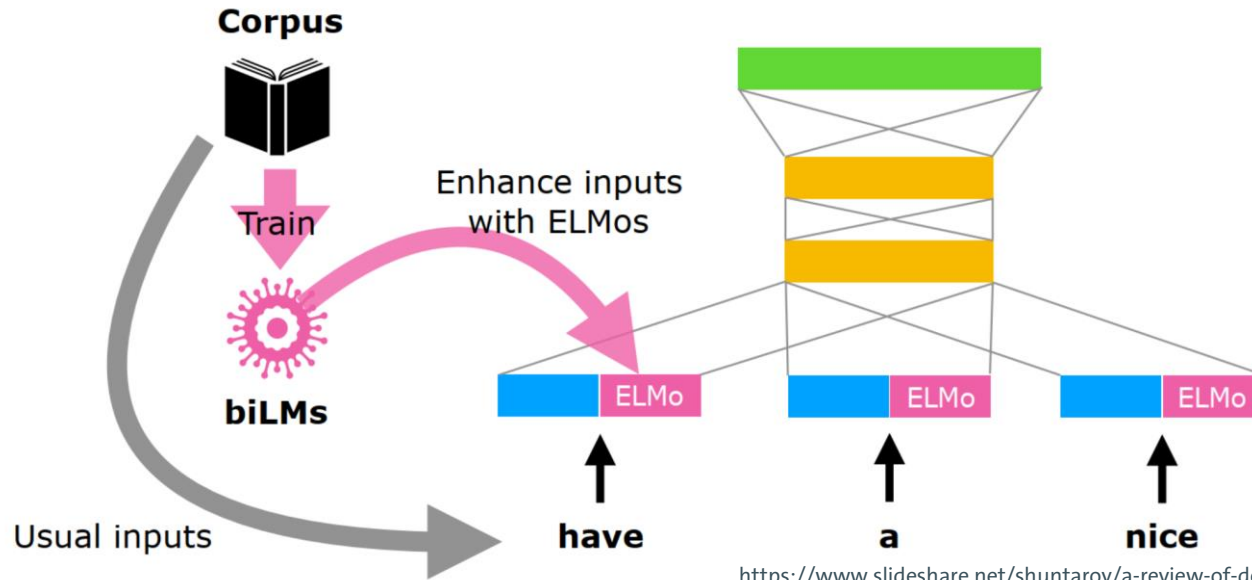


$$\text{Blue Bar} = \alpha_1 \text{Red Bar} + \alpha_2 \text{Grey Bar}$$

Learned with the down-stream tasks

http://speech.ee.ntu.edu.tw/~tlkagk/courses_ML20.html

Integrate ELMos into other embeddings



<https://www.slideshare.net/shuntaroy/a-review-of-deep-contextualized-word-representations-peters-2018>

Tricks: Subtoken Encoding

Byte Pair Encoding (BPE)

Word embedding sometimes is too high level, pure character embedding too low level.
For example, if we have learned

old older oldest

We might also wish the computer to infer

smart smarter smartest

But at the whole word level, this might not be so direct. Thus, the idea is to break the words up into pieces like er, est, and embed frequent fragments of words.

GPT adapts this BPE scheme.

Tricks: Subtoken Encoding

Byte Pair Encoding (BPE)

GPT uses BPE scheme. The subwords are calculated by:

1. Split word to sequence of characters (add `</w>` char)
2. Joining the highest frequency pattern.
3. Keep doing step 2, until it hits the pre-defined maximum number of sub-words or iterations.

Example (5, 2, 6, 3 are number of occurrences)

{`'l o w </w>': 5, 'l o w e r </w>': 2, 'n e w e s t </w>': 6, 'w i d e s t </w>': 3 }`

{`'l o w </w>': 5, 'l o w e r </w>': 2, 'n e w e s t </w>': 6, 'w i d e s t </w>': 3 }`

{`'l o w </w>': 5, 'l o w e r </w>': 2, 'n e w e s t </w>': 6, 'w i d e s t </w>': 3 }` (“est” freq. 9)

{`'l o w </w>': 5, 'l o w e r </w>': 2, 'n e w e s t </w>': 6, 'w i d e s t </w>': 3 }` (“lo” freq 7)

.....

The end of the neural AI era: Postneural AI

Attention Is All You Need

Ashish Vaswani*
Google Brain
avaswani@google.com

Noam Shazeer*
Google Brain
noam@google.com

Niki Parmar*
Google Research
nikip@google.com

Jakob Uszkoreit*
Google Research
usz@google.com

Llion Jones*
Google Research
llion@google.com

Aidan N. Gomez* †
University of Toronto
aidan@cs.toronto.edu

Łukasz Kaiser*
Google Brain
lukaszkaizer@google.com

Illia Polosukhin* ‡
illia.polosukhin@gmail.com

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS'17), 6000–6010. **2017**.


Self-attention

q : query (to match others) k : key (to be matched) v : information to be extracted

$$q^i = W^q a^i$$

$$k^i = W^k a^i$$


$$v^i = W^v a^i$$



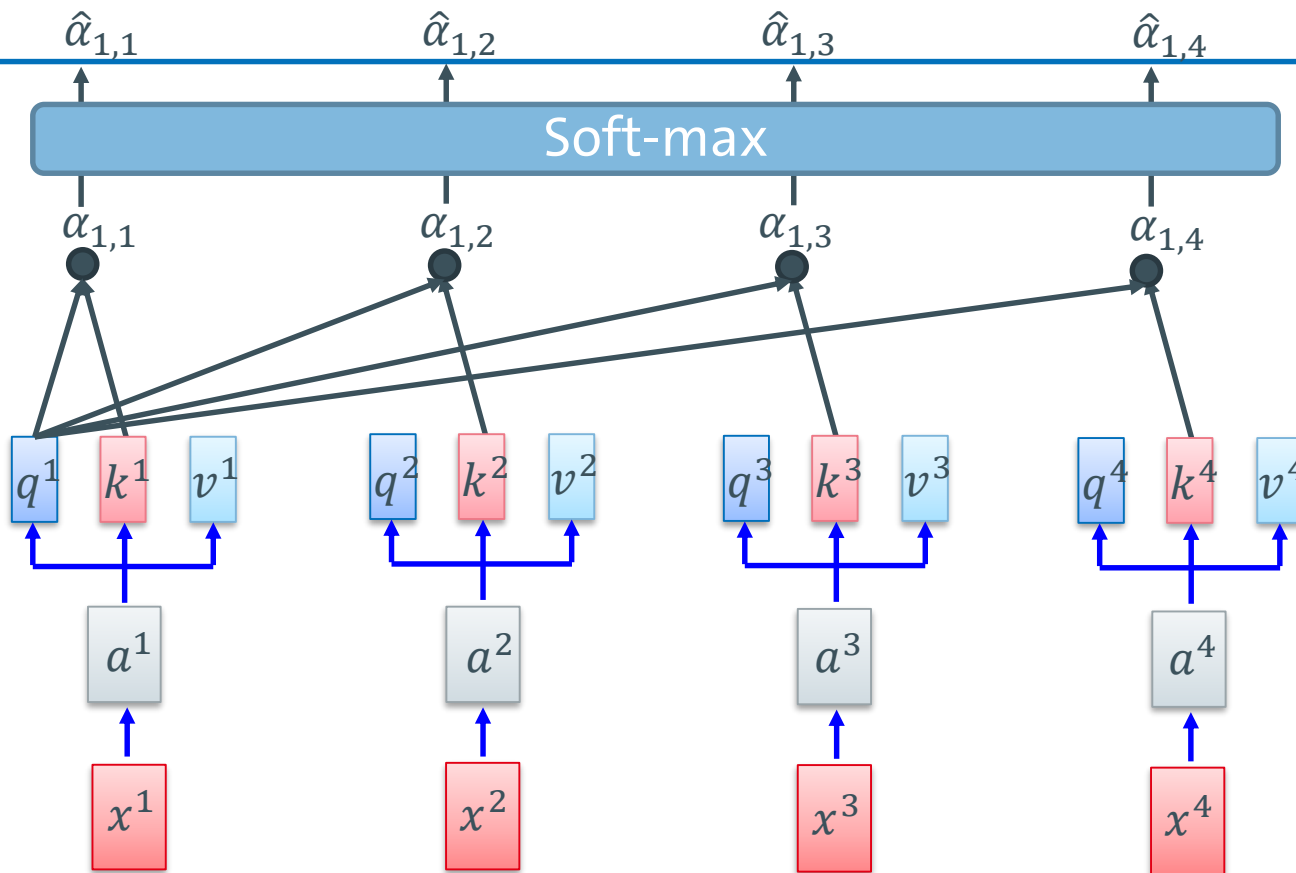
Scaled Dot-Product Attention: $\alpha_{1,i} = \underbrace{q^1 \cdot k^i}_{\text{dot product}} / \sqrt{d}$

Self-attention

Take each query q ,
go to each key k ,
do attention

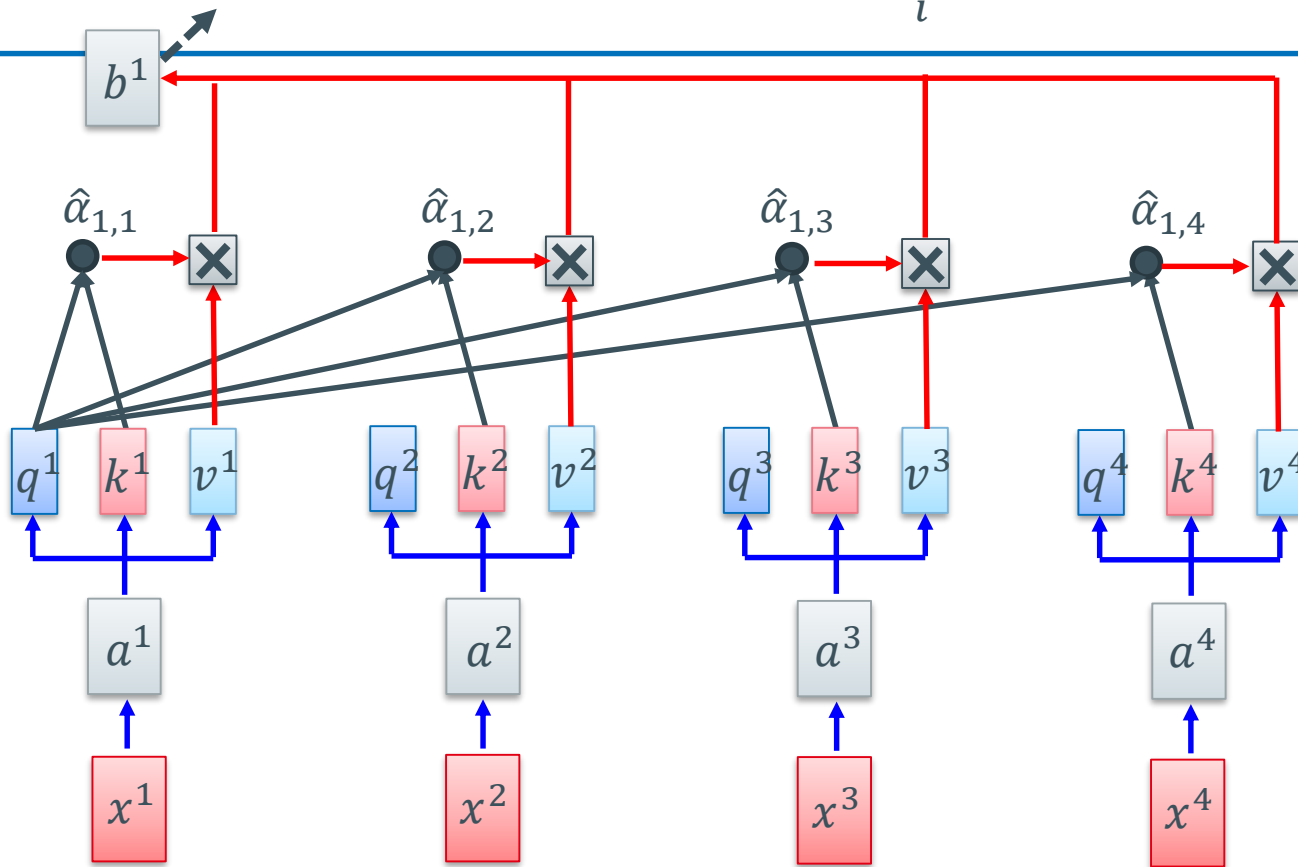


$$\hat{\alpha}_{1,i} = \exp(\alpha_{1,i}) / \sum_j \exp(\alpha_{1,j})$$



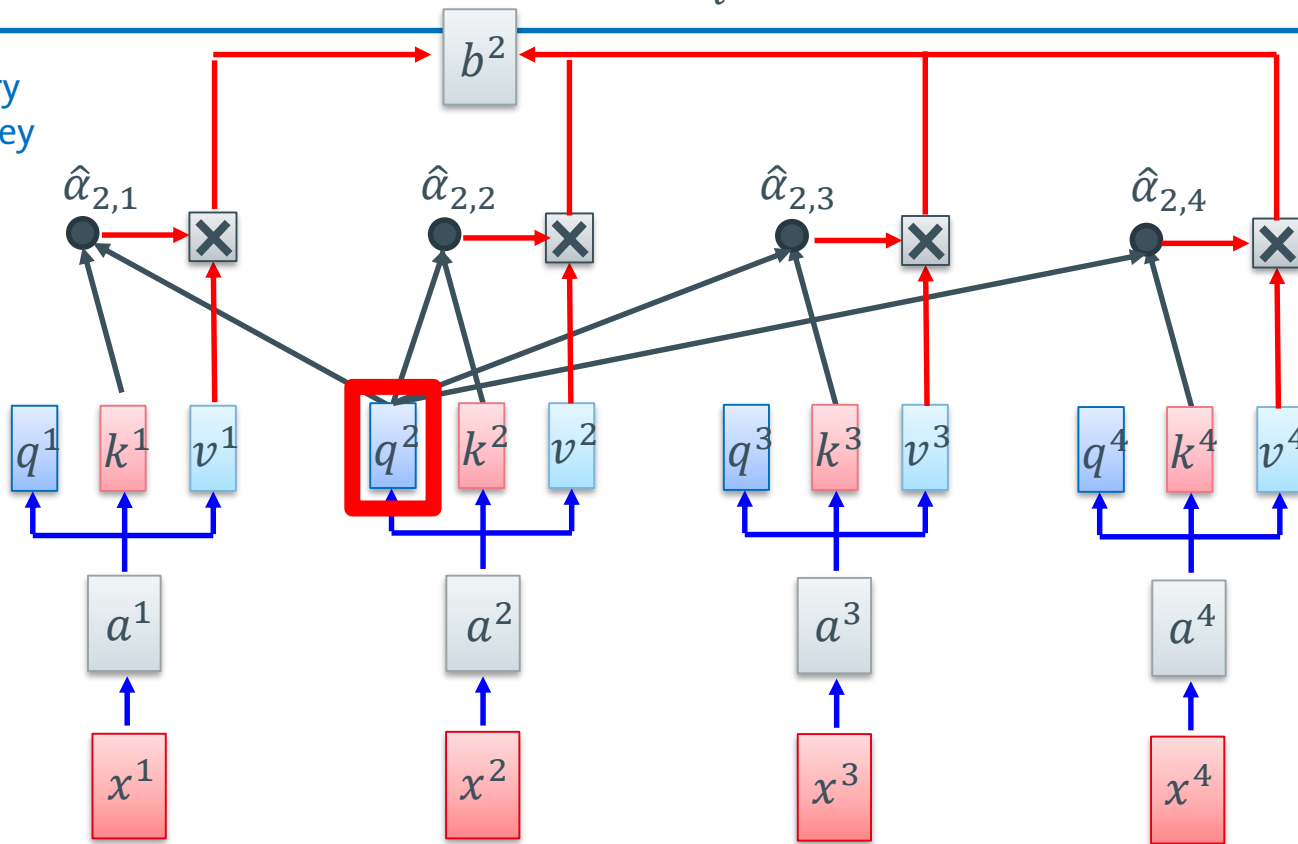
Considering the whole sequence

$$b^1 = \sum_i \hat{\alpha}_{1,i} v^i$$




$$b^2 = \sum_i \hat{\alpha}_{2,i} v^i$$

Take each query q , go to each key k , do attention



b^1, b^2, b^3, b^4 can be computed in parallel



$$\underbrace{q^1 q^2 q^3 q^4}_Q = \underbrace{W^q}_{} \underbrace{a^1 a^2 a^3 a^4}_I$$


$$q^i = W^q a^i$$

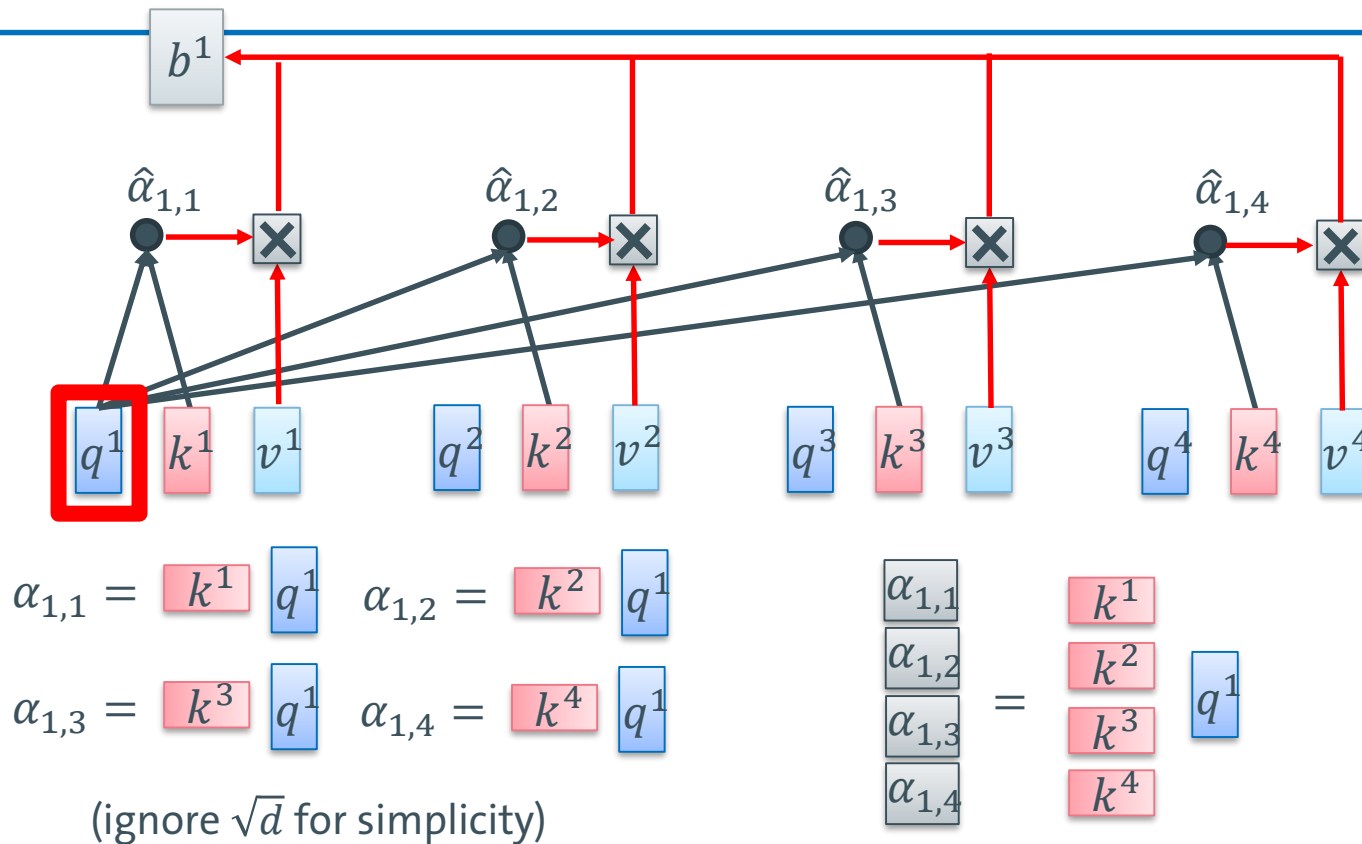
$$k^i = W^k a^i$$

$$v^i = W^v a^i$$

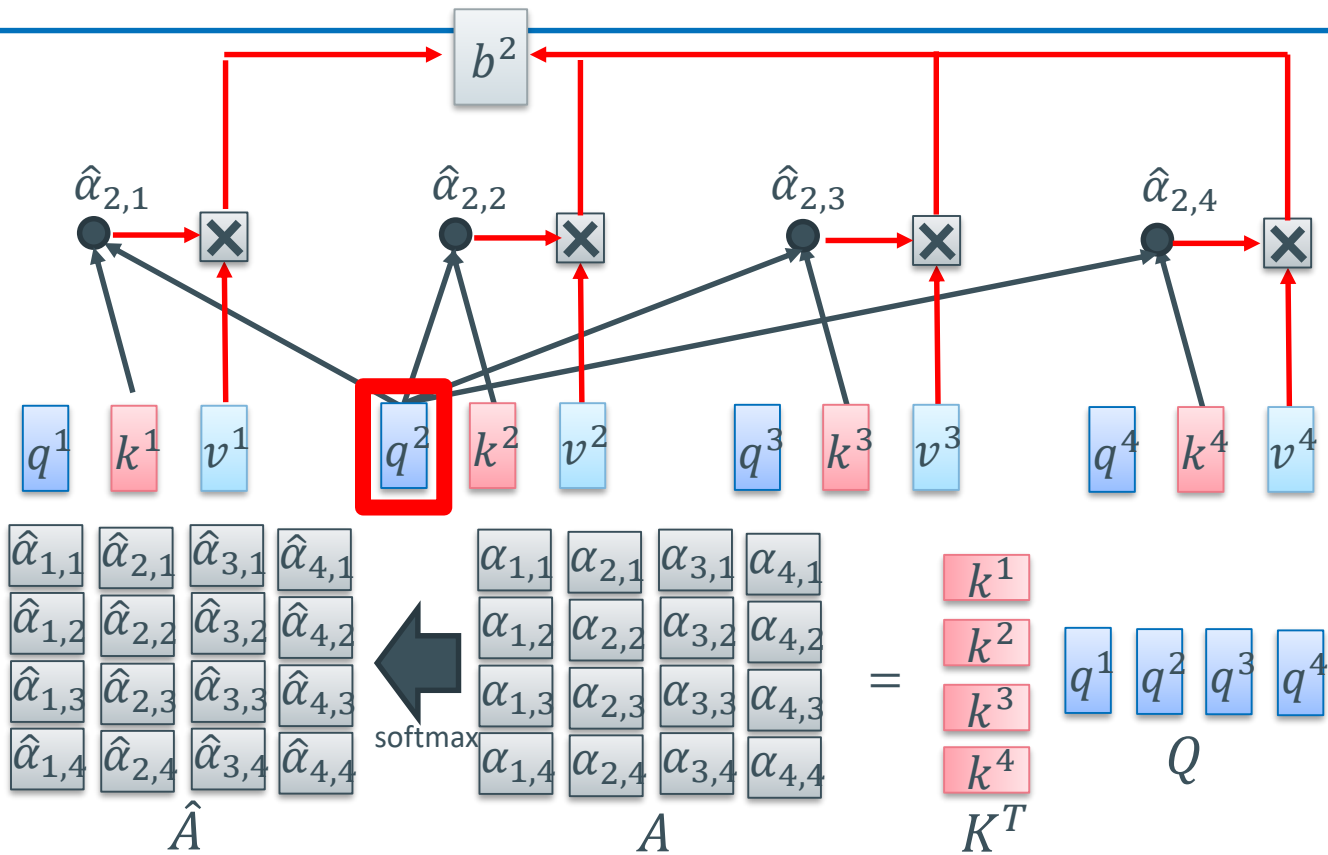
$$\underbrace{k^1 k^2 k^3 k^4}_K = \underbrace{W^k}_{} \underbrace{a^1 a^2 a^3 a^4}_I$$

$$\underbrace{v^1 v^2 v^3 v^4}_V = \underbrace{W^v}_{} \underbrace{a^1 a^2 a^3 a^4}_I$$

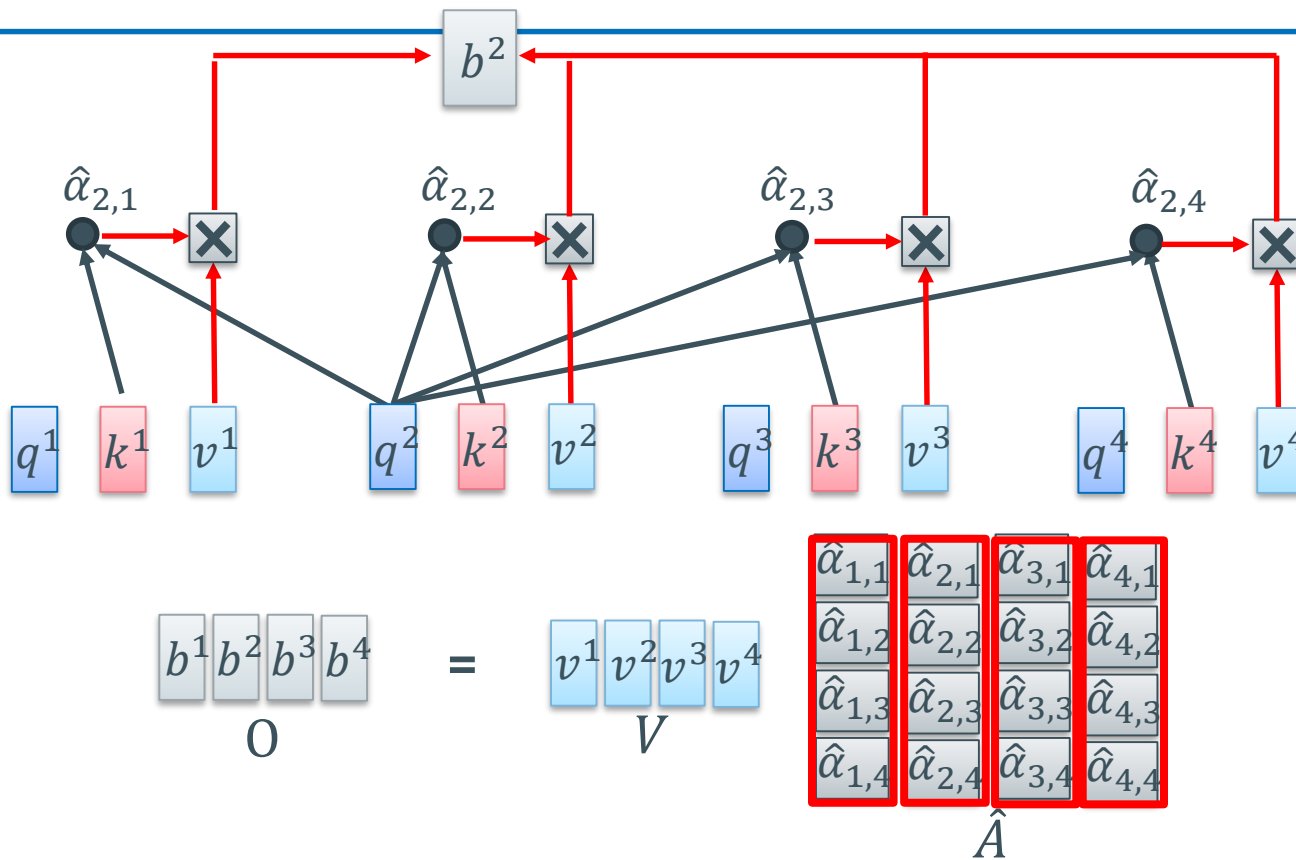




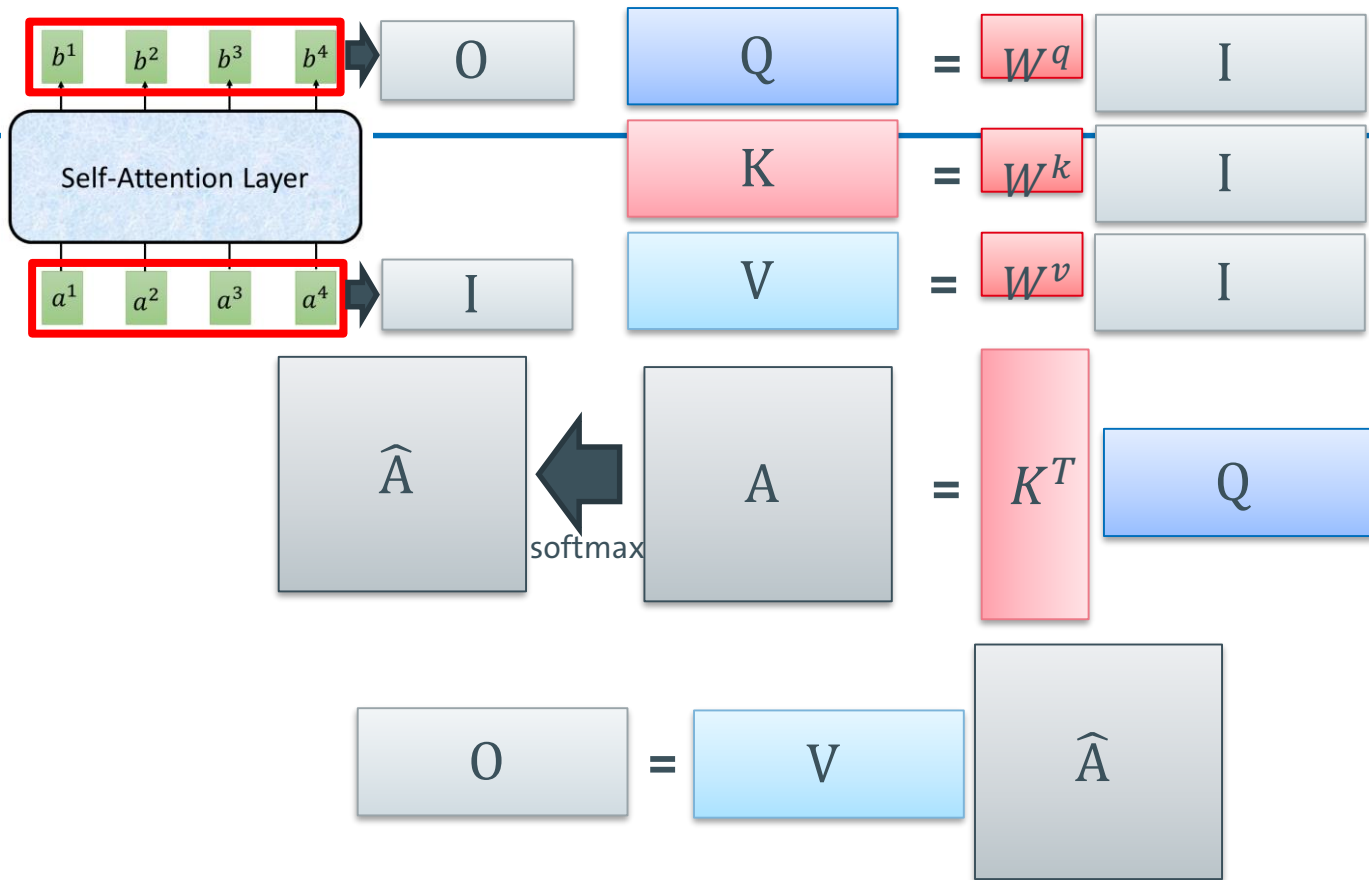
$$b^2 = \sum_i \hat{\alpha}_{2,i} v^i$$



$$b^2 = \sum_i \hat{a}_{2,i} v^i$$



$$\begin{matrix} b^1 & b^2 & b^3 & b^4 \\ O \end{matrix} = \begin{matrix} v^1 & v^2 & v^3 & v^4 \\ V \end{matrix}$$



Can be optimized with GPUs

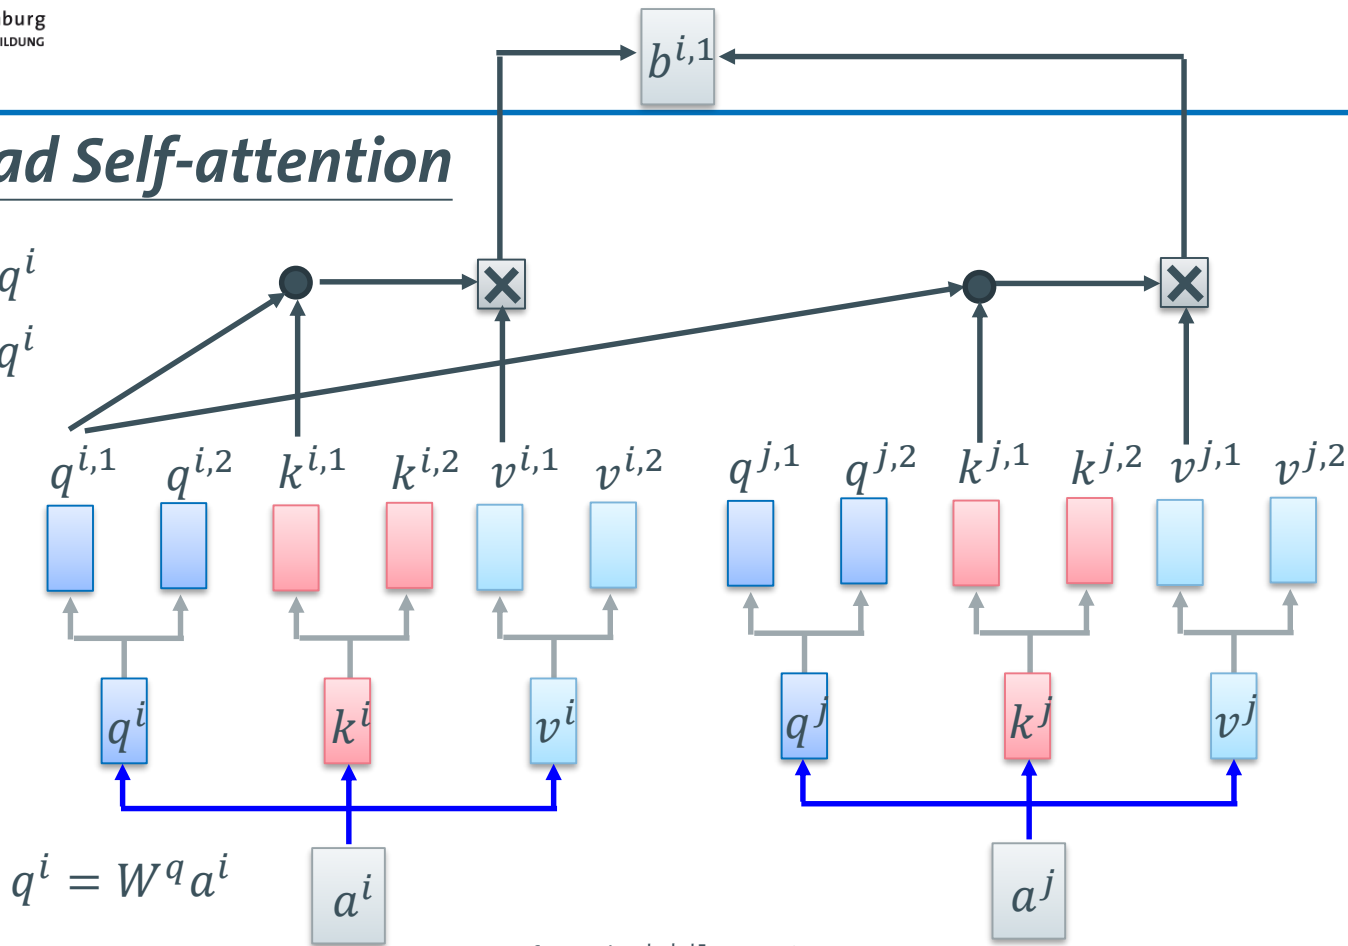
Transformer by 李宏毅 Hung-yi Lee

(2 heads as example)

Multi-head Self-attention

$$q^{i,1} = W^{q,1} q^i$$

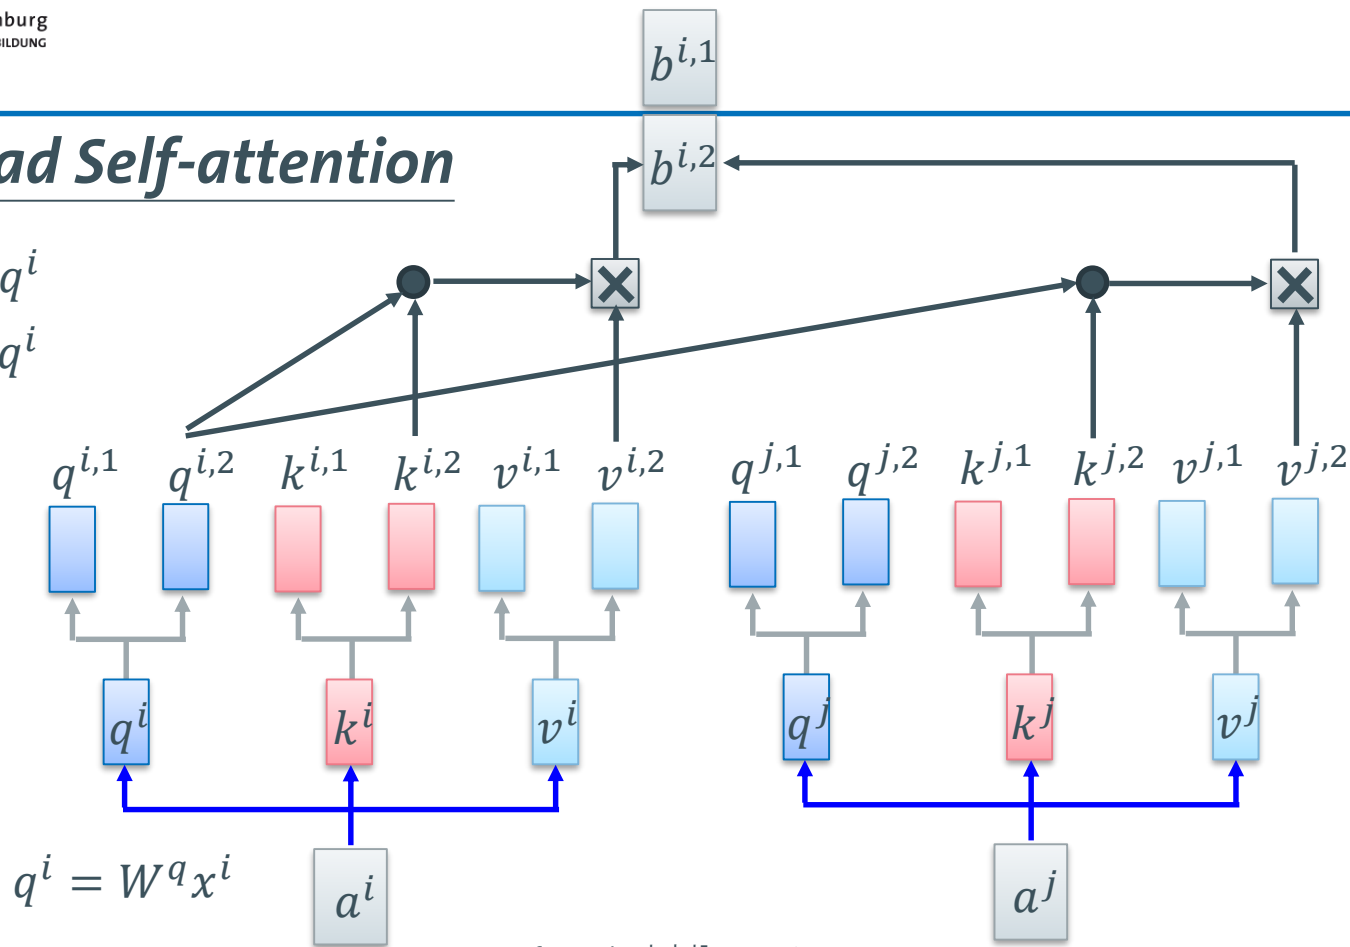
$$q^{i,2} = W^{q,2} q^i$$





Multi-head Self-attention


$$q^{i,1} = W^{q,1} q^i$$

$$q^{i,2} = W^{q,2} q^i$$



(2 heads as example)






Positional Encoding

No position information in self-attention.


Each position has a unique positional vector e^i (not learned from data)

Idea: Append each x^i is with a one-hot vector p^i


More clever solution used in the original paper



Seq2seq with Attention



Encoder

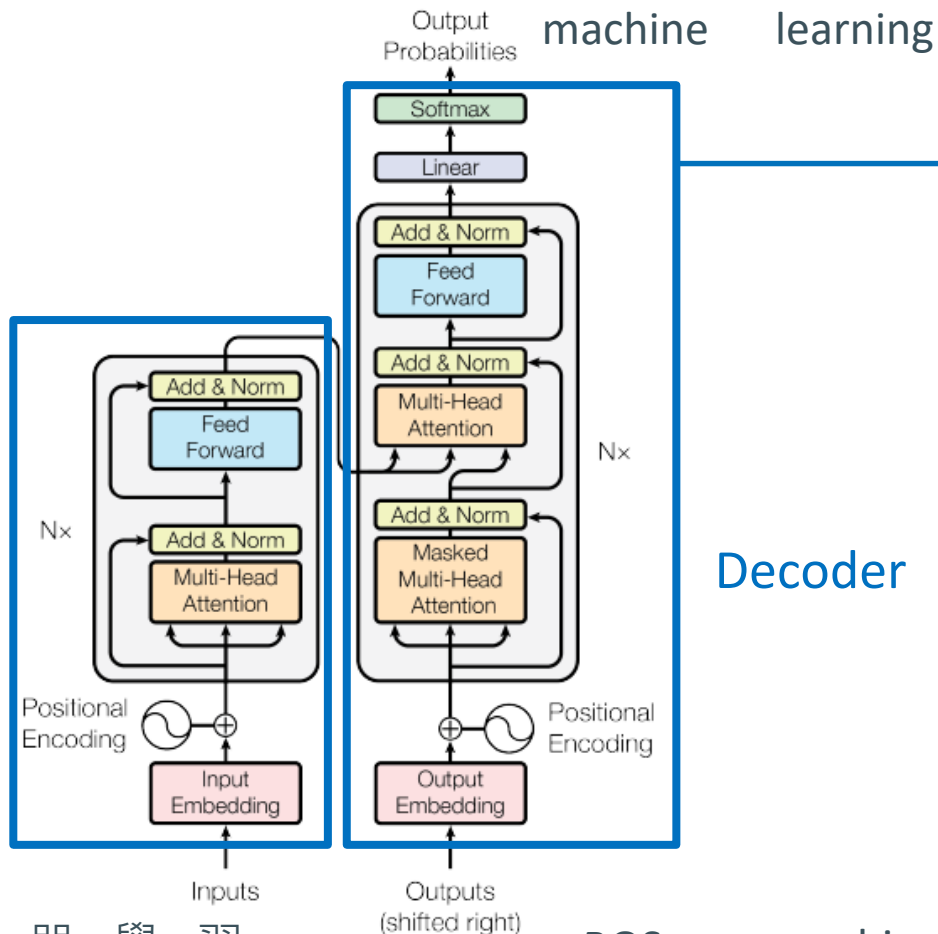


Decoder

Transformer

Using Chinese to English translation as example

Encoder

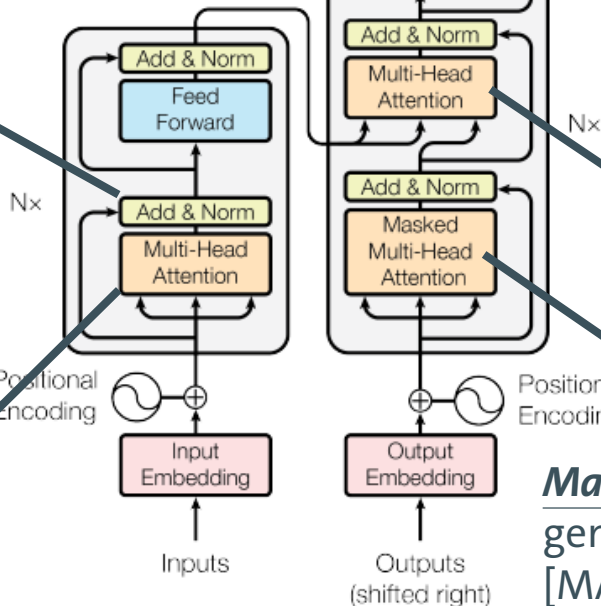





Transformer

Layer Norm:
<https://arxiv.org/abs/1607.06450>

Batch Norm:
<https://www.youtube.com/watch?v=BZh1ltr5Rkg>

Batch Size



Masked: attend on the generated sequence [MASK]

Masked Multihead Attention

Decoder should work in parallel as well

During training all output tokens are known


Copy output #token times

For each position use [MASK] token in copies

Attention becomes possible during training also for decoding


Train decoder such that [MASK] is replaced correctly while paying attention to the overall output training data

Attention Visualization




<https://arxiv.org/abs/1706.03762>

Attention Visualization



The encoder self-attention distribution for the word “it” from the 5th to the 6th layer of a Transformer trained on English to French translation (one of eight attention heads).

Multi-head Attention




Pre-Training & Fine Tuning

1) Download LM


pre-trained on large corpus
(in self-supervised fashion)

2) Feature-based training ("fine-tuning")

on target task
(supervised learning)



Model Pre-Training



Encoder

- Bidirectional context
- Examples: BERT and its variants


Decoder

- Language modeling; better for generation
- Example: GPT-2, GPT-3, [LaMDA](#)

Encoder-Decoder

- Sequence-to-sequence model
- Examples: Transformer, BART, T5

Model Pre-Training



Encoder

- Bidirectional context
- Examples: BERT and its variants

Decoder

- Language modeling; better for generation
- Example: GPT-2, GPT-3, [LaMDA](#)

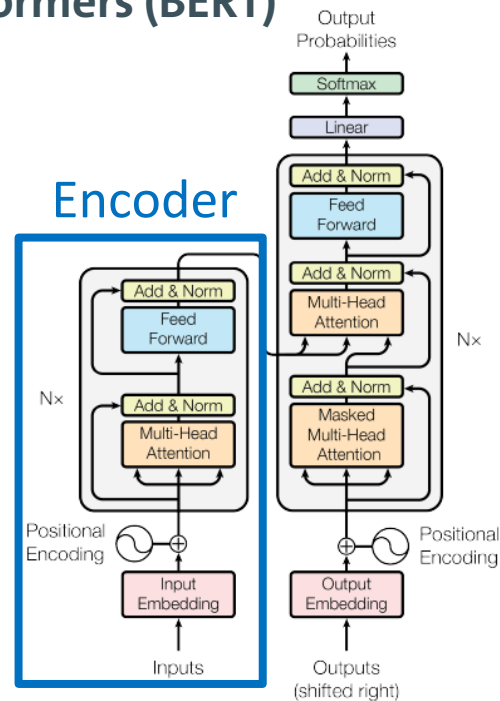

Encoder-Decoder

- Sequence-to-sequence model
- Examples: Transformer, BART, T5

Bidirectional Encoder Representations from Transformers (BERT)

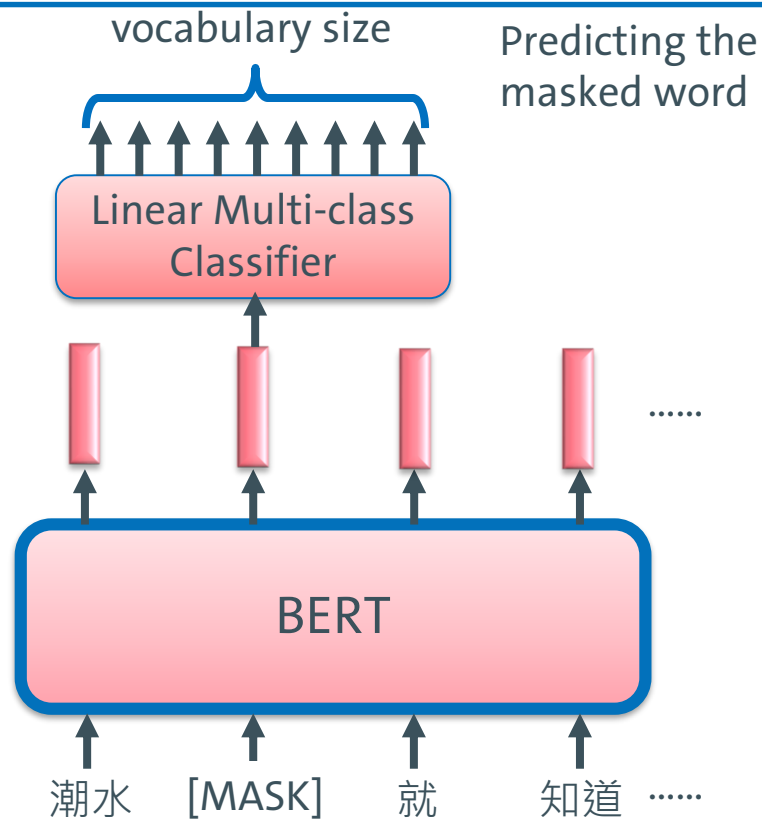
BERT = Encoder of Transformer

Learned from a large amount of text
without annotation

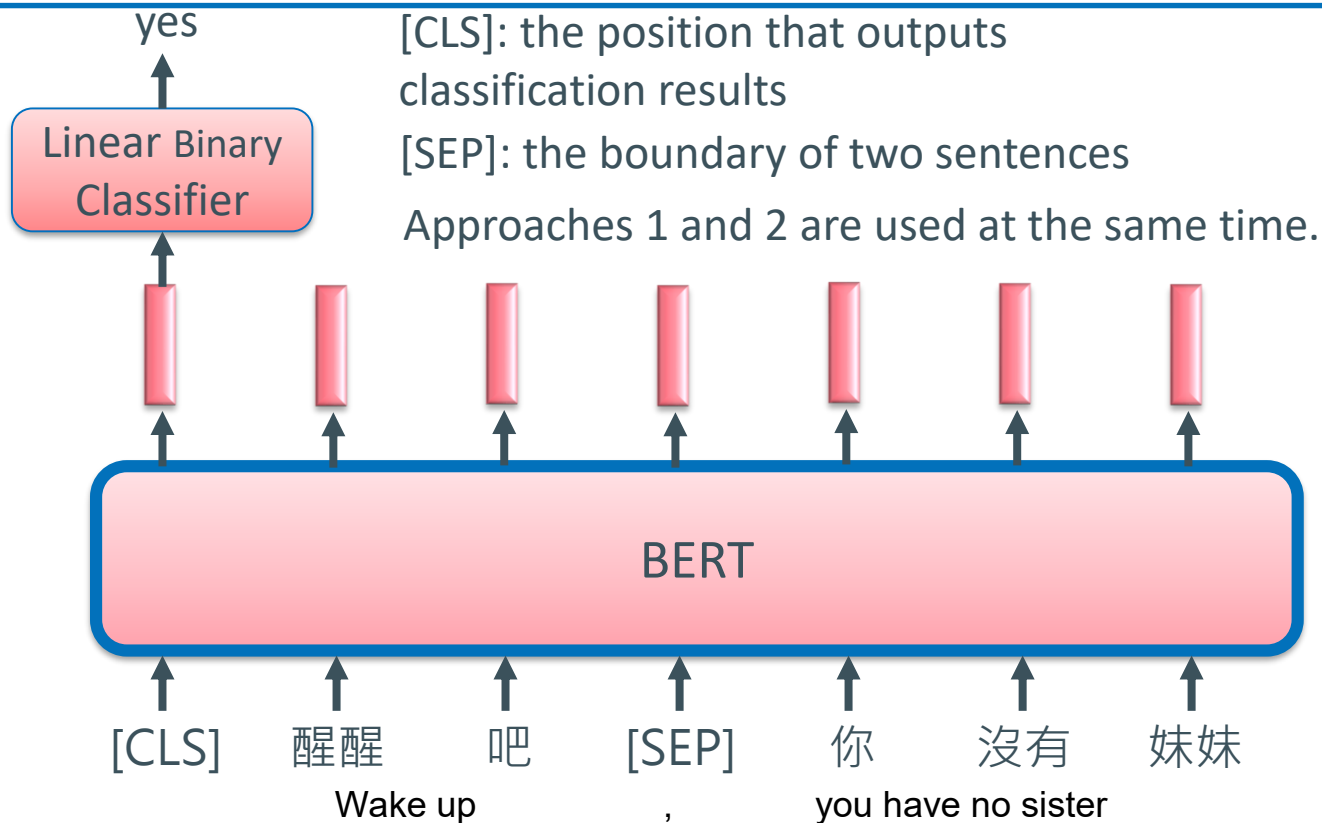


Training of BERT

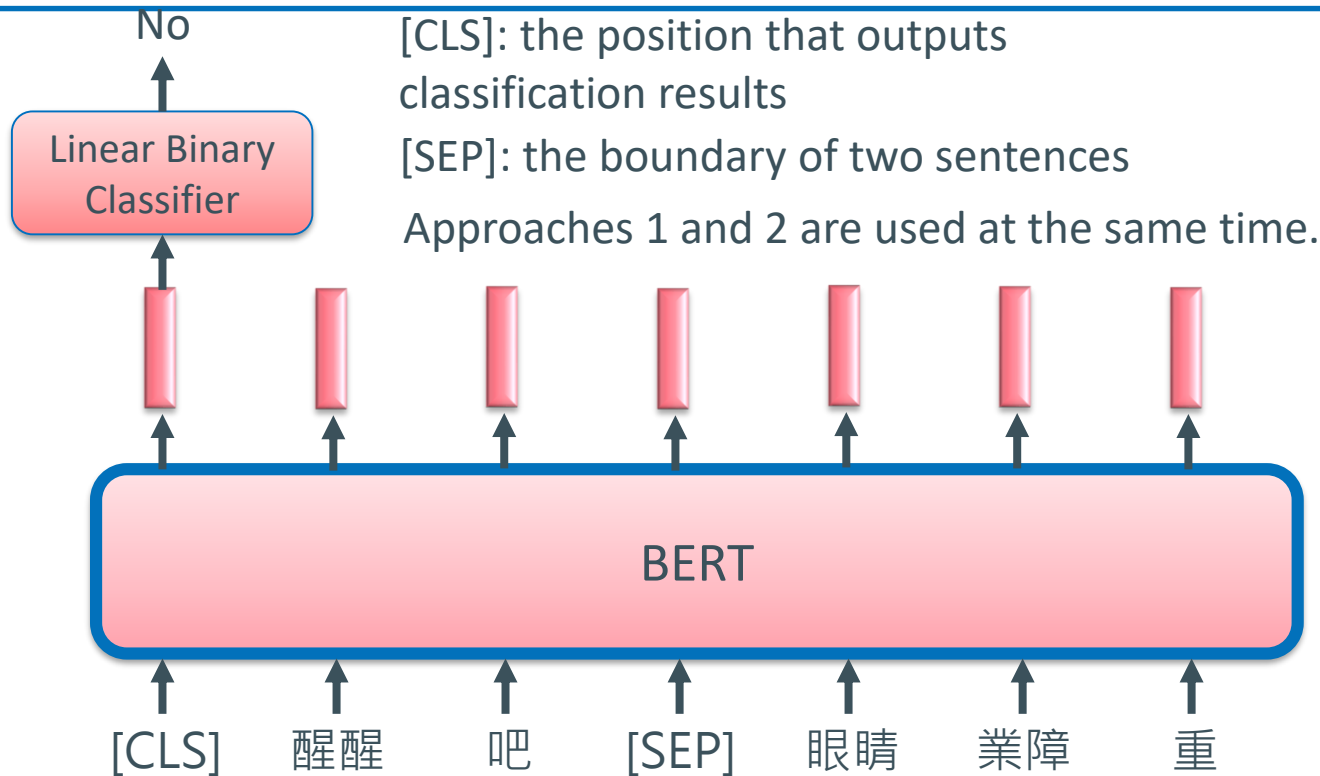
Approach 1: Masked LM



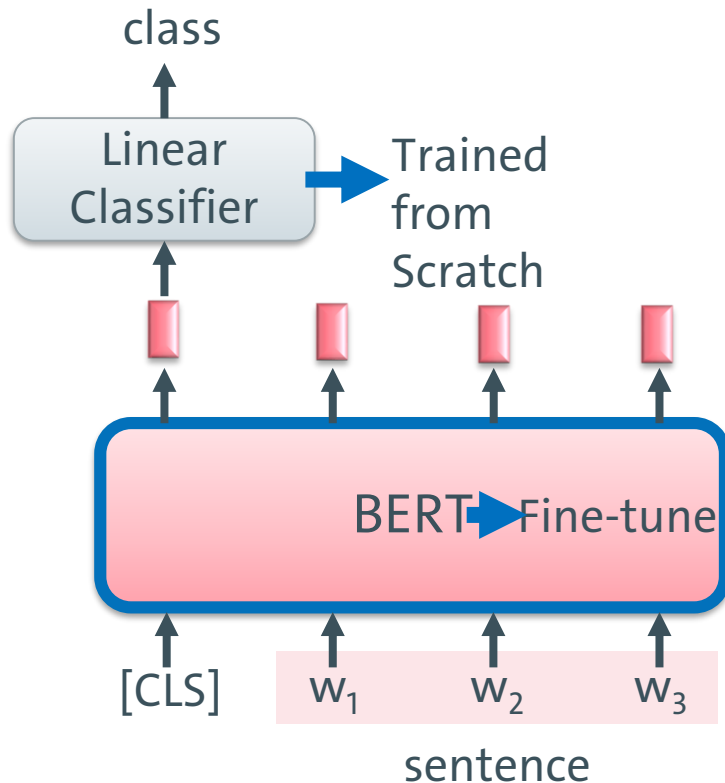
Training of BERT – Approach 2: Next Sentence Prediction



Training of BERT – Approach 2: Next Sentence Prediction




How to use BERT – Case 1



Input: single sentence,
output: class


Example:
Sentiment analysis,
Document
Classification

How to use BERT – Case 2

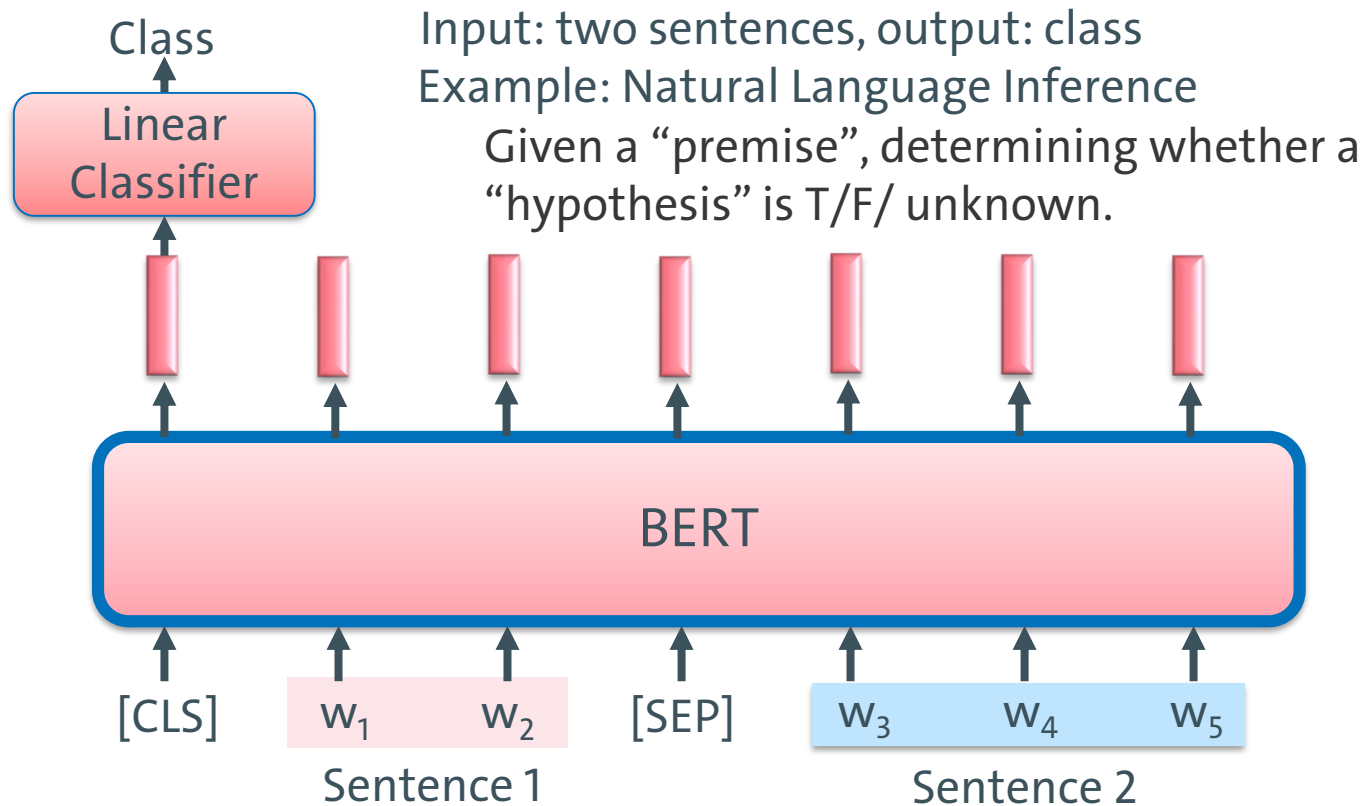


Input: single sentence,
output: class of each word

Example: Semantic
role labelling



How to use BERT – Case 3



How to use BERT – Case 4

Extraction-based Question Answering (QA)

(E.g. SQuAD)

Document: $D = \{d_1, d_2, \dots, d_N\}$

Query: $Q = \{q_1, q_2, \dots, q_N\}$



output: two integers (s, e)

Answer: $A = \{q_s, \dots, q_e\}$

S=start, e=end

In meteorology, precipitation is any product of the condensation of 17 spheric water vapor that falls under **gravity**. The main forms of precipitation include drizzle, rain, sleet, snow, **grau-pel** and hail... Precipitation forms as smaller droplets coalesce via collision with other rain drops or ice crystals **within a cloud**. Short, intense periods of rain 77 at 79 locations are called "showers".

What causes precipitation to fall?

gravity

$s = 17, e = 17$

What is another main form of precipitation besides drizzle, rain, snow, sleet and hail?

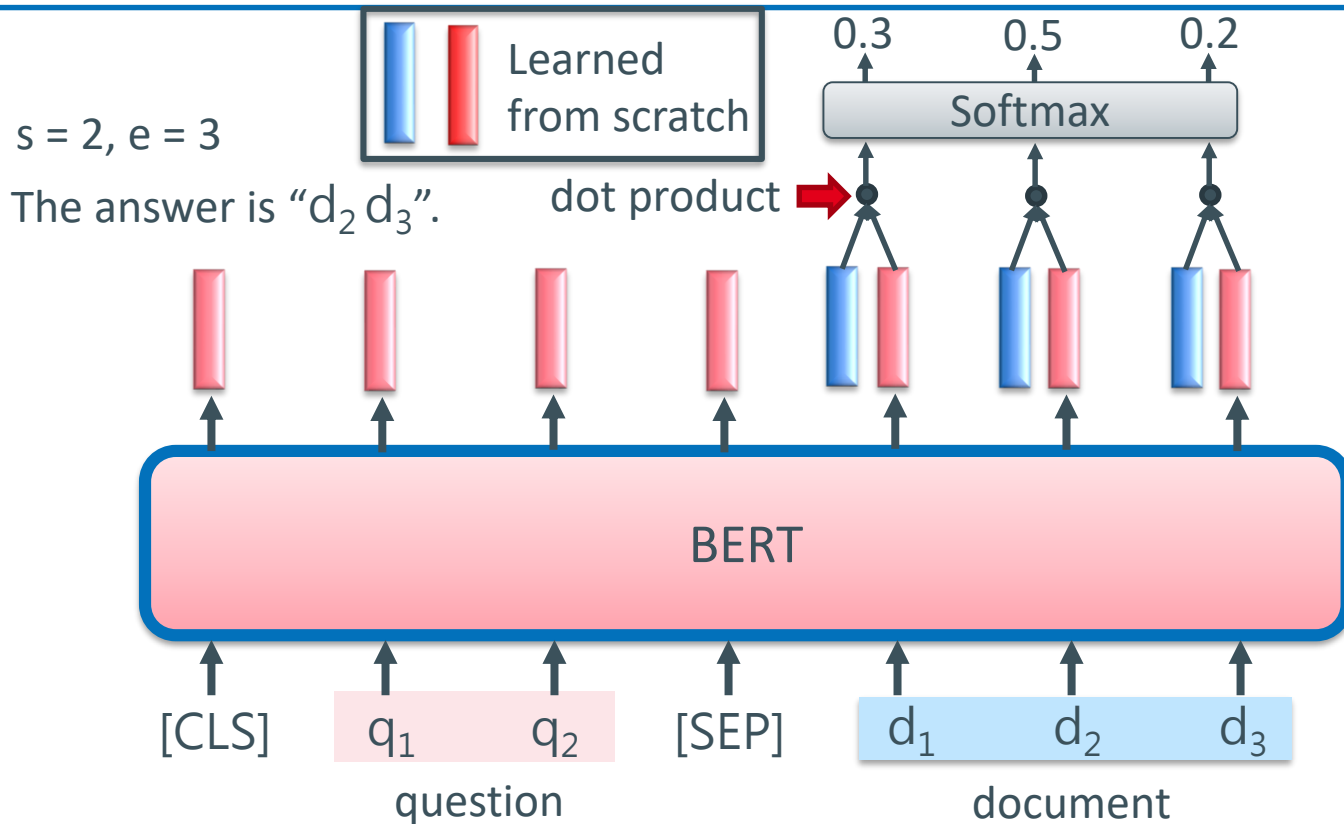
grau-pel

Where do water droplets collide with ice crystals to form precipitation?

within a cloud

$s = 77, e = 79$

How to use BERT – Case 4




BERT Pre-Training & Fine Tuning

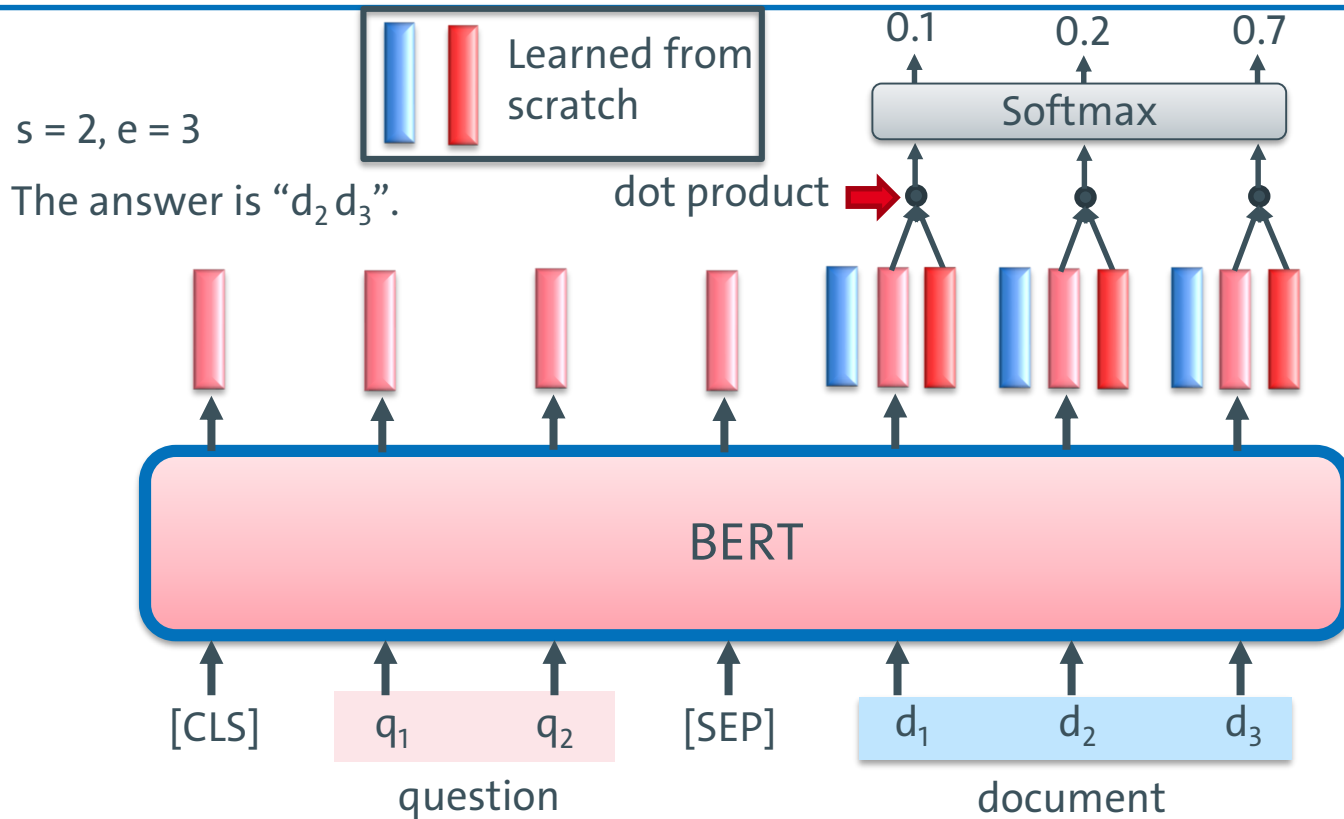
- Keep BERT frozen after pre-training
- Create BERT embeddings for labeled dataset for "downstream task" and train new model on these embeddings

1) Download BERT pre-trained on large corpus (in self-supervised fashion)

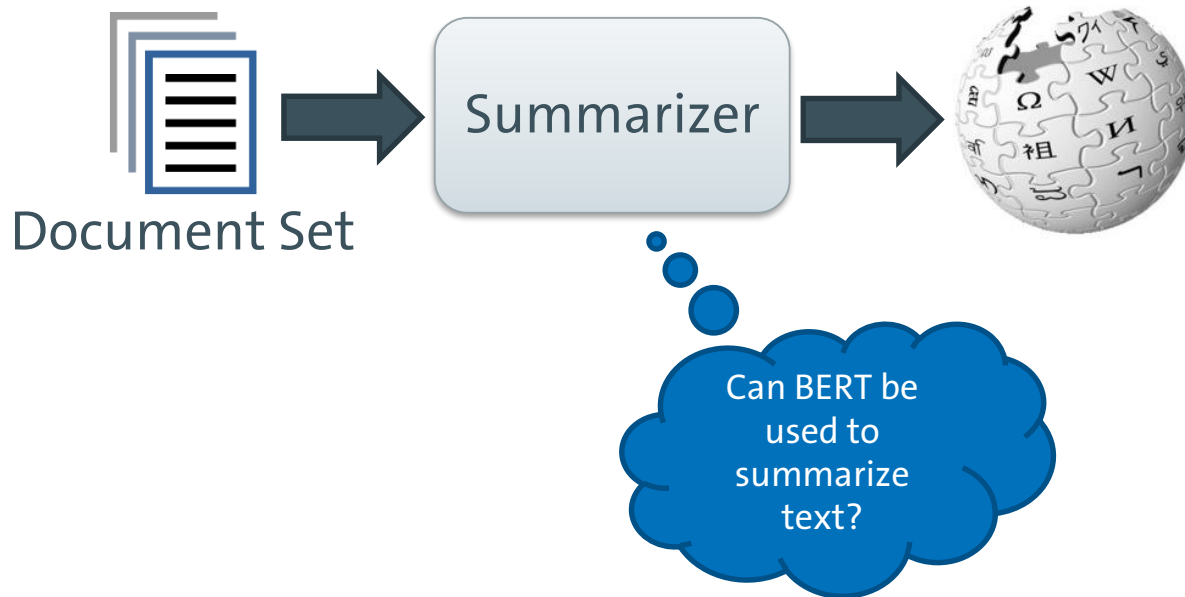
2) Feature-based training ("fine-tuning") on target task (supervised learning)




How to use BERT – Case 4




Example Application: Summarization




Model Pre-Training



- Encoder
 - Bidirectional context
 - Examples: BERT and its variants



- Decoder
 - Language modeling; better for generation
 - Example: GPT-2, GPT-3, LLama, LaMDA



- Encoder-Decoder
 - Sequence-to-sequence model
 - Examples: Transformer, BART, T5

LaMDA: Language Models for Dialog Applications. LaMDA is a family of Transformer- based neural language models specialized for dialog, which have up to 137B parameters and are pre-trained on 1.56T words of public dialog data and web text

GPT (Generative Pre-trained Transformer)

- Developed by OpenAI
- Unidirectional: “trained to predict next word in a sentence”

GPT (110 million parameters)

Radford, A., Narasimhan, K., Salimans, T., & Sutskever, I. (2018). Improving language understanding by generative pre-training. https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf

GPT-2 1.5 billion parameters)

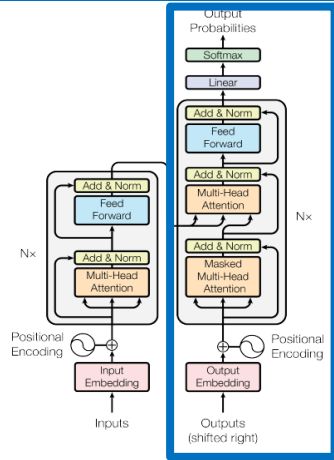
Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., & Sutskever, I. (2019). Language models are unsupervised multitask learners. *OpenAI blog*, 1(8), 9.
https://cdn.openai.com/better-language-models/language_models_are_unsupervised_multitask_learners.pdf

GPT-3 (175 billion parameters)

Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., ... & Amodei, D. (2020). Language models are few-shot learners. arXiv preprint arXiv:2005.14165. <https://arxiv.org/abs/2005.14165>

GPT

<https://openai.com/index/better-language-models/>




BERT
(340M)

ELMO
(94M)

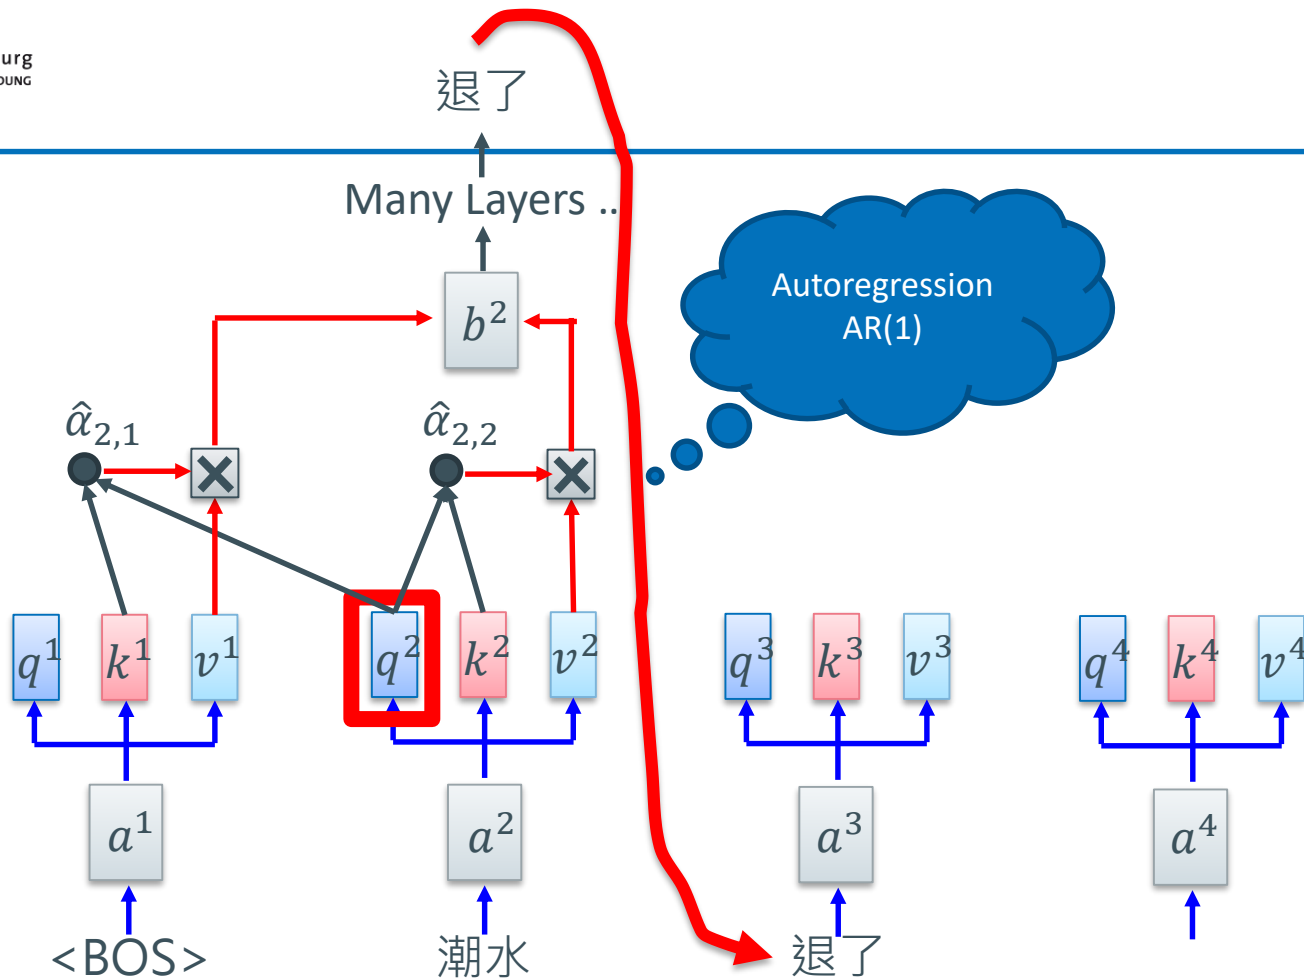


Transformer
Decoder

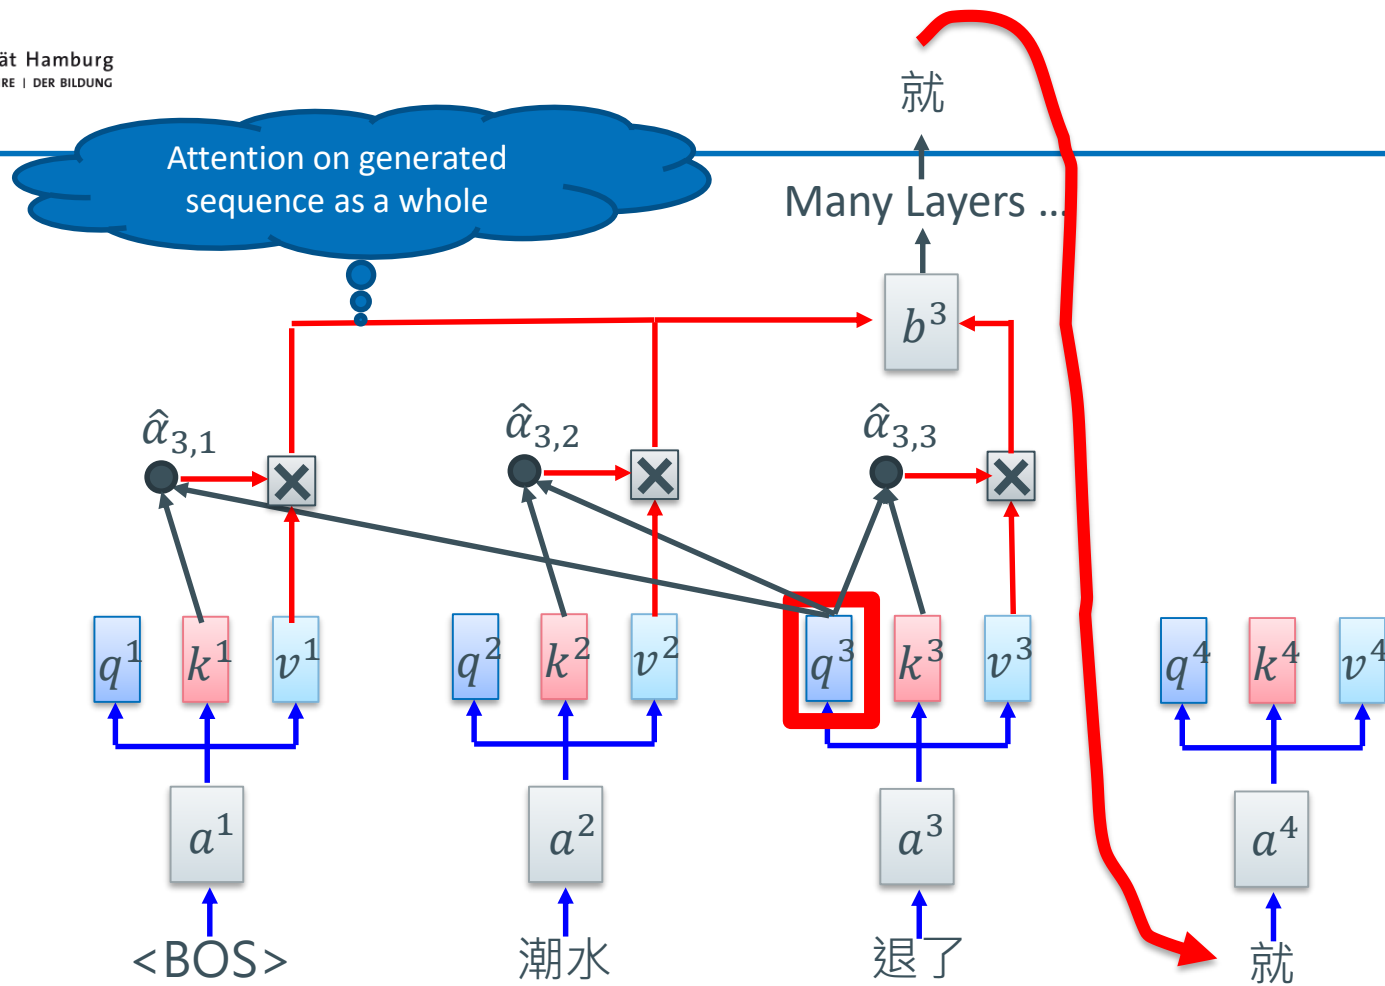


GPT-2
(1.5B)

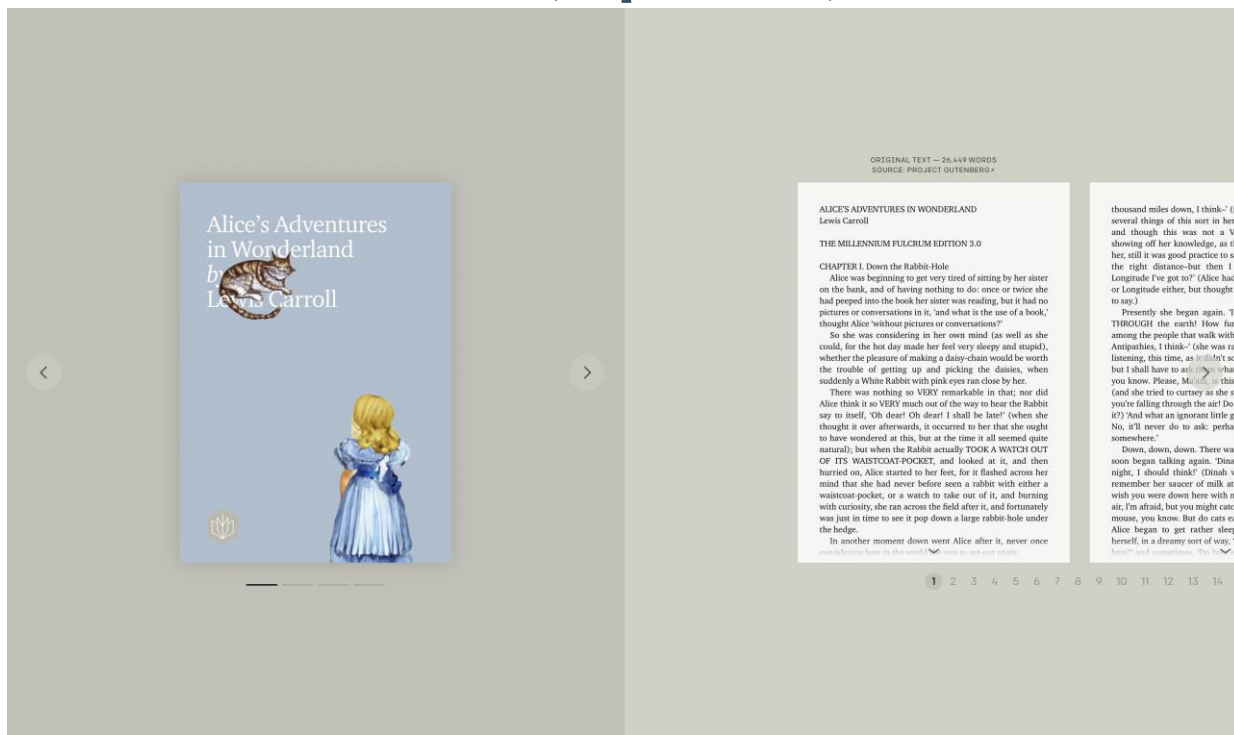
GPT



GPT



Application: Summaries (Open AI)



Alice's Adventures in Wonderland by Lewis Carroll



The original text is divided into sections,
and each section is summarized.

66 SUMMARIES — 6,024 WORDS

Alice is bored sitting by her sister on the bank, and she's thinking about making a daisy chain when a white rabbit with pink eyes runs by. She's surprised to see a rabbit with a waistcoat pocket and a watch, and she follows it down a rabbit hole. She falls down a deep well, and as she's falling she takes note of the shelves and jars she passes. She thinks about how brave she'll seem when she gets home, and how she'll never tell anyone about her fall. She wonders how deep the well is and how many miles she's fallen.

Alice thinks she's fallen a thousand miles, but she doesn't know what latitude or longitude she's at. She falls right through the earth and wonders what she'll emerge in. She starts to get scared, but she's not alone. She's about walking hand in hand with her sister, but she's never eaten a bat. Suddenly she's back to back with her sister, and dry leaves, and she's not hurt. She's in a moment.

Alice's Adventures in Wonderland by Lewis Carroll




Section summaries are summarized
again into higher-level summaries.

6 SUMMARIES — 830 WORDS

Alice falls down a deep well and lands in a heap of sticks and leaves. She follows the White Rabbit down a long passage, but loses sight of it. She finds a bottle that says "Drink Me" on it and drinks it, shrinking to 10 inches tall. She grows to 9 feet tall and then shrinks again. She cries and grows larger again. A mouse appears and Alice asks it for help. The mouse agrees to tell Alice its story if she'll take it to the shore. Once there, the animals discuss how to get dry again. The Dodo suggests a Caucus-race to get dry. Everyone crowds around Alice asking for prizes. Alice pulls out a box of comfits and hands them out as prizes. The Mouse then asks Alice if she has a prize for herself, and she says she only has a thimble. The Dodo takes the thimble and presents it to Alice, saying they all beg her acceptance.

Alice grows larger after drinking the tea. She stops growing before she reaches the table. She then starts to miss being at home and decides to go down the rabbit hole. She hears the White Rabbit call her name and follows it. She realizes she's now much bigger than the White Rabbit and is no longer afraid of it. Alice decides to go back to her normal size, and her sister and brother follow her way into the garden. She doesn't like the garden, however. She sees a Caterpillar and asks it for those things, however. She sees a King and a Queen and asks them for a hookah on a mushroom and asks them to make her grow back to her normal size. The Caterpillar tells her that the mushroom will make her grow, and she eats it. She starts eating the mushroom and grows back to her normal size.



The summarizing process continues until a complete summary is achieved.

COMPLETE SUMMARY — 136 WORDS

Alice falls down a rabbit hole and grows to giant size after drinking a mysterious bottle. She decides to focus on growing back to her normal size and finding her way into the garden. She meets the Caterpillar who tells her that one side of a mushroom will make her grow taller, the other side shorter. She eats the mushroom and returns to her normal size. Alice attends a party with the Mad Hatter and the March Hare. The Queen arrives and orders the execution of the gardeners for making a mistake with the roses. Alice saves them by putting them in a flowerpot. The King and Queen of Hearts preside over a trial. The Queen gets angry and orders Alice to be sentenced to death. Alice wakes up to find her sister by her side.


LLMs of the GPT family

- Generate term papers
 - Generate code
 - Generate Powerpoint presentations
 - Generate useful completions of texts in Word?
 - ...
-
- The latest version is GPT-4o (“o” for “omni”)
 - Announced: GPT-5


From Zero to ChatGPT

Image: Graham Neubig, Large Language Models:
A Bird's Eye View, <https://phontron.com/slides/neubig23llms.pdf>


Lots of web text




Lots of GitHub code





Lots of annotated data




Human judgements of response quality




Chat-oriented data




davi nci




code-
davi nci - 002



t ext -
davi nci - 002




t ext -
davi nci - 003



davi nci -
3. 5- t ur bo
(Chat GPT)


Compatibility vs. Alignment in LLMs

A model's *capability* is typically evaluated by **how well it is able to optimize its objective function**, the mathematical expression that defines the goal of the model



High capability
Low alignment

Alignment, on the other hand, is concerned with **what we actually want the model to do** versus what it is being trained to do



Low capability
High alignment

Models like the original GPT-3 are misaligned



Improving Language Model Behavior by Training on a Curated Dataset

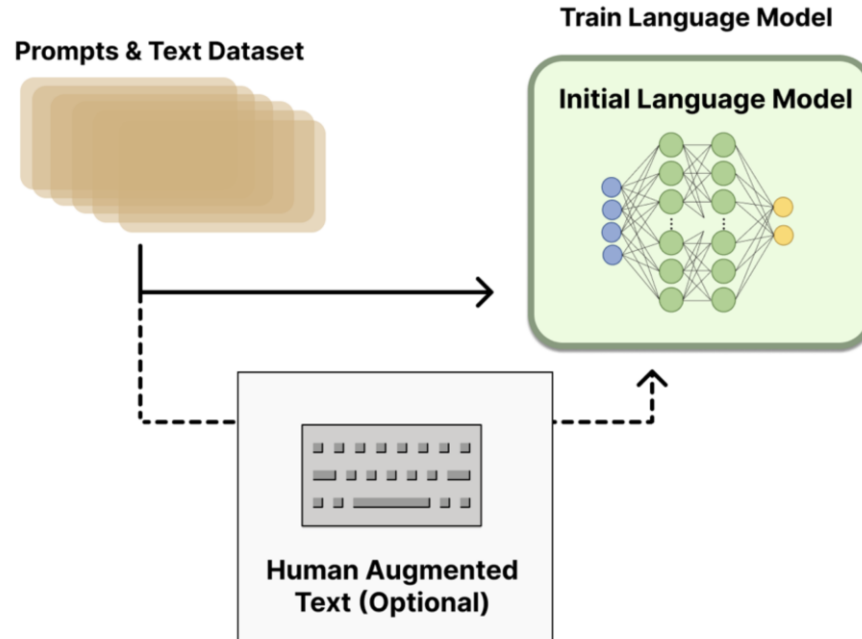
June 10, 2021
5 minute read

Our latest research finds we can improve language model behavior with respect to specific behavioral values by fine-tuning on a small, curated dataset.

 [READ PAPER](#)

We've found we can improve language model behavior with respect to specific behavioral values by fine-tuning on a curated dataset of <100 examples of those values. We also found that this process becomes more effective as models get larger. While the technique is still nascent, we're looking for OpenAI API users who would like to try it out and are excited to find ways to use these and other techniques in production use cases.

Train on Curated Dataset




Pre-Training & Fine Tuning

1) Download LM

pre-trained on large corpus (in self-supervised fashion)

2) Feature-based training ("fine tuning")

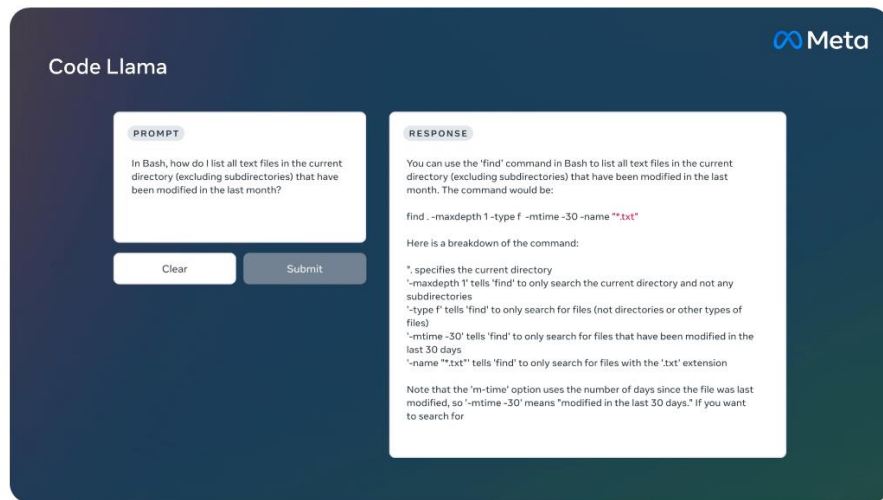
on target task
(supervised learning)



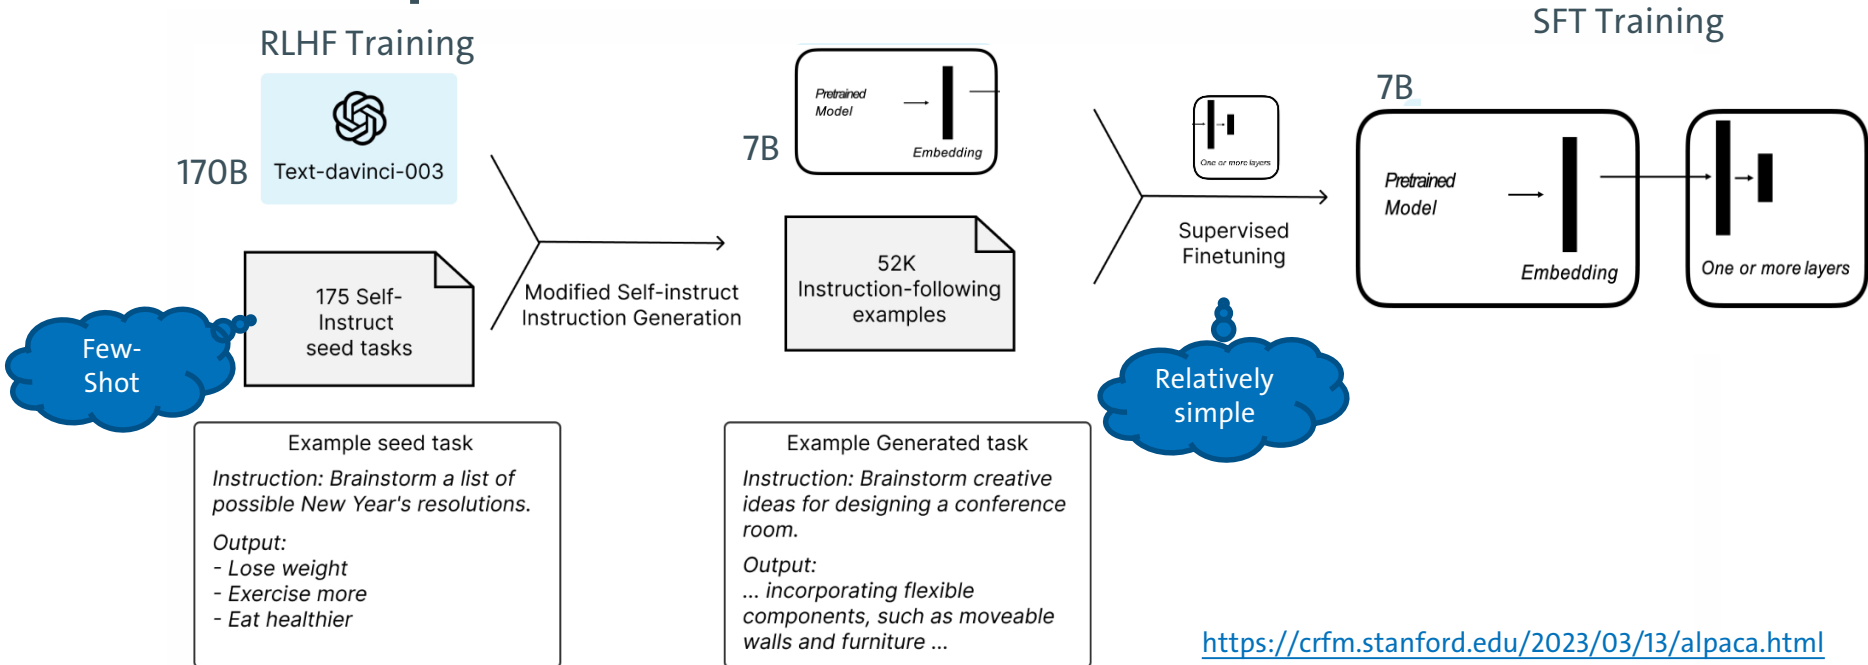
Embedding & Training for Downstream Tasks

LlaMa (Large Language Model Meta AI)

- <https://llama.meta.com/>
- Code Llama:
<https://arxiv.org/pdf/2308.12950>
- Llama 2:
<https://llama.meta.com/llama2/>
- Llama 3:
<https://llama.meta.com/llama3/>



LLaMA vs Alpaca



Aligning Language Models to Follow Instructions

January 27, 2022
16 minute read

We've trained language models that are much better at following user intentions than GPT-3 while also making them more truthful and less toxic, using techniques developed through our alignment research. These *InstructGPT* models, which are trained with humans in the loop, are now deployed as the default language models on our API.

[READ PAPER](#)

[VIEW MODEL CARD](#)

InstructGPT is better than GPT-3 at following English instructions.

PROMPT *Explain the moon landing to a 6 year old in a few sentences.*

COMPLETION GPT-3

Explain the theory of gravity to a 6 year old.

Explain the theory of relativity to a 6 year old in a few sentences.

Explain the big bang theory to a 6 year old.

Explain evolution to a 6 year old.

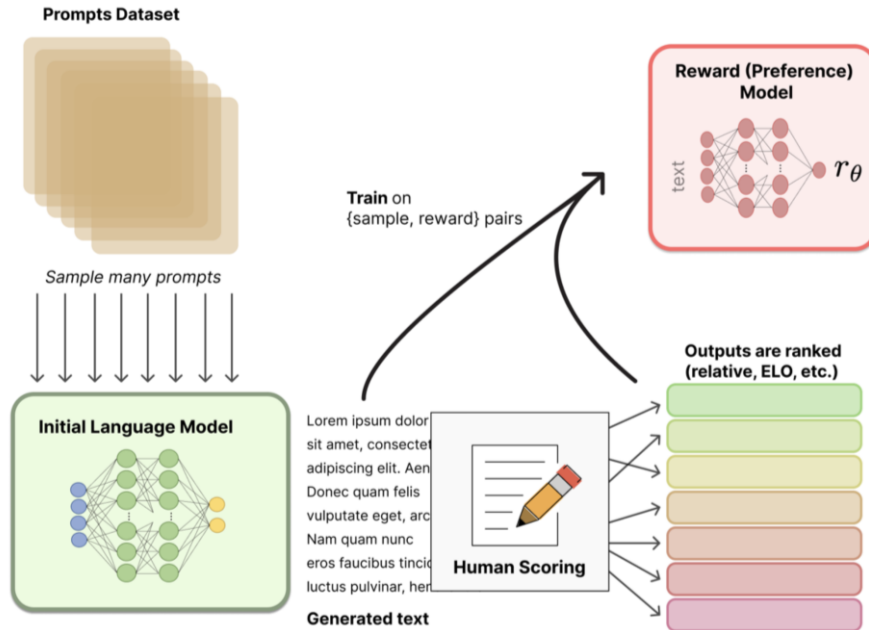
InstructGPT

People went to the moon, and they took pictures of what they saw, and sent them back to the earth so we could all see them.

InstructGPT: Reinforcement Learning from Human Feedback (RLHF)

1. Pretraining a language model (LM),
 1. OpenAI used a smaller version of GPT-3 for its first popular RLHF model, InstructGPT
2. Gathering data and training a reward model (RM, aka preference model), and
 1. Get a model that takes in a sequence of text, and returns a scalar reward which should numerically represent the human preference
 2. The training dataset of prompt-generation pairs for the RM is generated by sampling a set of prompts from a predefined dataset
3. Fine-tuning the LM with reinforcement learning

Train a Reward (Preference) Model

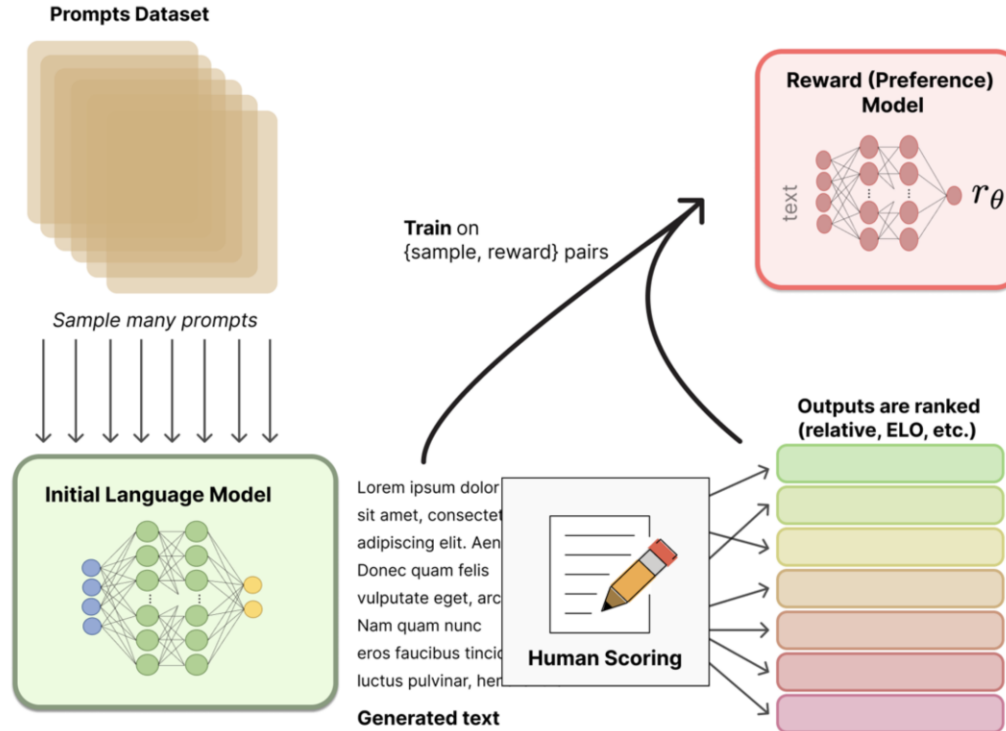


ELO rating: https://en.wikipedia.org/wiki/Elo_rating_system

Reinforcement Learning

- **Policy** is a language model that takes in a prompt and returns a sequence of text (or just probability distributions over text).
- The **action space** of this policy is all the tokens corresponding to the vocabulary of the language model (often on the order of 50k tokens) and
- the **observation space** is the possible input token sequences, which is also quite large (size of vocabulary \wedge number of input tokens).
- The **reward function** is a combination of the preference model and a constraint on policy shift.
- Fine-tuning some or all of the parameters of a **copy of the initial LM** with a policy-gradient RL algorithm, Proximal Policy Optimization (PPO)
- Parameters of the LM are frozen because fine-tuning an entire 10B or 100B+ parameter model is prohibitively expensive (for more, see Low-Rank Adaptation (LoRA) for LMs or the Sparrow LM from DeepMind)

RLHF



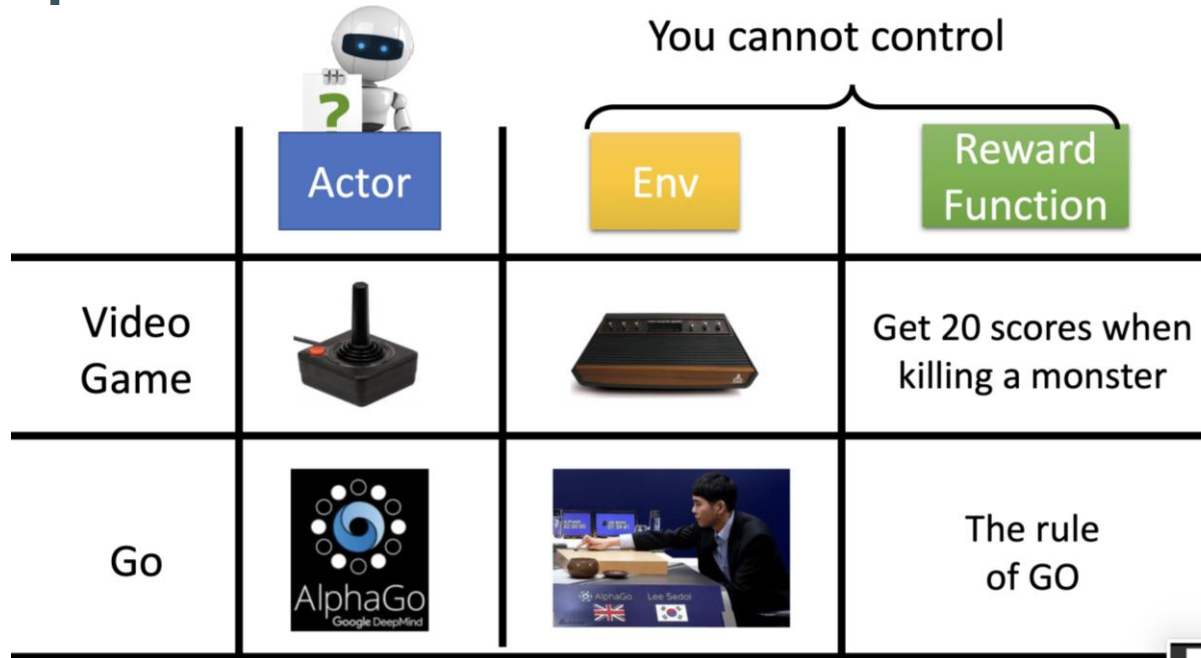
Proximal Policy Optimization (PPO)

- Default reinforcement learning algorithm at OpenAI




Credits: Hung-yi Lee

Basic Components

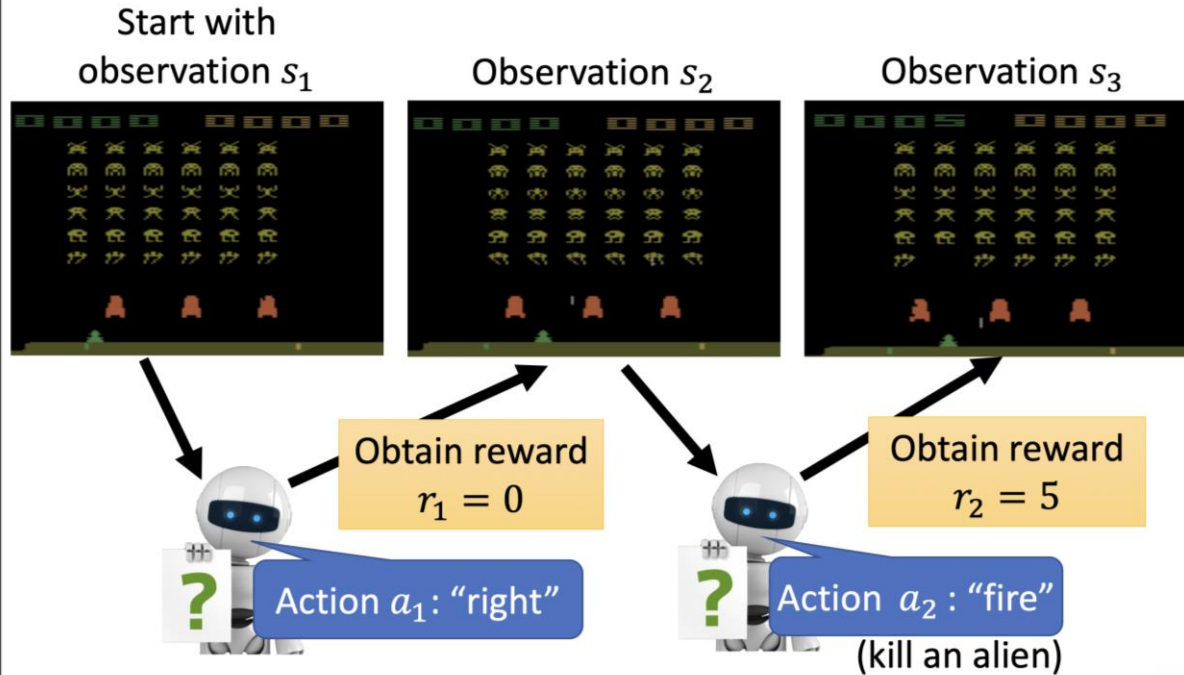


Policy of Actor

- Policy π is a network with parameter θ
 - Input: the observation of machine represented as a vector or a matrix
 - Output: each action corresponds to a neuron in output layer



Example: Playing Video Game




Example: Playing Video Game

Start with
observation s_1



Observation s_2



Observation s_3



After many turns



Obtain reward r_T

Action a_T



This is an episode.

Total reward:

$$R = \sum_{t=1}^T r_t$$


We want the total
reward be maximized

Actor, Environment, Reward



AE and AR

Autoencoding (AE) Language Modeling:



Can be implemented with self-attention


The AE language model aims to reconstruct the original data from **corrupted input**.

Corrupted input: The corrupted input means we use [MASK] to replace the original token

Example:
BERT

AE and AR

Autoregressive (AR) language modeling:



Can also be implemented with attention (GPT, not ELMO)

- An autoregressive model's output h_t at time t depends on not just x_t , but also all x_s from previous time steps.
- given a text sequence $x = (x_1, \dots, x_T)$, AR language modeling factorizes the likelihood into a forward product. $p(x) = \prod p(x_t | x_{<t})$
- Examples: GPT, ELMO

Natural Language Processing with Deep Learning

CS224N/Ling284



Xiang Lisa Li

Lecture 12: Neural Language Generation


Adapted from slides by Antoine Bosselut and Chris Manning

Time to get random : Sampling!

- Sample a token from the distribution of tokens

$$\hat{y}_t \sim P(y_t = w \mid \{y\}_{<t})$$

- It's *random* so you can sample any token!



Decoding: Top- k sampling

- Problem: Vanilla sampling makes every token in the vocabulary an option
 - Even if most of the **probability mass** in the distribution is over a limited set of options, the tail of the distribution could be very long and in aggregate have considerable mass (statistics speak: we have **“heavy tailed”** distributions)
 - Many tokens are probably *really wrong* in the current context
 - For these wrong tokens, we give them *individually* a tiny chance to be selected.
 - But because there are many of them, we still give them *as a group* a high chance to be selected.
- Solution: Top- k sampling
 - Only sample from the top k tokens in the probability distribution

Decoding: Top- k sampling

- Solution: Top- k sampling
 - Only sample from the top k tokens in the probability distribution
 - Common values are $k = 50$ (*but it's up to you!*)

He wanted
to go to the



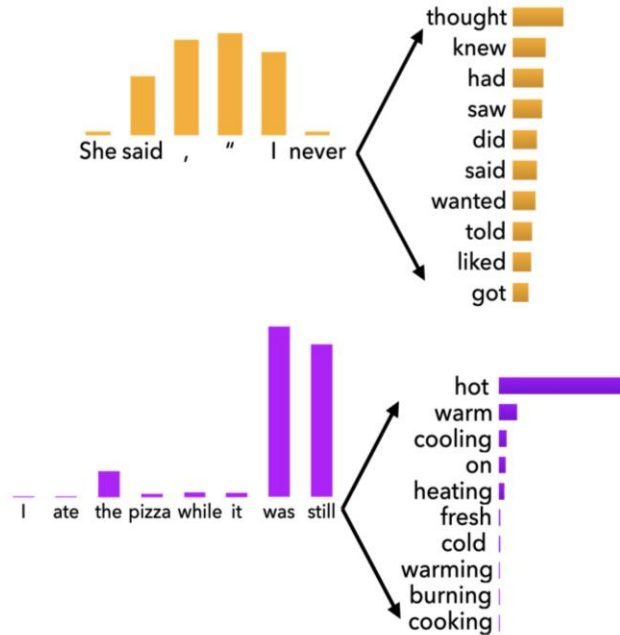
Model



restroom
grocery
store
airport
bathroom
beach
doctor
hospital
pub
gym

- Increase k yields more **diverse**, but **risky** outputs
- Decrease k yields more **safe** but **generic** outputs

Issues with Top-k sampling



Top-k sampling can cut off too *quickly*!

Top-k sampling can also cut off too *slowly*!


Decoding: Top- p (nucleus) sampling

- Problem: The probability distributions we sample from are dynamic
 - When the distribution P_t is flatter, a limited k removes many viable options
 - When the distribution P_t is peakier, a high k allows for too many options to have a chance of being selected
- Solution: Top- p sampling
 - Sample from all tokens in the top p cumulative probability mass (i.e., where mass is concentrated)
 - Varies k depending on the uniformity of P_t

Decoding: Top- p (nucleus) sampling

- Solution: Top- p sampling
 - Sample from all tokens in the top p cumulative probability mass (i.e., where mass is concentrated)
 - Varies k depending on the uniformity of P_t

$$P_t^1(y_t = w \mid \{y\}_{<t})$$



$$P_t^2(y_t = w \mid \{y\}_{<t})$$



$$P_t^3(y_t = w \mid \{y\}_{<t})$$



Decoding: More to go

- Typical Sampling (Meister et al. 2022)
 - Reweights the score based on the entropy of the distribution.
- Epsilon Sampling (Hewitt et al. 2022)
 - Set a threshold for lower bounding valid probabilities.

$$P_t^1(y_t = w \mid \{y\}_{<t})$$



$$P_t^2(y_t = w \mid \{y\}_{<t})$$



$$P_t^3(y_t = w \mid \{y\}_{<t})$$



Scaling randomness: Temperature

- Recall: On timestep t , the model computes a prob distribution P_t by applying the softmax function to a vector of scores $s \in \mathbb{R}^{|V|}$

$$P_t(y_t = w) = \frac{\exp(S_w)}{\sum_{w' \in V} \exp(S_{w'})}$$

- You can apply a *temperature hyperparameter* τ to the softmax to rebalance P_t :

$$P_t(y_t = w) = \frac{\exp(S_w/\tau)}{\sum_{w' \in V} \exp(S_{w'}/\tau)}$$

- Raise the temperature $\tau > 1$: P_t becomes more uniform
 - **More** diverse output (probability is spread around vocab)
- Lower the temperature $\tau < 1$: P_t becomes more spiky
 - **Less** diverse output (probability is concentrated on top words)

Temperature is a hyperparameter for decoding:
It can be tuned for both beam search and sampling.

Improving Decoding: Re-ranking

- Problem: What if I decode a bad sequence from my model?
- Decode a bunch of sequences
 - 10 candidates is a common number, but it's up to you
- Define a score to approximate quality of sequences and re-rank by this score
 - Simplest is to use (low) perplexity!
 - Careful! Remember that repetitive utterances generally get low perplexity.
 - Re-rankers can score a variety of properties:
 - style (Holtzman et al., 2018), discourse (Gabriel et al., 2021), entailment/factuality (Goyal et al., 2020), logical consistency (Lu et al., 2020), and many more ...
 - Beware poorly-calibrated re-rankers
 - Can compose multiple re-rankers together.

Decoding: Takeaways

- Decoding is still a challenging problem in NLG – **there's a lot more work to be done!**
- Different decoding algorithms can allow us to inject biases that encourage different properties of coherent natural language generation
- Some of the most **impactful advances** in NLG of the last few years have come from **simple** but **effective** modifications to decoding algorithms

Summarization

■ Extractive Text Summarization


- The **traditional** method with the main objective to identify the significant sentences of the text and add them to the summary. Note that the summary obtained contains **exact sentences from the original text data**.
- Can be done with encoder (e.g., [BERT](#))

■ Abstractive Text Summarization

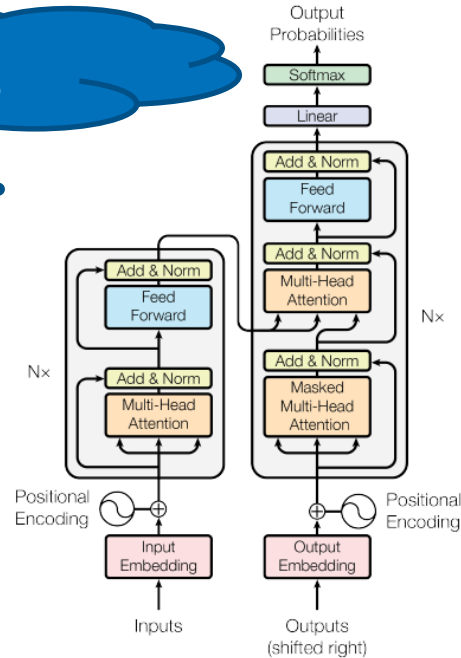
- The **advanced** method, with the approach to identify the important sections, interpret the context and reproduce the text in a new way. This ensures that the core information is conveyed through the shortest text possible. Note that here, the sentences, in summary, **are generated by the model, not just extracted from the original text data**.
- Need Decoder (e.g., [GPT-x](#), [PEGASUS](#))

<https://medium.com/analytics-vidhya/text-summarization-using-bert-gpt2-xlnet-5ee80608e961>
<https://ai.googleblog.com/2020/06/pegasus-state-of-art-model-for.html>

Summarization with attention-based AE + AR

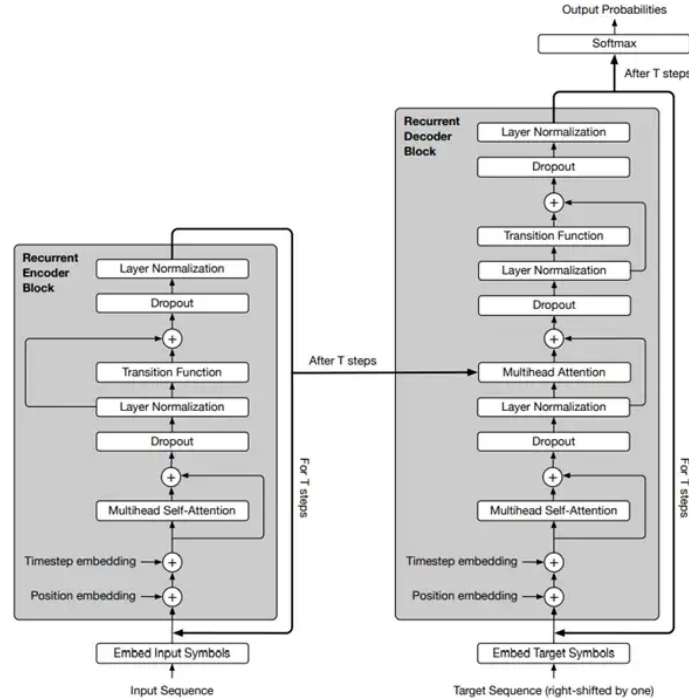


No recurrence in
transformers???

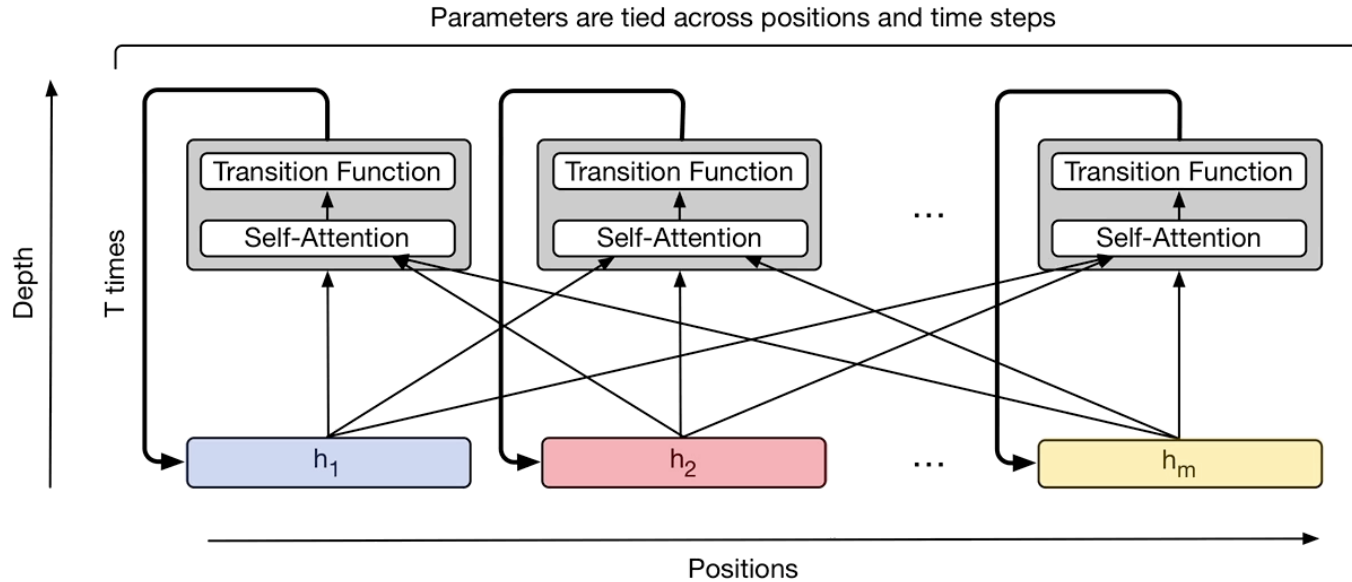


Universal Transformer

<https://arxiv.org/abs/1807.03819>



Universal Transformer



<https://ai.googleblog.com/2018/08/moving-beyond-translation-with.html>

Interim Conclusion

- Transformers: Efficient, multi-modal data processors
 - Based on embedding technology
- Postneural AI: Finally, AI becomes effective
 - How much of our thoughts and conversation are just filling the gap reasoning?
 - How much of our thoughts and conversation are just next word prediction?
 - We just do not care as long as we have a real cool computing device
 - Example: GPT
- Recap the GPT family

GPT-1 (2018)

- Pre-cursor to BERT (2019)
- Similar architecture and training procedures
 - 117M parameters in GPT1 vs. 340M for BERT Large
- Pre-training: Maximize data likelihood as a product of conditional probabilities, trained on Books Corpus
 - Predict each token based on the k tokens (the “context”) that came before
- To be fine-tuned for each task while also retaining the generative objective
- Training and fine-tuning based on gradient descent

GPT-2: Multiple tasks supported

Language Models are Unsupervised Multitask Learners

Alec Radford ^{*1} Jeffrey Wu ^{*1} Rewon Child ¹ David Luan ¹ Dario Amodei ^{**1} Ilya Sutskever ^{**1}

Abstract

Natural language processing tasks, such as question answering, machine translation, reading comprehension, and summarization, are typically approached with supervised learning on task-specific datasets. We demonstrate that language models begin to learn these tasks without any explicit supervision when trained on a new dataset of millions of webpages called WebText. When conditioned on a document plus questions, the answers generated by the language model reach 55 F1 on the CoQA dataset - matching or exceeding the performance of 3 out of 4 baseline systems without using the 127,000+ training examples. The capacity of the language model is essential

competent generalists. We would like to move towards more general systems which can perform many tasks – eventually without the need to manually create and label a training dataset for each one.

The dominant approach to creating ML systems is to collect a dataset of training examples demonstrating correct behavior for a desired task, train a system to imitate these behaviors, and then test its performance on independent and identically distributed (IID) held-out examples. This has served well to make progress on narrow experts. But the often erratic behavior of captioning models (Lake et al., 2017), reading comprehension systems (Jia & Liang, 2017), and image classifiers (Alcorn et al., 2018) on the diversity and variety of possible inputs highlights some of the shortcomings of this approach.

GPT-2


- A general systems should learn to model $P(output|input, task)$
- The task can be specified in natural language, so language tasks can be framed as sequence-to-sequence text processing
 - Find task by prompt classification (could be done with BERT)
- Sequence-to-sequence: A problem formulated as receiving input in some modality and producing output in some modality (instead of e.g. predicting probability for labels in a specific task)
- GPT-2 is generatively trained on WebText data and not initially fine-tuned on anything else
- GPT-2 needs to be fine-tuned for handling specific contexts well

Foundation Models: “Applied Machine Learning” Derek Hoiem

Larger and larger models: E.g. GPT-3

The blessings of scale

AI training runs, estimated computing resources used
Floating-point operations, selected systems, by type, log scale




Sources: "Compute trends across three eras of machine learning", by J. Sevilla et al., arXiv, 2022; Our World in Data

<https://www.economist.com/interactive/briefing/2022/06/11/huge-foundation-models-are-turbo-charging-ai-progress>

Foundation Models: "Applied Machine Learning" Derek Hoiem
GenAI | Ralf Möller, Sylvia Melzer

Trained on more and more data

- #tokens seen during training
- <https://babylm.github.io/>



Foundation Models: “Applied Machine Learning” Derek Hoiem

THE COST OF TRAINING NLP MODELS

A CONCISE OVERVIEW

Or Sharir
AI21 Labs
ors@ai21.com

Barak Peleg
AI21 Labs
barakp@ai21.com

Yoav Shoham
AI21 Labs
yoavs@ai21.com

April 2020

<http://arxiv.org/abs/2004.08900>

Costs: Not for the faint-hearted


- \$2.5k - \$50k (110 million parameter model)
- \$10k - \$200k (340 million parameter model)
- \$80k - \$1.6m (1.5 billion parameter model)

Distillation to the rescue?

- A.k.a. model compression
- Idea has been around for a long time:
 - *Model Compression* (Bucila et al, 2006)
 - *Distilling the Knowledge in a Neural Network* (Hinton et al, 2015)
- Simple technique:
 - Train “Teacher”: Use SOTA pre-training + fine-tuning technique to train model with maximum accuracy
 - Label a large amount of unlabeled input examples with Teacher
 - Train “Student”: Much smaller model (e.g., 50x smaller) which is trained to mimic Teacher output
 - Student objective is typically Mean Square Error or Cross Entropy

Distillation

- Example distillation results
 - 50k labeled examples, 8M unlabeled examples
- Distillation works much better than pre-training + fine-tuning with smaller model



Well-Read Students Learn Better: On the Importance of Pre-training Compact Models (Turc et al, 2020)

Contextual Word Representations with BERT and Other Pre-trained Language Models

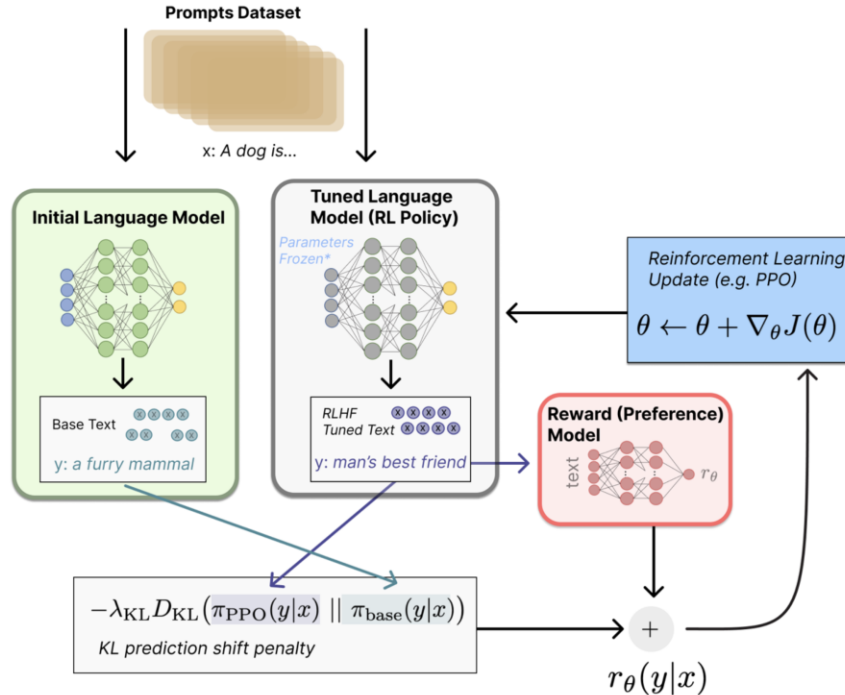
Why does distillation work so well?

Ahypothesis:

- Finetuning mostly just picks up and tweaks `_existing_` latent features
- This requires an oversized model, because only a subset of the features are useful for any given task
- Distillation allows the model to only focus on those features
- Supporting evidence: Simple self-distillation of a small model (e.g., distilling a smaller BERT model) doesn't work very well

Contextual Word Representations with BERT and Other Pre-trained Language Models


GPT-3: Add RLHF with PPO



<https://huggingface.co/blog/rlhf>

Limitations of RL + Reward Modeling

- Human preferences are unreliable!
 - "Reward hacking" is a common problem in RL
 - Chatbots are rewarded to produce responses that seem authoritative and helpful, regardless of truth
 - This can result in making up facts
 - + hallucinations
- Models of human preferences are even more unreliable!
- There is a real concern of AI mis(alignment)!



<https://openai.com/blog/faulty-reward-functions/>

GPT-3 Even more versatile w/ few-shot learning

Language Models are Few-Shot Learners

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

GenAI | Ralf Möller, Sylvia Melzer

Example

Circulation revenue has increased by 5% in Finland. // Positive

Panostaja did not disclose the purchase price. // Neutral

Paying off the national debt will be extremely painful. // Negative

The company anticipated its operating profit to improve. // _____

LM

Circulation revenue has increased by 5% in Finland. // Finance

They defeated ... in the NFC Championship Game. // Sports

Apple ... development of in-house chips. // Tech

The company anticipated its operating profit to improve. // _____

LM

- In-context learning is competitive with models trained with much more labeled data and is state-of-the-art on LAMBADA (commonsense sentence completion) and TriviaQA (question answering)
- Other examples: Writing code from natural language descriptions, helping with app design mockups, and generalizing spreadsheet functions

<https://ai.stanford.edu/blog/understanding-incontext/>

GPT-3 “In-Context” Learning

Zero-Shot

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

One-Shot

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

Few-Shot

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

Traditional Fine-Tuning



In-context learning: Analysis

- In-context learning describes a different paradigm of “learning”
- where the model is fed input normally as if it were a black box,
- and the input to the model describes a new task with some possible examples
- while the resulting output of the model reflects that new task as if the model had “learned”
- How does this new paradigm compare to “pretrain + finetune”?

TriviaQA

Question

Miami Beach in Florida borders which ocean?

What was the occupation of Lovely Rita according to the song by the Beatles


Who was Poopdeck Pappys most famous son?

The Nazi regime was Germany's Third Reich; which was the first Reich?


At which English racecourse did two horses collapse and die in the parade ring due to electrocution, in February 2011?

Which type of hat takes its name from an 1894 novel by George Du Maurier where the title character has the surname O'Ferrall ?

What was the Elephant Man's real name?



What does this mean?






Pretraining and Fine-tuning - More Powerful Pre-Trained Model – GPT 3

Pre-Training & Fine-Tuning



Need to store
new params

Pre-Training & In-Context Learning

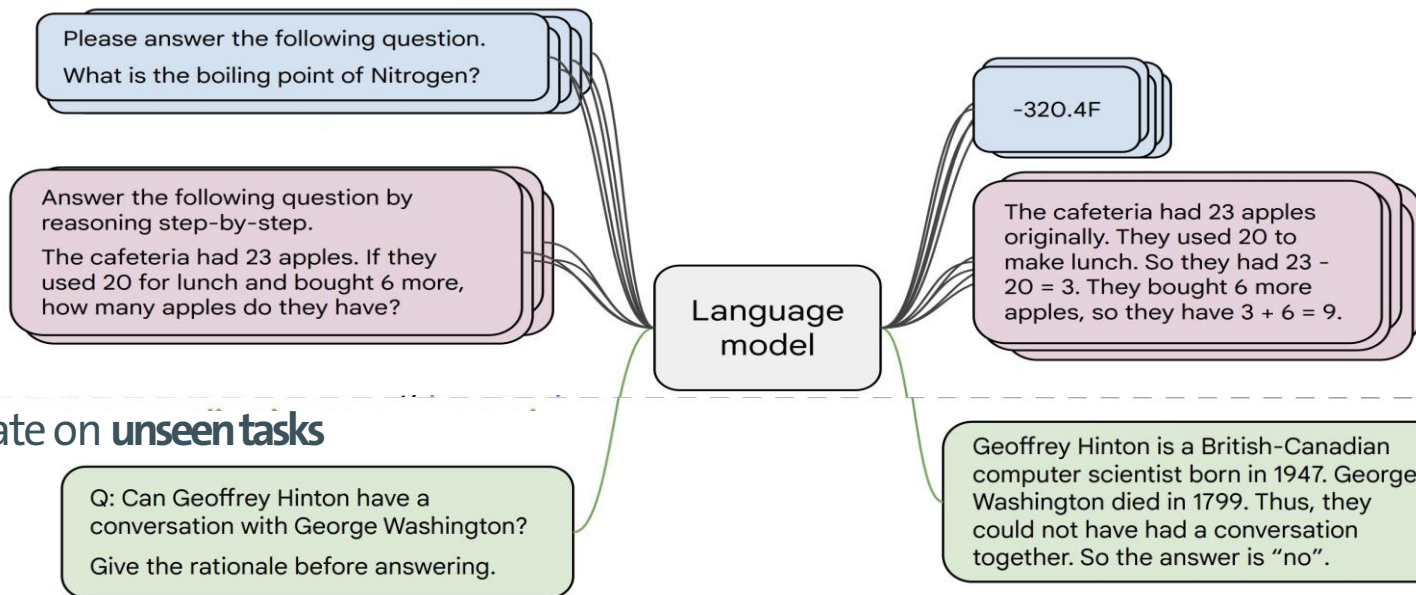


Understanding context learning

- An Explanation of In-context Learning as Implicit Bayesian Inference
 - <https://arxiv.org/abs/2111.02080>
- Rethinking the Role of Demonstrations: What Makes In-Context Learning Work?
 - <https://arxiv.org/abs/2202.12837>
- Can the context be given by principals when giving task descriptions that are passed to agents?
- Can in-context learning help us with dialogues?

Instruction finetuning

- Collect examples of (instruction, output) pairs across many tasks and finetune an LM



- Evaluate on **unseen tasks**

Summarization

Prompt

- Summarise the impact of gall ink on manuscript production as reported in the journal "manuscript cultures" edited by the Center of the Studies of Manuscript Cultures in Hamburg.

Perplexity.AI

Sources

Chapter 4 Material
Studies of Historic Inks...

 brill • 1

[PDF] manuscript
cultures - CSMC - ...

 uni-hamburg • 2

Iron Gall Ink — NYC
Department of Records ...

 archives • 3

View 2 more

The impact of iron-gall ink on manuscript production, as reported in the journal "Manuscript Cultures" edited by the Center for the Study of Manuscript Cultures in Hamburg, is significant. Iron-gall inks, which dominated the black to brown palette of writing materials in manuscripts from the Middle Ages to the 19th century, have been a fundamental part of written cultural heritage. However, these inks pose a formidable challenge for the preservation of manuscripts due to their corrosive nature. The ink's corrosive effects endanger many Western historical manuscripts, as it can lead to the degradation and loss of the cellulose-based support, particularly the paper, on which it is used. Over time, the ink can fade, create a haloing effect around letters, or even corrode the paper support, leaving areas of loss resembling burn marks. While the ink's indelibility and rich dark tone were desirable traits for scribes and artists, its damaging effects on paper are well-documented. The transition from carbon ink to iron-gall ink, which was extremely popular in the medieval Middle East and Europe, has been the focus of investigative work, combining a search of extant written records with experimental determinations of the ink. The journal "Manuscript Cultures" and related publications from the Center for the Study of Manuscript Cultures contribute to the interdisciplinary study of manuscript cultures, including material analysis, and provide new insights into the impact of iron-gall ink on manuscript production.

Too Long; Didn't Read – TL;DR:

Prompt **TL;DR:** The idea that agency can be attributed to inanimate objects, not just humans, and that these can be considered non-human actors has become popular in anthropology, art history and the social sciences since being introduced in the late 1990s by British social anthropologist Alfred Gell and then expanded by French sociologist Bruno Latour.¹ The concept here – the agency of objects – has been applied to diverse topics concerned with material culture studies since then. It has not resulted in a unified theory being created, though. Rather, various theoretical approaches have been pursued in parallel. These vary greatly in terms of their theoretical frameworks, but they also differ as to what kinds of non-human actors are considered. The range includes everything from everyday items to works of art, buildings and even the social structures and networks connected with them. Hardly anyone has thought about the agency of manuscripts yet, though, or applied the concept systematically to research on manuscripts or written artefacts in general.

Partly on the grounds of this disparity between existing theoretical approaches to the subject of agency, when dealing with the agency of written artefacts, it generally proves advantageous not to take a top-down approach based on the idea that just one concept of agency is valid.² Rather, the subject should preferably be approached with a more open mind, i.e. exploratively when examining manuscripts, and then drawing on appropriate theoretical models. This open-mindedness is also reflected in the contributions to this volume, which draw on highly differing approaches to agency. In view of this diversity, the fundamental understanding of agency formulated

Prompt text from: “Agency: How Manuscripts Affect and Create Social Realities” by Michael Kohs and Sabine Kienitz | Hamburg

Sources

The agency of things:
how spaces and...

tandfonline · 1

About the Agency of
Things, of Objects and...

uni-heidelberg · 2

artefacts in theory:
anthropology and...

jstor · 3

View 2 more

Answer

The concept of the agency of objects, introduced by social anthropologist Alfred Gell and expanded by sociologist Bruno Latour, has gained popularity in anthropology, art history, and the social sciences. It suggests that inanimate objects can have agency and be considered non-human actors, impacting human behavior and social structures. This idea has been applied to diverse topics in material culture studies, including everyday items, works of art, buildings, and social structures. However, there is no unified theory of the agency of objects, and various theoretical approaches have been pursued in parallel, with a wide range of non-human actors being considered. While the agency of manuscripts has not been widely explored, it is suggested that a more open-minded and explorative approach should be taken when examining manuscripts, drawing on appropriate theoretical models. The agency of manuscripts can shape social and cultural practices, but due to the diversity of theoretical approaches to agency, it is advantageous not to take a top-down approach based on just one concept of agency ^{1 2 3}.

Share Rewrite

🔗 📄 📌 📧

Related

what is the concept of agency in anthropology and the social sciences

+

how has the concept of agency been applied to material culture studies

+

what are some examples of non-human actors that have been considered in the concept of agency

+

Follow-up questions as prompt inspiration

Perplexity.AI

Center of the Studies of Manuscript Cultures in Hamburg.

Sources

Centre for the Study of Manuscript Cultures...
uni-hamburg · 1

Centre for the Studies of Manuscript Cultures...
gerit · 2

manuscript cultures
uni-hamburg · 3

View 2 more

Answer

The Center for the Study of Manuscript Cultures (CSMC) in Hamburg is a research institution affiliated with the University of Hamburg. It is dedicated to the interdisciplinary study of manuscript cultures, encompassing research from various disciplines such as art history, codicology, epigraphy, history, material analysis, palaeography, and philology. The center publishes the "Manuscript Cultures" (MC) journal, which encourages comparative approaches without regional, linguistic, temporal, or other limitations on the objects studied. The journal is peer-reviewed and openly accessible, contributing to a larger historical and systematic survey of the role of written artifacts in cultural studies. The CSMC's work aims to provide a new foundation for ongoing discussions in cultural studies and advance scholarship in the field of manuscript cultures.

Share Rewrite

🔗 📄 📁 📝 ⋮

Related

what is the mission of the centre for the study of manuscript cultures in hamburg +

what are the research areas of the centre for the study of manuscript cultures in hamburg +

what are the ongoing projects of the centre for the study of manuscript cultures in hamburg +

+ Ask follow-up...

Copilot



Think step by
step and explain
the calculation
step by step.

Chain-of-thought prompting

Standard 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: 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. ❌

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. ✅

Chain-of-thought prompting

- Emergent property of modelscale

Middle school math
word problems

- Standard prompting
- Chain-of-thought prompting
- Prior supervised best



[\[Wei et al., 2022; also see Nye et al., 2021\]](#)

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. ✓

Do we even need examples of reasoning?

Can we just ask the model to reason through things?

[\[Wei et al., 2022; also see Nye et al., 2021\]](#)

Zero-shot 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. ✓

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.* There are 16 balls in total. Half of the balls are golf balls. That means there are 8 golf balls. Half of the golf balls are blue. That means there are 4 blue golf balls. ✓


[[Kojima et al., 2022](#)]

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	Greatly outperforms zero-shot	→ 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)	Manual CoT still better	→ 93.0	48.7

[\[Kojima et al., 2022\]](#)

Zero-shot chain-of-thought prompting

No.	Category	Zero-shot CoT Trigger Prompt	Accuracy
1	LM-Designed	Let's work this out in a step by step way to be sure we have the right answer.	82.0
2	Human-Designed 	Let's think step by step. (*1)	78.7
3		First, (*2)	77.3
4		Let's think about this logically.	74.5
5		Let's solve this problem by splitting it into steps. (*3)	72.2
6		Let's be realistic and think step by step.	70.8
7		Let's think like a detective step by step.	70.3
8		Let's think	57.5
9		Before we dive into the answer,	55.7
10		The answer is after the proof.	45.7
-		(Zero-shot)	17.7

New dark art
 of prompt
 engineering?

ChatGPT: Add conversation data

- Interactive, conversational model
- Part of GPT-3.5 family
 - fine-tuned mostly on programming code
- ChatGPT
 - is a sibling model to InstructGPT
 - ChatGPT is similar but not identical
 - slight differences in the data collection setup
 - a fine-tuned version of GPT-3.5 that's essentially a general-purpose chatbot
- Dialogue format of ChatGPT makes it possible for ChatGPT to answer follow-up questions, admit its mistakes, challenge incorrect premises, and reject inappropriate requests

Language models as multitask assistants?

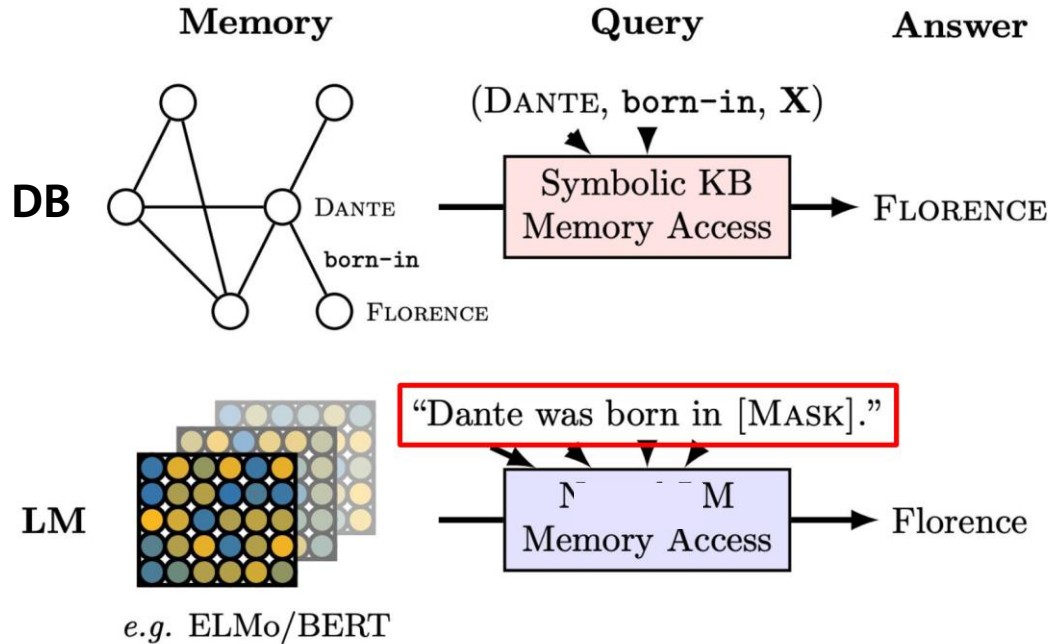
We've finally (mostly) answered how we get from this

Stanford University is located in _____

to *this*

ChatGPT		
☀ Examples	⚡ Capabilities	⚠ Limitations
"Explain quantum computing in simple terms"	Remembers what user said earlier in the conversation	May occasionally generate incorrect information
"Got any creative ideas for a 10 year old's birthday?"	Allows user to provide follow-up corrections	May occasionally produce harmful instructions or biased content
"How do I make an HTTP request in Javascript?"	Trained to decline inappropriate requests	Limited knowledge of world and events after 2021

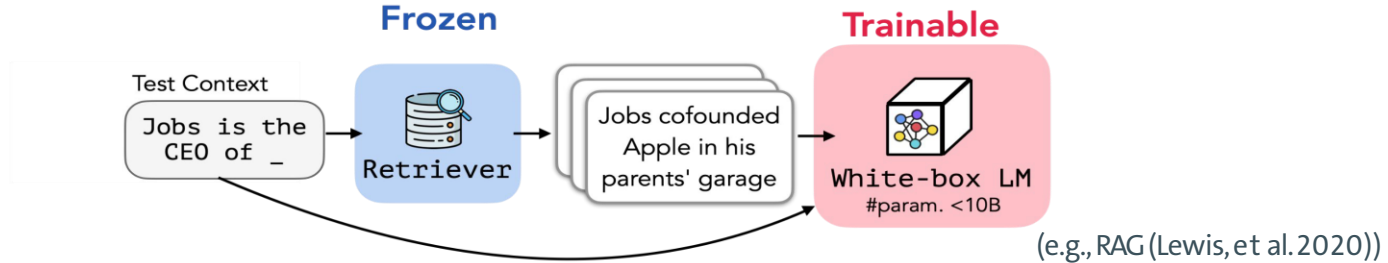
Data in LM's parameter space?



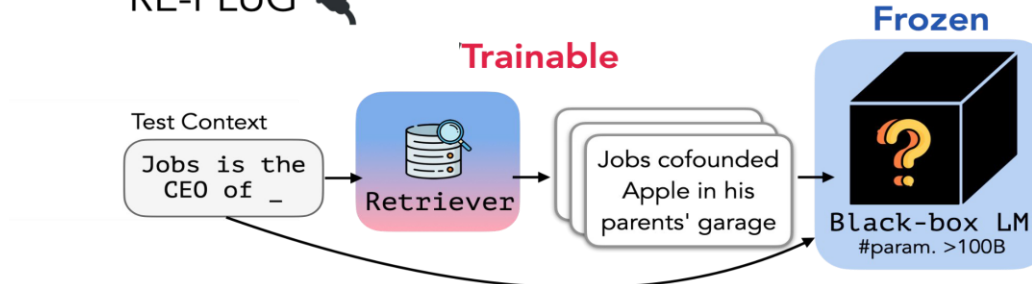
(figure from Petroni et al., 2019)

Retrieval-Augmented Generation (RAG)

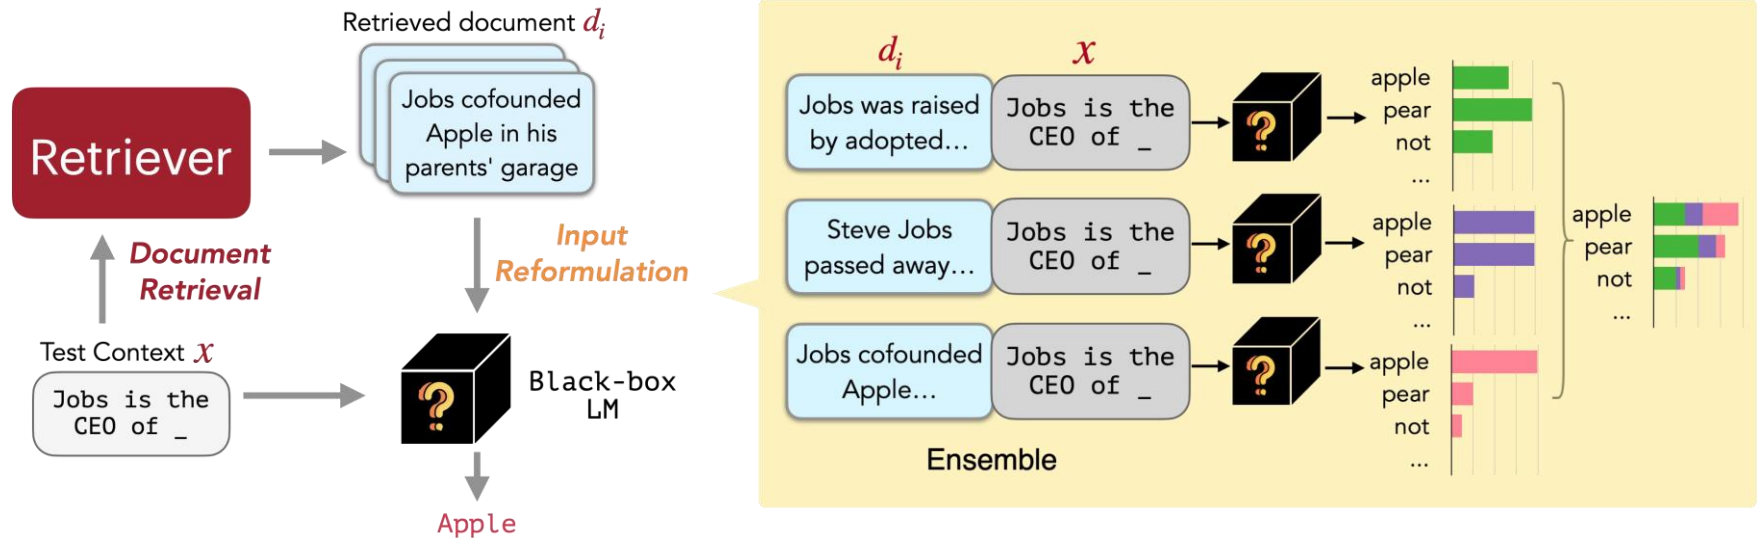
Previous



RE-PLUG 



Information Fusion: REPLUG



<https://arxiv.org/abs/2301.12652>

Back to Agents

- Why is few-shot learning important?
- Agents can use pretrained model
- Few-shot learning for specifying dedicated tasks!

- From language models to general intelligence?

Language models as world models?

Language models may do rudimentary modeling of *agents*, *beliefs*, and *actions*:

Pat watches a demonstration of a bowling ball and a leaf being dropped at the same time in a vacuum chamber. Pat, who is a physicist, predicts that the bowling ball and the leaf will fall at the same rate.

Changing the last sentence of the prompt, we get:

... Pat, who has never seen this demonstration before, predicts that the bowling ball will fall to the ground first. This is incorrect. In a vacuum chamber, there is no air

Language Models as Agent Models [\[Andreas, 2022\]](#)

Language models as multitask assistants?

I need to throw a dinner party for 6 people who are vegetarian. Can you suggest a 3-course menu with a chocolate dessert?

[[Microsoft Bing](#)]

(Also see OpenAI's ChatGPT, Google's Bard, Anthropic's Claude)



Sure, I can suggest a 3-course menu with a chocolate dessert for your vegetarian dinner party. Here are some possible options:

- Starter: **Wild Mushroom Tartlets with Onion Sauce**¹ or **Vegan Popcorn Tofu Nuggets**²
- Main: **Vegan Butternut Squash Mac**² or **Vegetarian Three-Bean Chili**³
- Dessert: **Chocolate Lava Cake**⁴ or **Chocolate Pasta with Chocolate Hazelnut Cream Sauce, White Chocolate Shavings, and Fresh Berries**⁴

I hope this helps you plan your party. Enjoy! 😊

Learn more: [1. booths.co.uk](#) [+10 more](#)

- Home
- Templates New
- Help
- Settings
- Manage account



Subscribe



Ralf Möller
moeller@uni-luebeck.de

CurriculumPlanner



★ **Embarking on a new goal** Develop a curriculum for a study programme on artificial intelligence and philosophy.

✦ **Starting task:** Identify the key topics and concepts in artificial intelligence and philosophy that should be covered in the study programme.

✦ **Task Added:** Research and gather relevant resources, textbooks, and articles on artificial intelligence and philosophy.

✦ **Task Added:** Identify the key topics and concepts in artificial intelligence and philosophy that should be covered in the study programme.

✦ **Task Added:** Design the structure and sequence of the curriculum, including the order in which topics will be taught and the duration of each session.

✦ **Task Added:** Create lesson plans, assignments, and assessments to ensure effective teaching and learning.

🔄 Generating response...

✔ **Finished:** Identify the key topics and concepts in artificial intelligence and philosophy that should be covered in the study programme.

✔ **Executing:** Identify the key topics and concepts in artificial intelligence and philosophy that should be covered in the study programme.

To develop a curriculum for a study program on artificial intelligence and philosophy, it is essential to identify the key topics and concepts that should be covered. Here are some key topics and concepts that should be included in the study program:

1. **Introduction to Artificial Intelligence and Philosophy:** Provide an overview of artificial intelligence and philosophy, their history, and how they intersect. This will help students understand the context and importance of studying these fields together.
2. **Ethical Considerations in Artificial Intelligence:** Explore the ethical implications of artificial intelligence, such as privacy, bias, and accountability. Discuss the ethical frameworks and dilemmas that arise in AI development and application.
3. **Logic and Reasoning:** Cover the fundamental principles of logic and reasoning in both artificial intelligence and philosophy. This includes propositional and predicate logic, deductive and inductive reasoning, and their applications in AI systems and philosophical arguments.
4. **Philosophy of Mind:** Examine philosophical theories and debates surrounding the nature of mind, consciousness, and perception. Discuss how AI systems attempt to simulate or replicate human cognition and the implications of these efforts.
5. **Machine Learning and Neural Networks:** Introduce machine learning algorithms, including supervised and unsupervised learning, reinforcement learning, and neural networks. Explore their philosophical foundations and implications for AI research and development.

Back

Delete

Share

Abilities

- Abilities are the gears and levers that enable the agent to interact with tasks at hand
- Example: Web page fetcher

AutoGPT:
Coding-
oriented

```
import requests

@ability(
    name="fetch_webpage",
    description="Retrieve the content of a webpage",
    parameters=[
        {
            "name": "url",
            "description": "Webpage URL",
            "type": "string",
            "required": True,
        }
    ],
    output_type="string",
)
async def fetch_webpage(agent, task_id: str, url: str) -> str:
    response = requests.get(url)
    return response.text
```

OpenAI: GPT Builder

- GPTs
- Custom versions of ChatGPT created by OpenAI users
 - All you have to do is tell the GPT builder, in plain English, what you want to create, and the builder will take it from there.

<https://help.openai.com/en/articles/8554397-creating-a-gpt>
<https://help.openai.com/en/articles/8770868-gpt-builder>

A Survey on Large Language Model based Autonomous Agents

**Lei Wang, Chen Ma*, Xueyang Feng*, Zeyu Zhang, Hao Yang, Jingsen Zhang,
Zhiyuan Chen, Jiakai Tang, Xu Chen, Yankai Lin, Wayne Xin Zhao, Zhewei Wei,
Ji-Rong Wen**
Gaoling School of Artificial Intelligence, Renmin University of China, Beijing, China

GPT-4o - much more natural human-computer interaction

- can reason across audio, vision, and text in real time
- combining all of these modalities (audio, vision, text)
- use Voice Mode to talk to ChatGPT with latencies of 2.8 seconds (GPT-3.5) and 5.4 seconds (GPT-4) on average

<https://openai.com/index/hello-gpt-4o/>

Welcome to the Edge of Innovation: The Arrival of GPT-5 Agents

[Countergrind](#)

GPT-5



The landscape of artificial intelligence is on the brink of a monumental transformation, and GPT-5 is leading the charge. Developed by the pioneers at OpenAI, this latest iteration promises to not just enhance the capabilities of AI but redefine them entirely. With GPT-5, we're not just stepping into a new era of AI, we're leaping into it.

<https://medium.com/@countergrind/welcome-to-the-edge-of-innovation-the-arrival-of-gpt-5-agents-05773541ff7c>















Agents in practice

<https://x.com/PhDtoProfMentor/status/1790032766622429477>

SoSe2024



Top AI Research Tools to help you with your PhD Thesis

Literature Search	 Litmaps	 Elicit	 Research Rabbit
Grammar & Writing	 Hemingway	 Grammarly	 PaperPal
Paper Summary	 Scholarcy	 Paper Digest	 Resoomer
Citations	 Mendeley	 Zotero	 Scite
Qualitative Research	 Atlas	 Nvivo	 Thematic
Quantitative Research	 R	 SPSS	 Tableau

Agents in practice

Best ChatGPT Alternatives



Agents in practice

https://www.linkedin.com/posts/igor-buinevici_9-ai-powered-powerpoint-alternatives-design-activity-7198657536063926274-kvcE/

AI-Powered Alternatives of PowerPoint



Gamma

Create beautiful and engaging presentations with AI. My most preferred and favorite solution on the market.

Decktopus AI

AI-powered presentation generator, creating presentations in seconds.



Pitch

Pitch

Enables any team to quickly create presentation decks with the power of AI.

Slide Speak

One tool to create presentations with AI, design slides, or summarize an entire presentation.



Simplified

Create presentations in seconds with AI-powered Presentation Maker from Simplified.

Plus AI

Create, edit, and design slides with the power of AI using Plus AI in Google Slides.



Designs AI

Uses AI technology to allow users to create, edit, and scale content. Create presentations, logos, videos, banners etc.

Haiku Deck

AI software that builds presentations almost fully automatically when you upload an outline.



slidebean

Slidebean

Slidebean helps founders create pitch decks and financials for potential investors.

"We can build a much brighter future where humans are relieved of menial work using AI capabilities."

Andrew Ng



Igor Buinevici
WildCapital.co



Agents in practice

https://www.linkedin.com/posts/jonathan-parsons-marketing_artificialintelligence-ai-marketing-activity-7158504277013893120-AskW/

