

Exploring the mechanisms underlying LLMs like ChatGPT

Lisa Meeden
Professor of Computer Science
Swarthmore College

Outline of the talk

- Discuss a word game, try it ourselves, and test how well ChatGPT can play it
- Learn how [word embeddings](#) are created to represent text for LLMs
- Learn how LLMs generate text through [transformers](#) and [self-attention](#)
- Reflect on LLMs

How to play NYT game Connections

Find groups of four items that share something in common

- Pick four items, and check if your guess is correct
- Try to find all four groups without making four mistakes

Category examples:

- FISH: Bass, Flounder, Salmon, Trout
- FIRE _____: Ant, Drill, Island, Opal

Each puzzle has exactly one solution. Watch out for words that seem to belong to multiple categories.

Connections

October 17, 2023

Create four groups of four!

GHOST

LEMON

GARDEN

DESERT

STAR

RASPBERRY

CANDY

BOMB

FLOP

HISS

IGNORE

BOO

BOTTOM

DUD

JEER

JILT

RUDELY BREAK OFF CONTACT

DESERT, GHOST, IGNORE, JILT

TAUNTS

BOO, HISS, JEER, RASPBERRY

FAILURE

BOMB, DUD, FLOP, LEMON

ROCK ___

BOTTOM, CANDY, GARDEN, STAR

Let's see how OpenAI's ChatGPT (3.5) handles this game

LM

Let's play a game called connections. I will give you 16 words. Your goal is to sort these words into 4 groups of 4 by finding the connections between them. You can give me one group of 4 at a time and I will tell you if you are correct. Here are the words: ghost, lemon, garden, desert, star, raspberry, candy, bomb, flop, hiss, ignore, boo, bottom, dud, jeer, jilt.



Certainly, let's play "Connections" with these words. I'll start by proposing a group of 4 words:   


Group 1: ghost, boo, hiss, jeer These 3 words are part of the "Taunts" group

Do these words form a correct group, or should I try a different combination?

LM

This is close. Only one word is incorrect.



Thank you for the feedback. Let's revise the group:   

Group 1: ghost, boo, hiss, jeer Note that these are the same 4 words as above

Provide a hint



For group 1, try replacing the word ghost with something else.



I see, let's modify Group 1:



Group 1: boo, hiss, jeer, flop

Now, for the next group:

Group 2: lemon, garden, raspberry, candy

Do these words form a correct group, or should I make further adjustments?



No, flop is not right for group 1. Try group 1 again.



I apologize for the error. Let's try again for Group 1:



Group 1: boo, hiss, jeer, jilt

Found the right concept for the group



This still is one word away from being correct for group 1. Can you describe why you think the words boo, hiss, jeer, and jilt are related?



My previous attempts at grouping were based on a connection I perceived among those words, which was related to negative or disapproving vocal expressions or actions. However, considering your hint and reassessing the words, I now see a different connection:



Group 1: ghost, desert, star, bomb [Completely remakes group 1](#)

This group includes words that can be associated with the concept of "sky" or celestial objects. Now, for the next group:

Even with another hint it cannot correctly complete the group



Do you know that the word raspberry can also be a disapproving vocal expression? Try adding this into your original group 1.



I see your point about "raspberry" as a disapproving vocal expression. Let's modify the original Group 1 to include it:



Group 1: ghost, boo, hiss, raspberry

Doesn't remember that I told it to remove ghost from group

Now, for the next group:

Group 2: lemon, garden, candy, flop

When asked, says these are related as dessert flavors

Do these words form a correct group, or should I make further adjustments?

Let's learn more about how LLMs like ChatGPT work

- First we will see how words are transformed into numeric data suitable for deep learning
- As you learn more about how LLMs work, think about why the word game Connections might be challenging for LLMs

Article in the *Financial Times*

Artificial Intelligence

Generative AI exists because of the transformer

This is how it works

By Visual Storytelling Team and Madhumita Murgia in London SEPTEMBER 12 2023

All of the visualizations included below (except where otherwise noted) are from this article

Word embeddings

To write text, LLMs must first translate words into a language they understand.

We go to work by train

First a block of words is broken into **tokens** — basic units that can be encoded. Tokens often represent fractions of words, but we'll turn each full word into a token.

We go to work by train

Questions about LLMs

— What are word embeddings and how are they created?

- Continuous Bag of Words is one method

one-hot representation

categorical, for ex

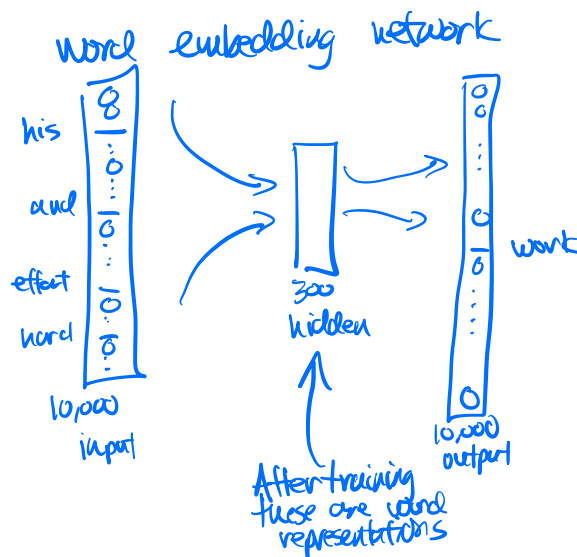
cat vs dog
1 0 0 1

ex 10 digits

1 0 0 0 0 0 0 0 0 0 = zero
0 1 0 0 0 0 0 0 0 0 = one
0 0 1 0 0 0 0 0 0 0 = two
etc.

Use only 10,000 most common English words
one particular location is set to 1 for every
single word

Then we process many texts where the goal is to predict
the center word from the surrounding words
"his hard work and effort"

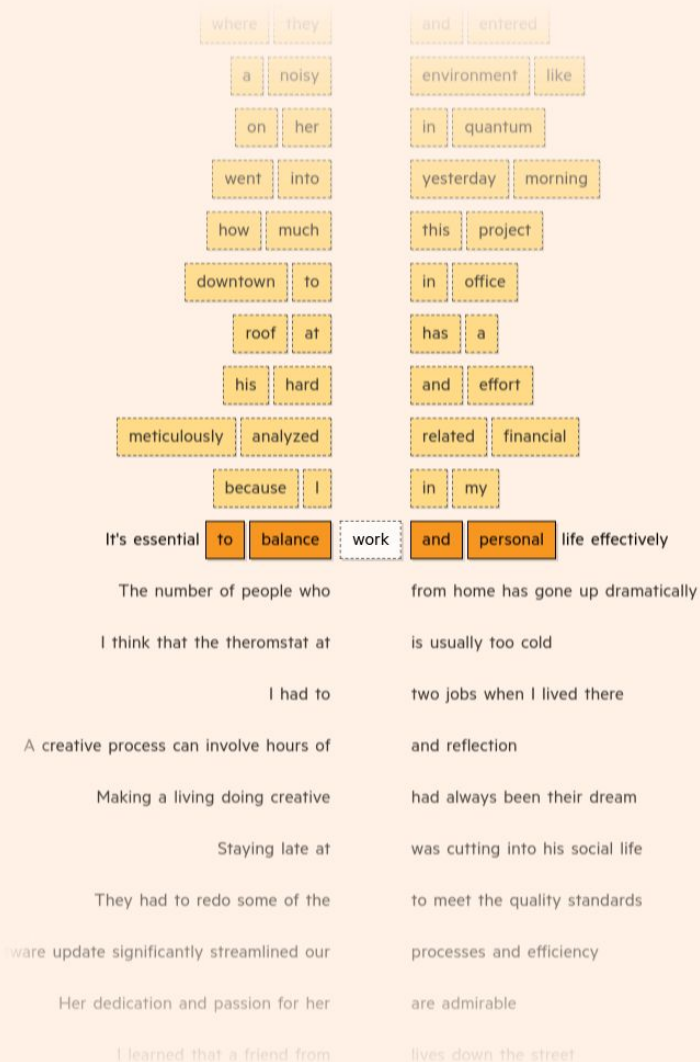


The network develops
hidden layer representations
that reflect the contexts
in which the words
occur.

Thus semantically similar
words have similar
hidden layer representations.

Word embeddings

In order to grasp a word's meaning, **work** in our example, LLMs first observe it in context using enormous sets of training data, taking note of **nearby words**. These datasets are based on collating text published on the internet, with new LLMs trained using billions of words.

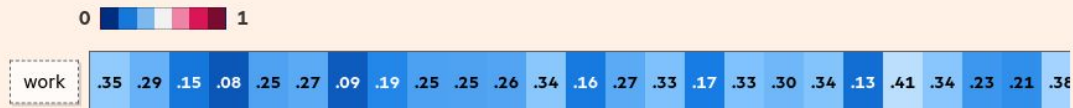


Word embeddings

Eventually, we end up with a huge set of the words found **alongside** **work** in the training data, as well as those that **weren't** found near it.

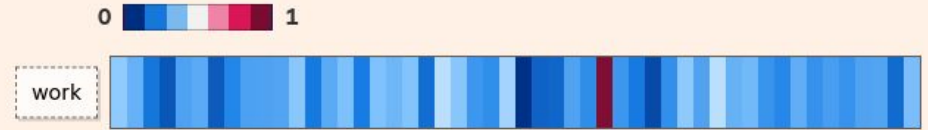
work	her	work	meet
work	atmosphere	work	dove
work	are	work	zebra
work	the	work	processes
work	to	work	streamlined
work	of	work	and
work	admirable	work	polka
work	our	work	for

As the model **processes** this set of words, it produces a vector — or list of values — and adjusts it based on each word's proximity to **work** in the training data. This vector is known as a word embedding.



Word embeddings

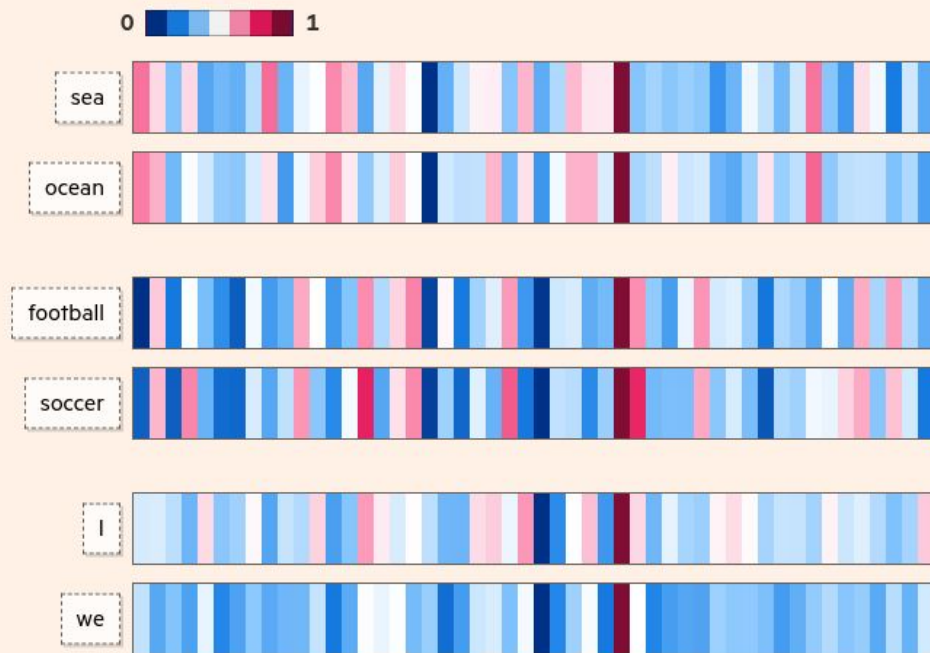
A word embedding can have hundreds of values, each representing a different aspect of a word's meaning. Just as you might describe a house by its characteristics — type, location, bedrooms, bathrooms, storeys — the values in an embedding quantify a word's linguistic features.



Word embeddings

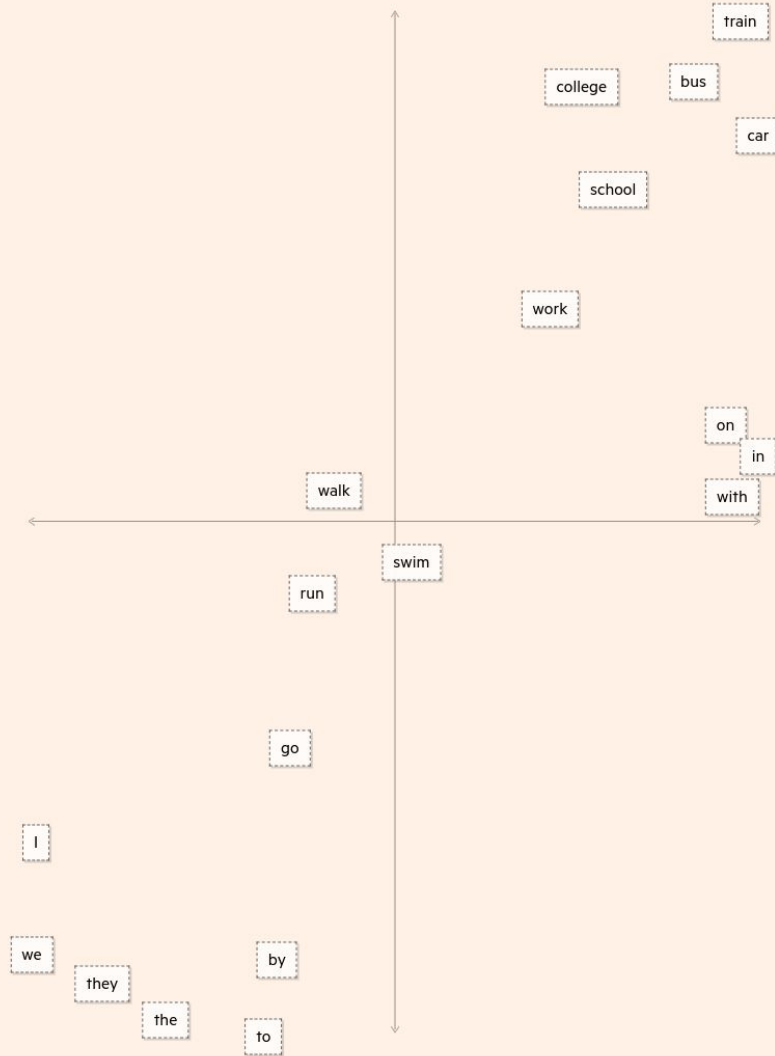
The way these characteristics are derived means we don't know exactly what each value represents, but words we expect to be used in comparable ways often have similar-looking embeddings.

A pair of words like **sea** and **ocean**, for example, may not be used in identical contexts ('all at ocean' isn't a direct substitute for 'all at sea'), but their meanings are close to each other, and embeddings allow us to quantify that closeness.



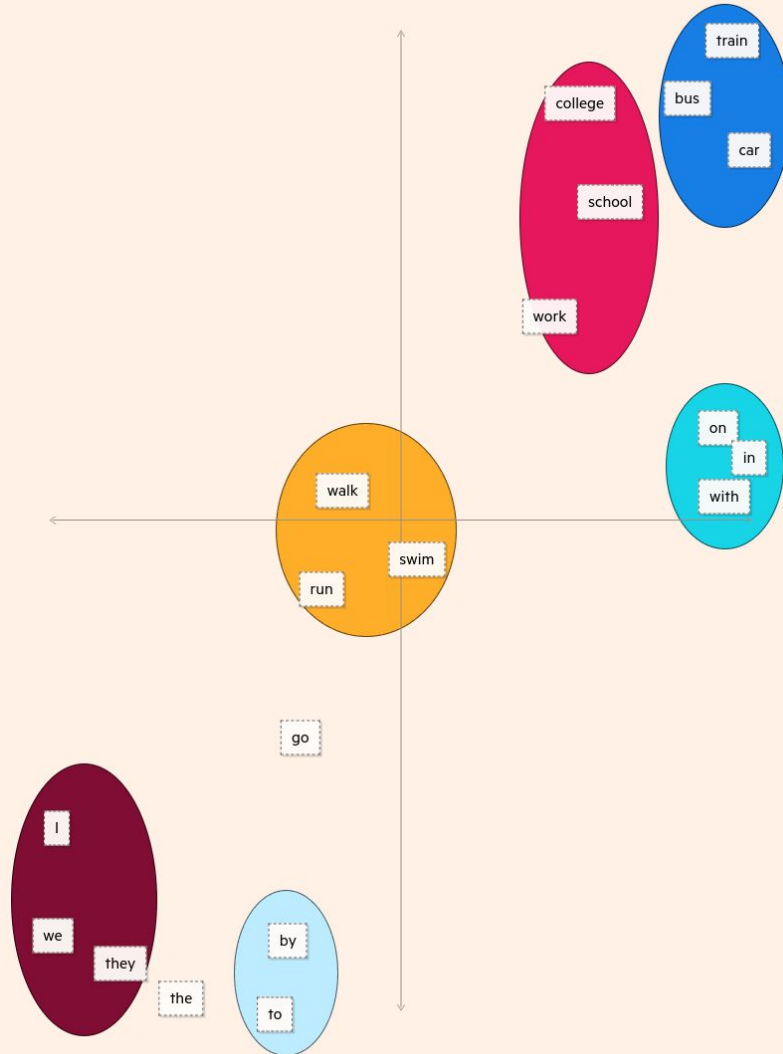
Word embeddings

By reducing the hundreds of values each embedding represents to just two, we can see the distances between these words more clearly.



Word embeddings

We might spot clusters of **pronouns**, or modes of **transportation**, and being able to quantify words in this way is the first step in a model generating text.



Let's learn about how LLMs generate text

- Transformers
- Self-attention
- These concepts can be applied to more than just text—it also works to produce pictures, music, and computer code

Transformers

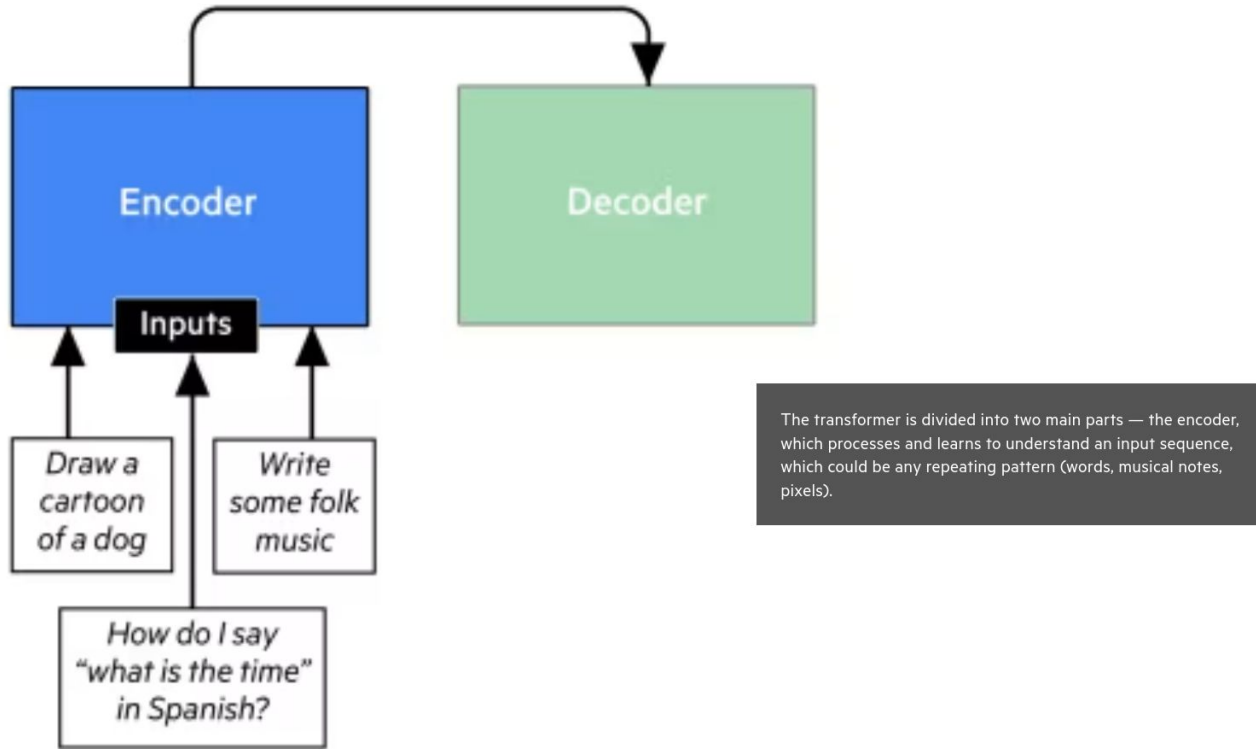


Image the Financial Times article entitled "Transformers: The Google scientists who pioneered an AI revolution"

Transformers

And the decoder, which produces an output sequence (a sentence, a piece of music, a picture).

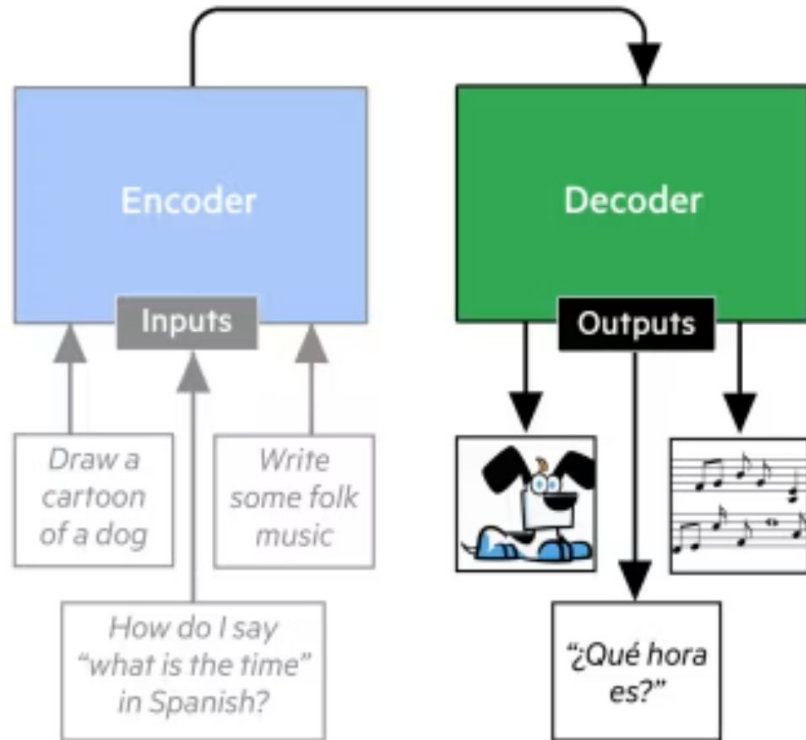


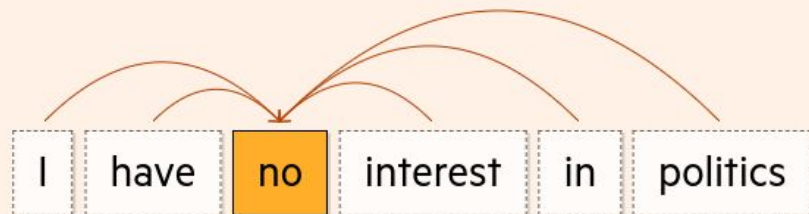
Image from the Financial Times article entitled "Transformers: The Google scientists who pioneered an AI revolution"

Transformers and self-attention

A key concept of the transformer architecture is self-attention. This is what allows LLMs to understand relationships between words.

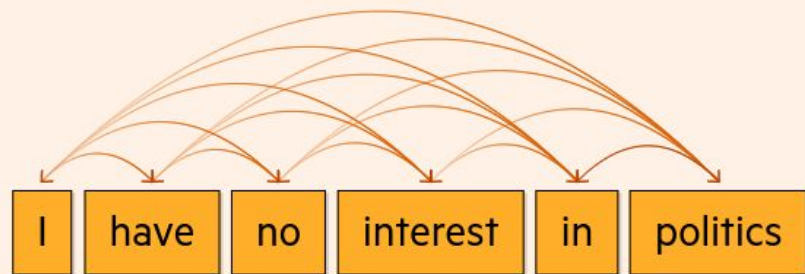
I have no interest in politics

Self-attention looks at each **token** in a body of text and decides which others are most important to understanding its meaning.



Transformers and self-attention

With self-attention, the transformer computes all the words in a sentence at the same time. Capturing this context gives LLMs far more sophisticated capabilities to parse language.



Transformers and self-attention

In this example, assessing the whole sentence at once means the transformer is able to understand that **interest** is being used as a noun to explain an individual's take on politics.



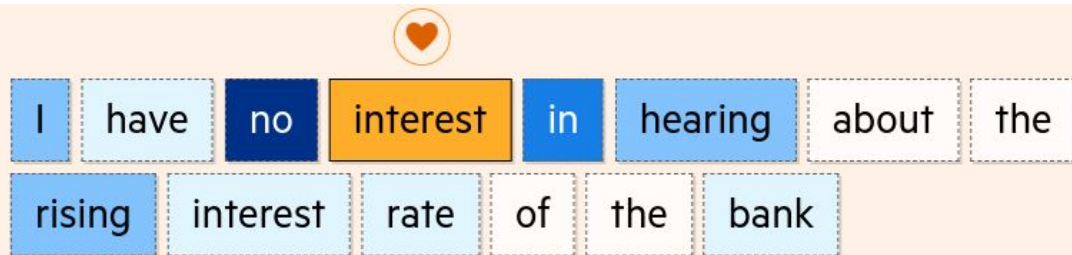
... the model understands **interest** is now being used in a financial sense.



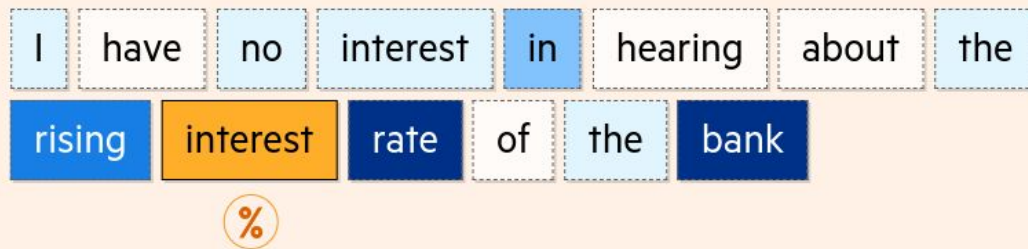
Transformers and self-attention

And when we combine the sentences, the model is still able to recognise the correct meaning of each word thanks to the attention it gives the accompanying text.

For the first use of interest, it is **no** and **in** that are most attended.



For the second, it is **rate** and **bank**.



Language challenge for AI

- Binary choice questions
- Two entities are mentioned in the question
- A pronoun is used to refer to one of them
- What is the correct referent of the pronoun?
- Changing just one word in the statement, changes the correct referent
- How will LLMs do at this task?

For example

- The dog chewed the bone because **it** was **hungry**.
- The dog chewed the bone because **it** was **delicious**.

What does the pronoun “it” refer to in each of these sentences?

How self-attention helps focus on the correct words

In the following sentence, self-attention is able to calculate that **it** is most likely to be referring to **dog**.

The **dog** **chewed** the bone because **it** was
hungry

And if we alter the sentence, swapping **hungry** for **delicious**, the model is able to recalculate, with **it** now most likely to refer to **bone**.

The **dog** **chewed** the **bone** because **it** was
delicious

Another example

- Sam painted a picture of shepherds with sheep, but **they** ended up looking like **dogs**.
- Sam painted a picture of shepherds with sheep, but **they** ended up looking like **golfers**.

What does the pronoun “they” refer to in each of these sentences?

ChatGPT doesn't fare as well in this example



In the sentence: Sam tried to paint a picture of shepherds with sheep, but they ended up looking like dogs, who does "they" refer to?



In the sentence "Sam tried to paint a picture of shepherds with sheep, but they ended up looking like dogs," the pronoun "they" refers to the shepherds and the sheep collectively. It suggests that the entire scene, including both the shepherds and the sheep, ended up looking like dogs in the painting.



Here it suggests the correct referent, but it's explanation is disappointing

LM

In the sentence "Sam tried to paint a picture of shepherds with sheep, but they ended up looking like golfers, who does "they" refer to?



In the sentence "Sam tried to paint a picture of shepherds with sheep, but they ended up looking like golfers," the pronoun "they" refers to the shepherds. This sentence indicates that the shepherds ended up looking like golfers in the painting, not the sheep.



LM

Why might shepherds look like golfers?



Shepherds might look like golfers in a painting or visual representation due to the artist's skill or the way they have depicted the individuals. It could be the result of artistic interpretation, style, or errors in the depiction. The comparison between shepherds and golfers could be unintentional and merely an outcome of the artist's technique or a misinterpretation of the intended subject matter. Artistic representations can sometimes lead to unexpected or humorous visual resemblances that were not the artist's original intention.



Though the desired information does seem to be present in the model

LLMs face the *Symbol grounding problem* (see Steve Harnad). They are *not embodied* -- only experience of the world is thru text.



What do shepherds hold that might look like a golf club?



Shepherds typically hold a shepherd's crook or staff, which might, in some artistic representations or from a certain angle, bear a resemblance to a golf club. Both the shepherd's crook and a golf club are long, slender objects that the individuals hold, but their shapes and purposes are quite different. A shepherd's crook is traditionally used by shepherds for tending to and guiding their sheep, while a golf club is used in the sport of golf to strike a ball into a hole. The similarity in appearance might be purely coincidental or a result of artistic interpretation, as mentioned earlier.



Transformers and self-attention

The benefits of self-attention for language processing increase the more you scale things up. It allows LLMs to take **context** from beyond sentence boundaries, giving the model a greater understanding of how and when a word is used.

The **dog** chewed the bone because it was delicious.

In a quiet little village nestled amidst picturesque landscapes, there lived a faithful canine with a red collar that adorned his neck like a crown. This charming creature was a beloved member of the Johnson family. With a gleaming coat of golden fur and eyes that sparkled with warmth and affection, he won the hearts of everyone who crossed his path. The red collar became a symbol of his loyalty and an emblem of the countless adventures that awaited him. From the day Ben trooped into the Johnsons' lives, he brought an abundance of joy and laughter. His days were filled with frolics in the nearby park, chasing butterflies, and playing fetch with the children. In the afternoons, he would faithfully accompany Mr. Johnson on his walks around the neighborhood, sniffing the scents of the world with curious enthusiasm. The townsfolk admired his friendly nature and the undeniable bond he shared with his family. Whenever he went, the red collar shone like a beacon, a reminder of the unconditional love he offered to those who embraced him. Every day, Mrs. Johnson ensured that Ben's dinner was prepared with love and care, with a wholesome blend of nutritious ingredients. Ben would patiently sit and wag his tail eagerly, until the clock struck six. As the aroma of his favorite meal wafted through the air, he couldn't contain his excitement. His red collar jingled with each step he took towards his feeding bowl, a sound that had become synonymous with the pensive anticipation of dinner time. Beyond the boundaries of the town, Ben's escapades expanded into a realm of wild exploration. He roamed through vast meadows and ventured into dense forests, his red collar contrasting against the vibrant hues of nature. On one such adventure, he met a pack of fellow canines, and together, they formed an inseparable bond. They navigated through the wilderness, encountering thrilling encounters with other animals, all while sharing tales of loyalty and bravery under the watchful stars. As time passed, Ben grew older, and the years began to leave their mark on his once vibrant fur. Though his steps may have slowed, his spirit remained unwearying. The red collar, now slightly faded, continued to adorn his neck, a symbol of the unforgettable memories he had woven into the fabric of his family's life. As he approached the twilight of his days, the Johnsons made sure to reciprocate the love and care he had bestowed upon them throughout the years. They cherished every moment, knowing that the time they spent together was a precious gift, and they were determined to make it memorable.

had a red collar

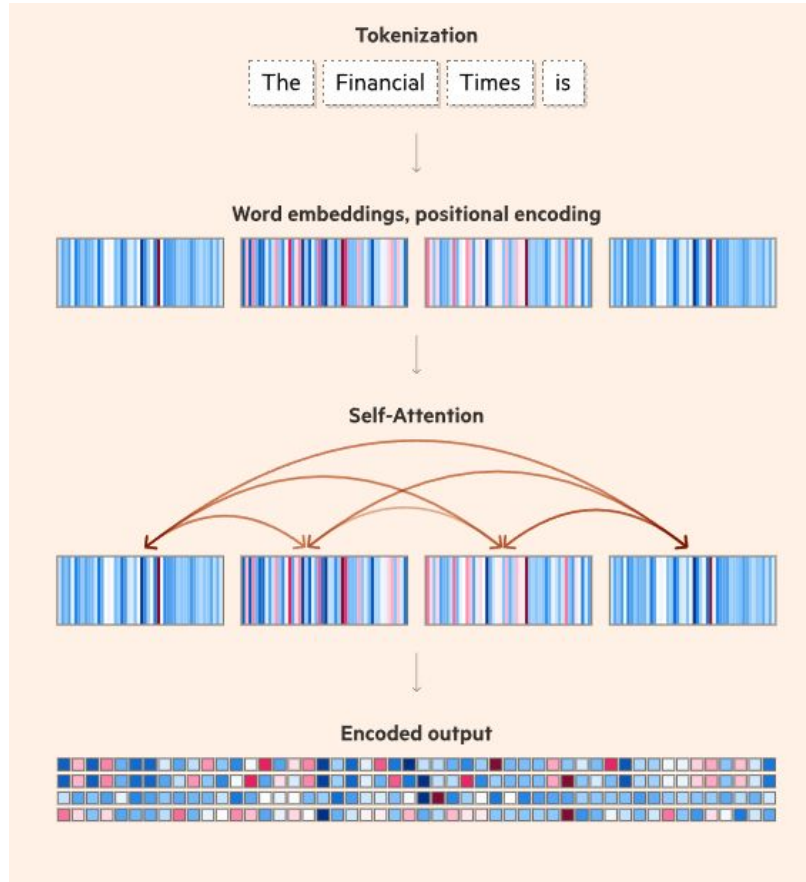
ate dinner at 6 pm

In a quiet little village nestled amidst picturesque landscapes, there lived a delightful canine named Luna, whose presence brought an unending sense of joy to all. Luna was a beautiful mix of Labrador and Border Collie, with a silky black coat that shimmered under the sun and a pair of striking amber eyes that sparkled with intelligence. The town's residents couldn't help but smile whenever they caught a glimpse of Luna's wagging tail and the exuberance in her every step. Her playful energy was infectious, drawing people from all walks of life to her side, eager to bask in the warmth of her companionship. Luna was adorned with the red collar that had become a symbol of her loyalty and an emblem of the countless adventures they embarked upon. She would join them in the meadows, chasing butterflies, and rolling in the soft grass. Her presence at the village square became a delightful spectacle, as she delighted in meeting new faces and playfully nuzzling anyone willing to indulge her in games of fetch or tug-of-war. Luna's boundless energy was a testament to the sheer happiness she brought to the lives of those who loved her. One of Luna's favorite pastimes was when she would play fetch with the children. Every afternoon, she would see a ball flying through the air, and her little friends would soon come running towards her. With a bright red ball gripped firmly in their hands, they would take turns throwing it far into the distance, and Luna would dash like the wind to retrieve it, her tail wagging back and forth in sheer delight. The children's laughter filled the air as they cheered her on, and the bond between Luna and her young playmates grew stronger with each game. Through their innocent games of fetch, they learned the value of camaraderie and the joy of giving and receiving unconditional love.

was his owner's best friend

loved playing fetch

Generating text



stochastic vs deterministic

LLMs as "stochastic parrots"

Generating text

At its simplest, the model's aim is now to **predict** the next word in a sequence and do this repeatedly until the output is complete.

The Financial Times is ...

about

more

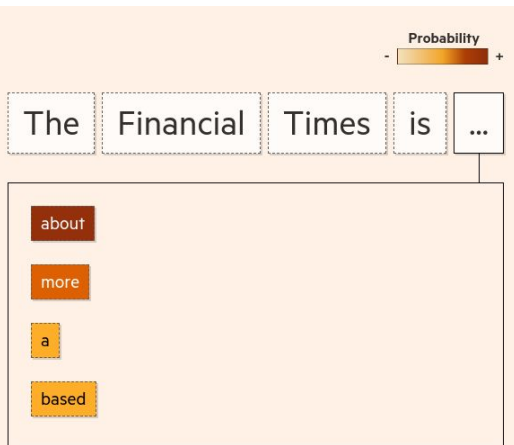
a

based

Paper: "On the dangers of stochastic parrots: Can large language models be too big?"

Generating text

To do this, the model gives a **probability score** to each token, which represents the likelihood of it being the next word in the sequence.




And it continues to do this until it is happy with the text it has produced.



Generating text

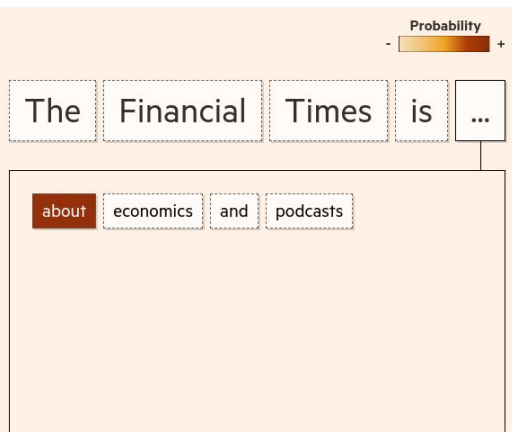
But this method of predicting the following word in isolation — known as “greedy search” — can introduce problems. Sometimes, while each individual token might be the next best fit, the full phrase can be less relevant.

Not necessarily always wrong, but perhaps not what you'd expect either.

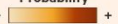
Probability 

The Financial Times is ...

about economics and podcasts



With beam search, the model is able to consider multiple routes and find the best option.

Probability 

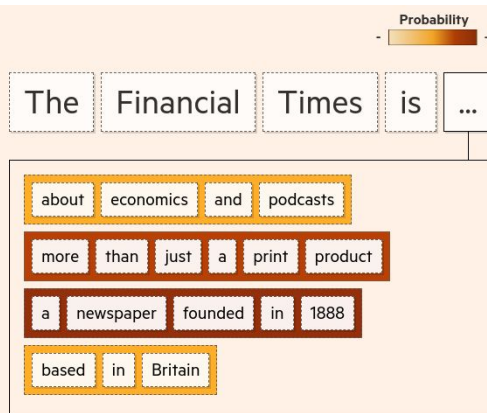
The Financial Times is ...

about economics and podcasts

more than just a print product

a newspaper founded in 1888

based in Britain



Generating a paragraph of text



Write me a short paragraph about my cat Pumpkin. She is an orange tabby and a little bit plump. She was born in October. She doesn't like my dog Cosmo, who chases her around the house. She loves to eat and demands to be fed promptly at 6 every morning by meowing loudly outside my bedroom door.



Pumpkin, my delightful orange tabby cat, is a charming bundle of feline personality. Her autumnal coat gives her an air of warmth and comfort. Slightly plump and irresistibly cute, she was born in October, and her birthday suits her well. Despite her plump appearance, she's remarkably agile when darting away from my dog, Cosmo, who delights in chasing her around the house. Pumpkin has a strong penchant for food, and she isn't shy about voicing her desire. Each morning at 6, she punctuates the silence with loud meows outside my bedroom door, a daily reminder of her unwavering breakfast schedule. She might be finicky around Cosmo, but she's the undisputed queen of our morning routine.



What makes the Connections game challenging for LLMs?

- The player is only provided with 16 words without any context
- The transformer and self-attention mechanisms need context to help disambiguate the words and determine their intended meaning

Summarizing capabilities of LLMs like ChatGPT

- They are not search engines—they are instead pattern matching engines
- They write text that is reasonably coherent and plausible, but there are no guarantees that it is factually correct
- They often fabricate information, generating made up names, dates, article titles, and quotes
- They have been trained on text scraped from the internet, and thus absorb the cultural biases inherent within that text