Impact of ChatGPT and Large Language Models on Higher Education

Lane Lawley

4/7/2023

What *is* language modeling?

Language is a complicated thing.

What *is* language modeling?

Language is a complicated thing.

When computer scientists encounter complicated things…

…we pretend they're probability distributions.

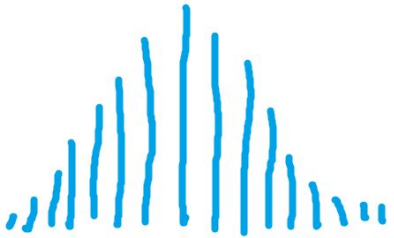

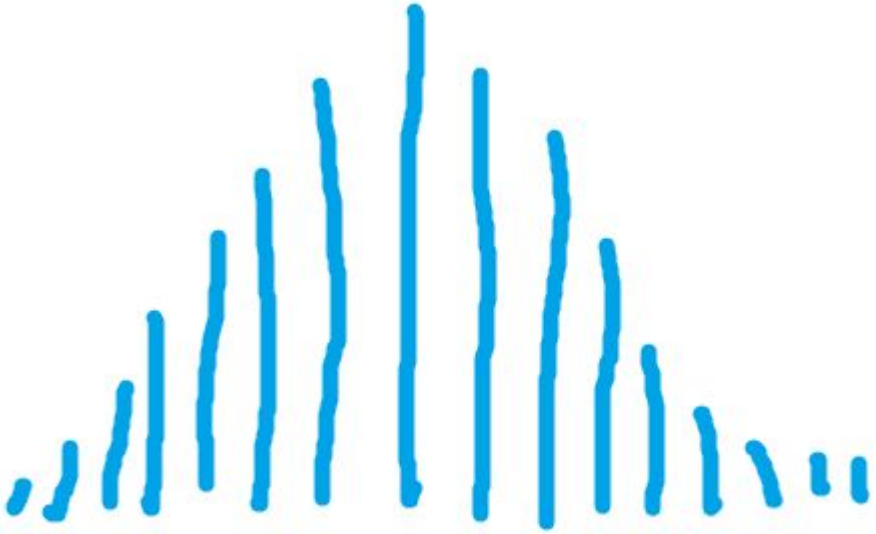

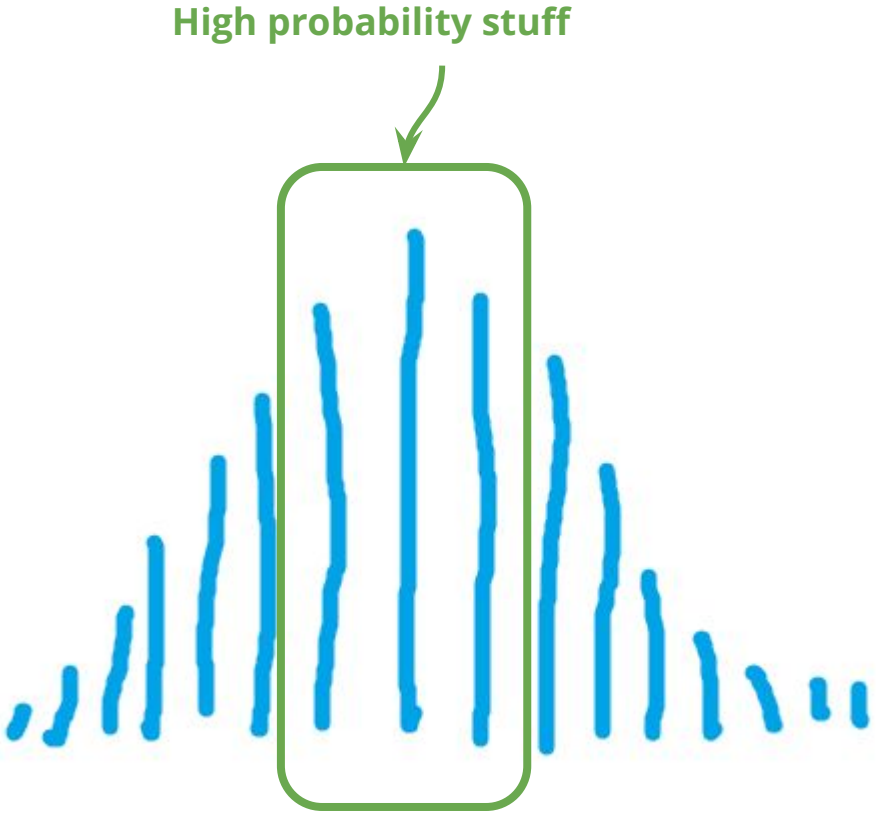
High probability stuff

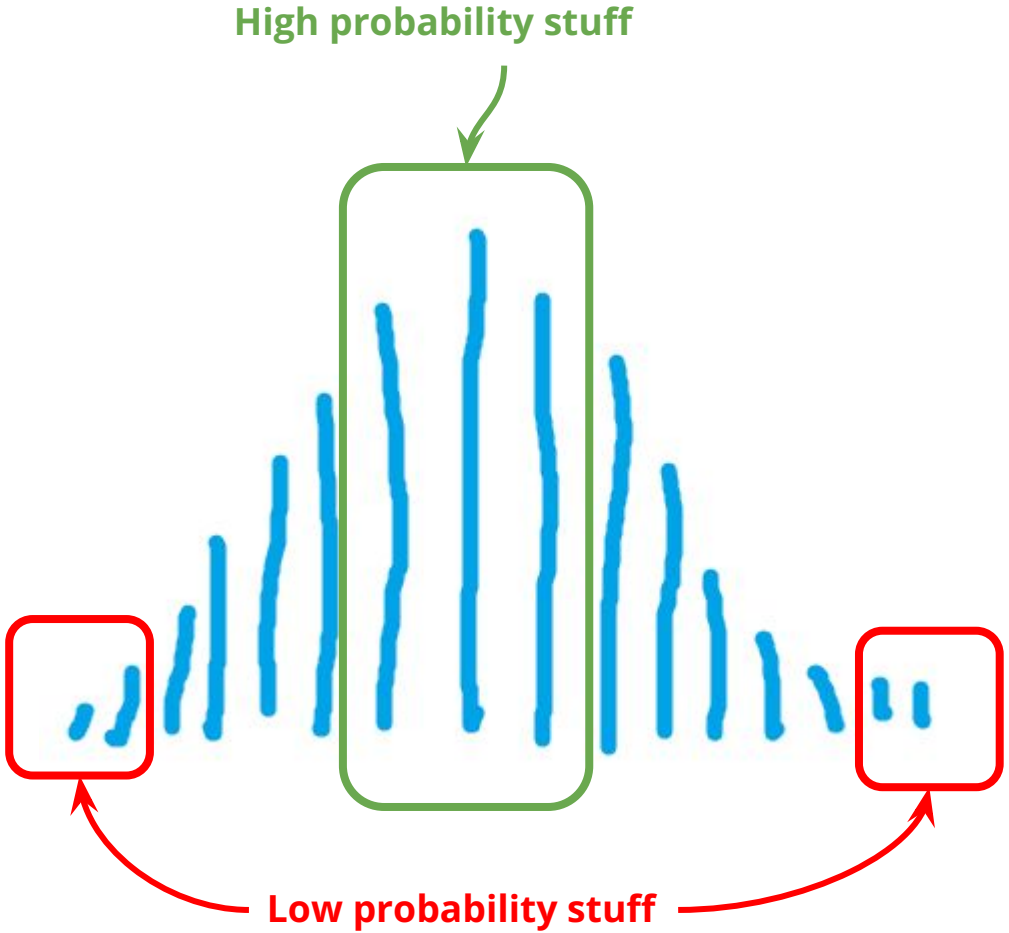
High probability stuff
Low probability stuff

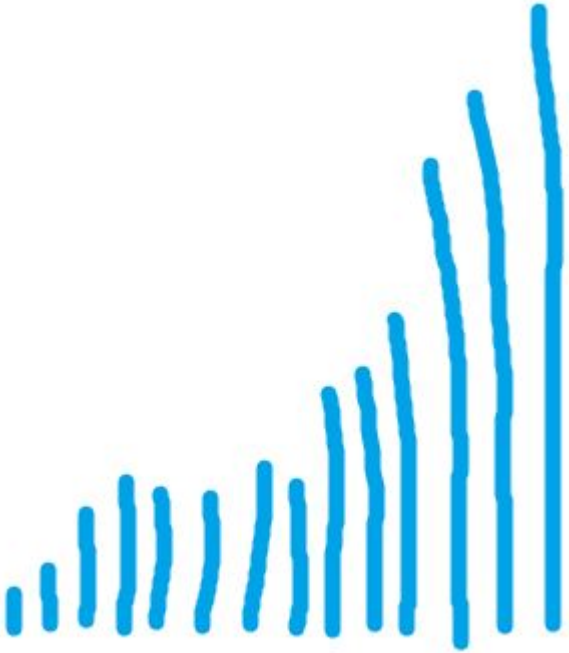

“Two plus two equals four!”

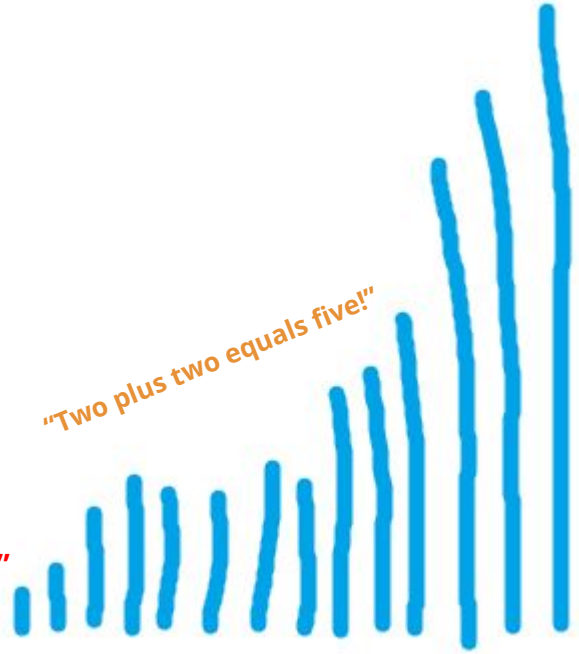

“Aasd;uhdfagasd”

Problem:

This distribution shows **14 bars**.

There are a lot more than 14 sentences!

“Two plus two equals four!”

Problem:

This distribution shows **14 bars**.

There are a lot more than 14 sentences!

Solution:

Define the distribution of sentences *in terms of a simpler distribution*.

Defining Sentence Probability

I promise this math is leading somewhere soon!

Defining Sentence Probability

P(One plus two equals three.)

=

P(One)

x P(plus | One)

x P(two | One plus)

x P(equals | One plus two)

x P(three. | One plus two equals)

Defining Sentence Probability

P(One plus two equals three.)

=

P(One)

x P(plus | One)

x P(two | One plus)

x P(equals | One plus two)

x P(three. | One plus two equals)

Now, we just have to define
P(single word **|** words before it**)**!

Defining Sentence Probability

P(One plus two equals three.)

=

P(One)

x P(plus | One)

x P(two | One plus)

x P(equals | One plus two)

x P(three. | One plus two equals)

Now, we just have to define **P(**single word | words before it**)**!

If we limit the number of words before it, let's say to **2,000 words** or so, we're closer to being able to put it in a computer!

Putting it in a computer

...But there are still *a lot* of combinations of 2,000 words.

We can either define P(word | 2,000 words) for *all* of those combinations, or we can **define an approximation**.

Core ideas of the approximation

1. The approximation is some **unknown function** that maps a set of 2,000 words onto **a probability distribution over the next word.**

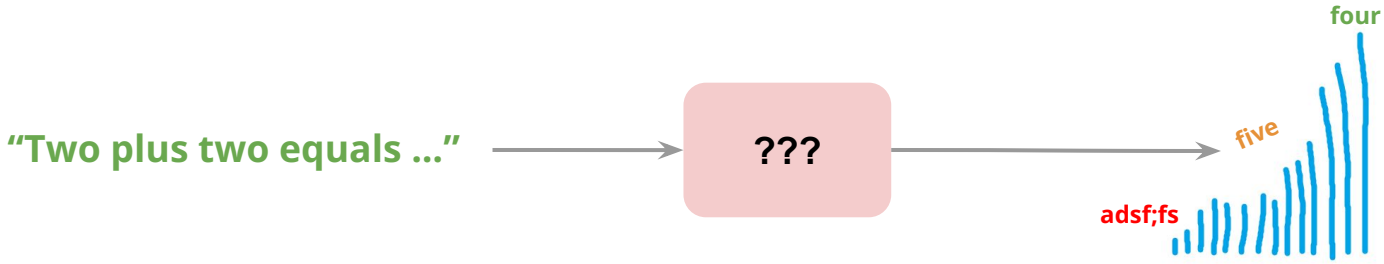

2. If we **learn** a version of this unknown function from *a lot* of **training data**, we hope it will **generalize** to unseen combinations of 2,000 words.

Say hello to the approximation!

???

=

GPT

Sidebar

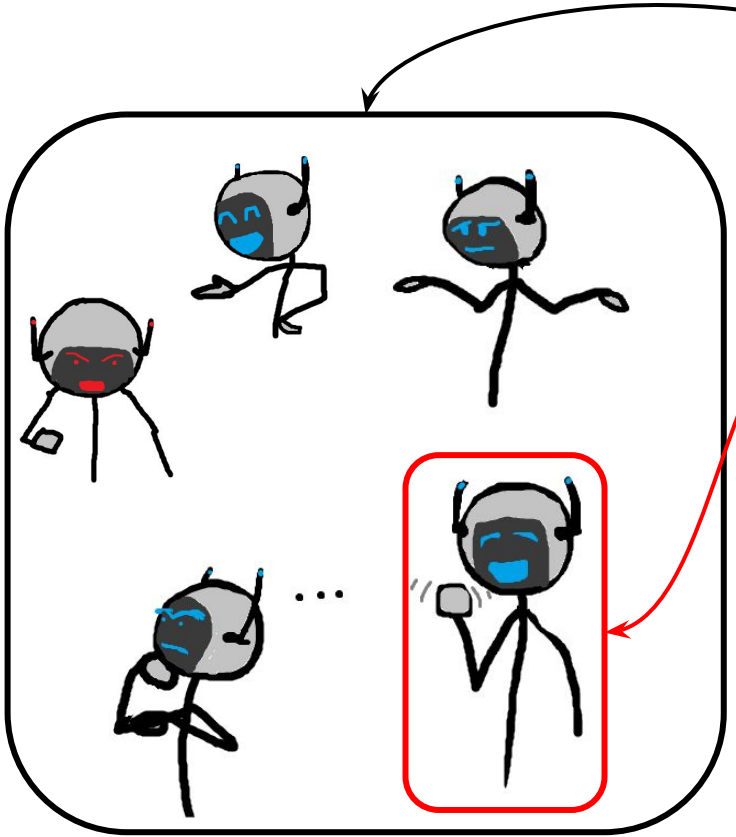

GPT is a *family of models*.

They all work the same way, but generally just use different training data.

ChatGPT is *one instance of a GPT model*, hyper-trained on chat data.

I know it's pedantic, but I wish people would stop using "ChatGPT" to refer to all GPT-family models.

GPT: a really, really good next-word predictor

What is
`P(COVID-19 | John caught the disease)`?

Uhh...it's *obviously* 0.000345.

Generating new sentences

GPT takes the “**context**” (~2,000 words), and maps that into a

probability distribution over the next word.

I love hanging out with my

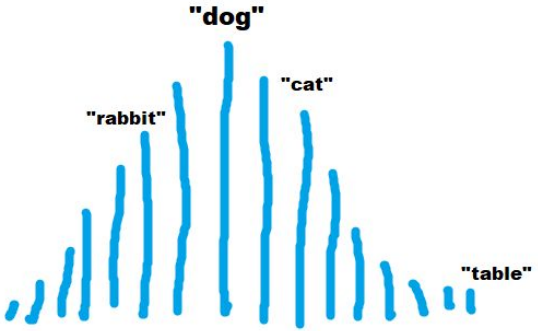

Generating new sentences

To generate more than one word,

the word it chooses is *added to the context*.

The process repeats until you've generated enough!

I love hanging out with my

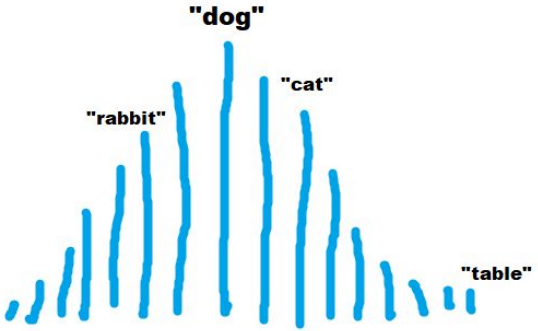

Generating new sentences

Remember: at each step, GPT could choose *any* word from the distribution.

I love hanging out with my

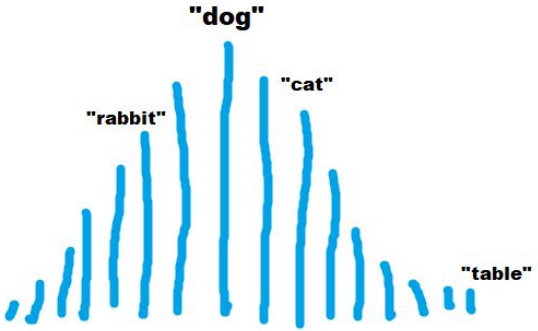

Generating new sentences

Remember: at each step, GPT could choose *any* word from the distribution.

One strategy is to always pick the most likely word.

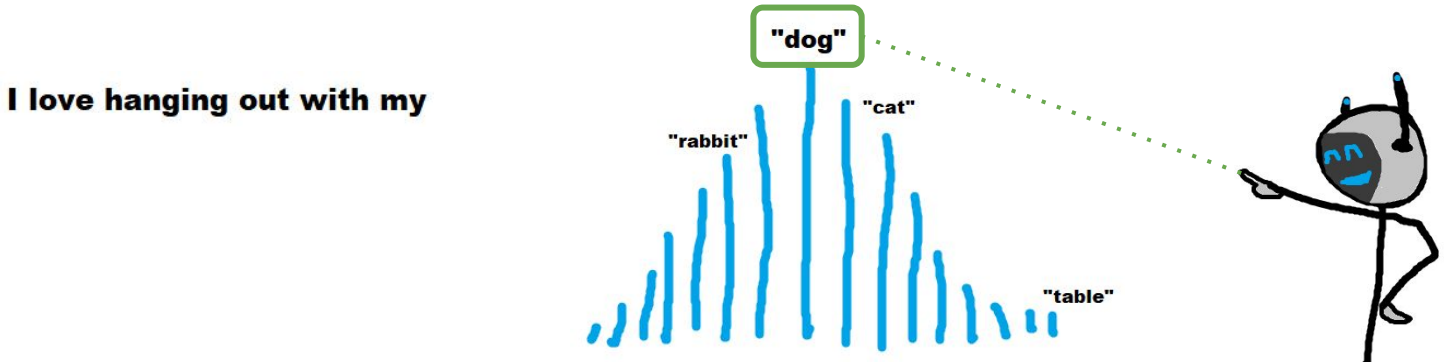

Generating new sentences

Remember: at each step, GPT could choose *any* word from the distribution.

One strategy is to always pick the most likely word.

But you could also choose lower-probability words sometimes!

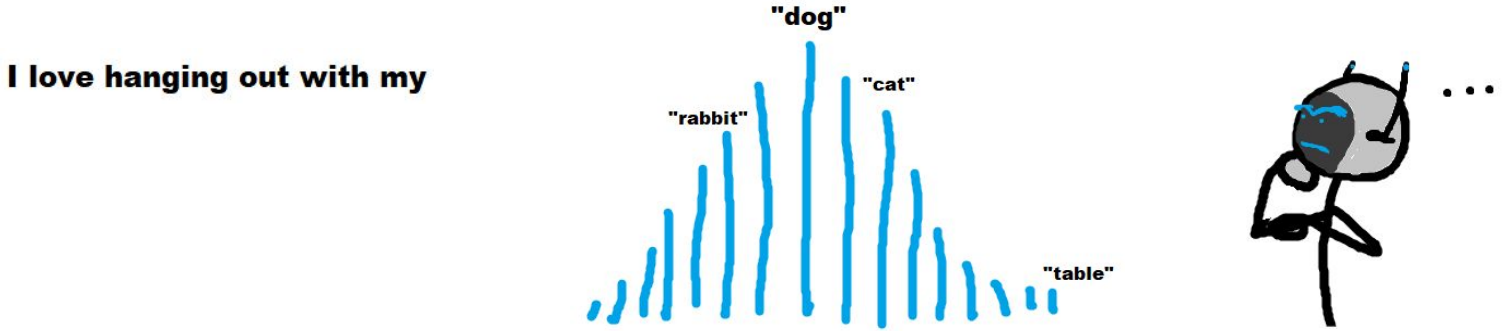

Generating new sentences

Remember: at each step, GPT could choose *any* word from the distribution.

One strategy is to always pick the most likely word.

But you could also choose lower-probability words sometimes!

I love hanging out with my

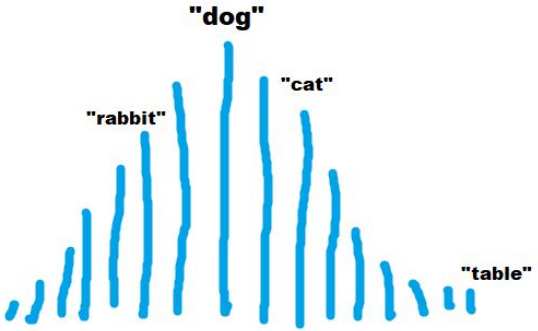

Generating new sentences

Remember: at each step, GPT could choose *any* word from the distribution.

One strategy is to always pick the most likely word.

But you could also choose lower-probability words sometimes!

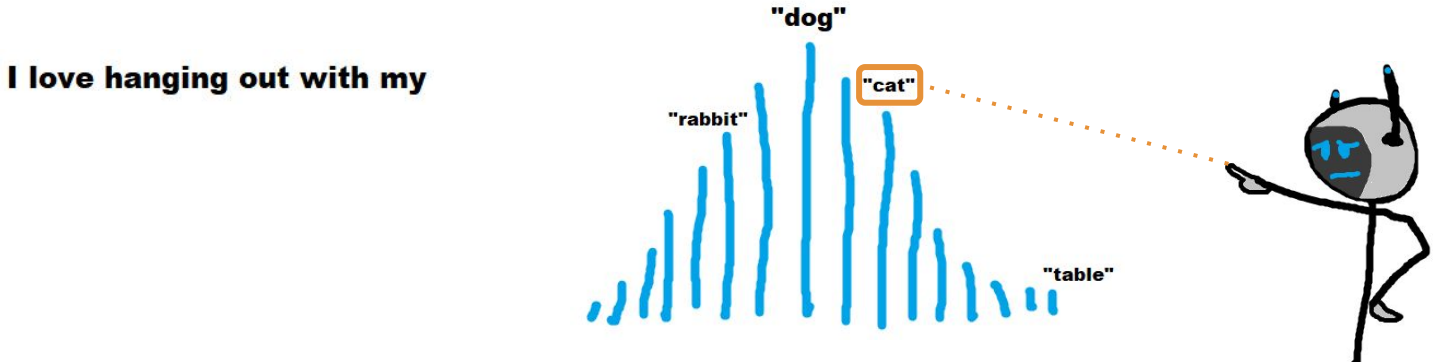

Prompting

Language models predict **P(**word **|** *past context***)**.

The careful design of that past context, a.k.a. the *prompt*, can accomplish a lot:

Prompting

Language models predict **P(**word **|** *past context***)**.

The careful design of that past context, a.k.a. the *prompt*, can accomplish a lot:

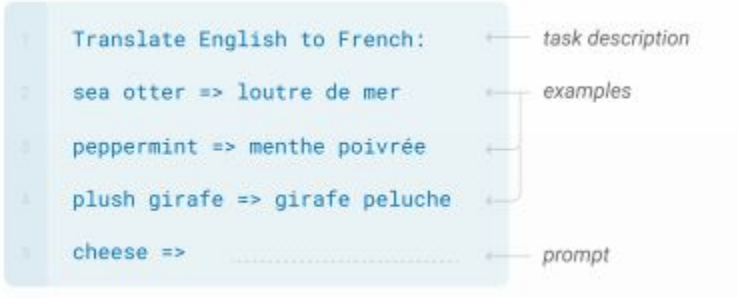

Prompting

Language models predict **P(**word **|** *past context***)**.

The careful design of that past context, a.k.a. the *prompt*, can accomplish a lot:

P(word **|** Translate English to French:
sea otter => loutre de mer
peppermint => menthe poivré
plush girafe => girafe peluche
cheese =>**)**

Prompting

Language models predict **P(**word **|** *past context***)**.

The careful design of that past context, a.k.a. the *prompt*, can accomplish a lot:

P(word **|** Translate English to French:
sea otter => loutre de mer
peppermint => menthe poivré
plush girafe => girafe peluche
cheese =>**)**

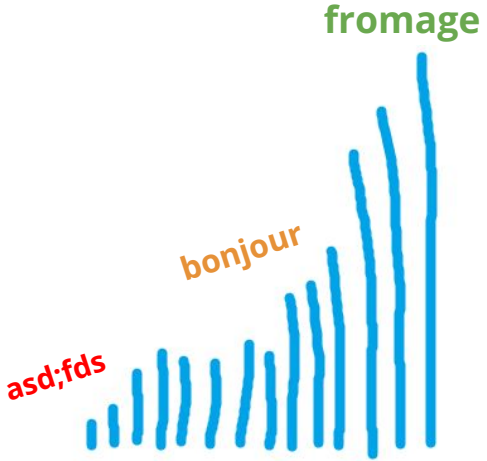

Prompting

Language models predict **P(**word | *past context*)**.**

The careful design of that past context, a.k.a. the *prompt*, can accomplish a lot:

P(word | Translate English to French:
sea otter => loutre de mer
peppermint => menthe poivré
plush girafe => girafe peluche
cheese =>**)**

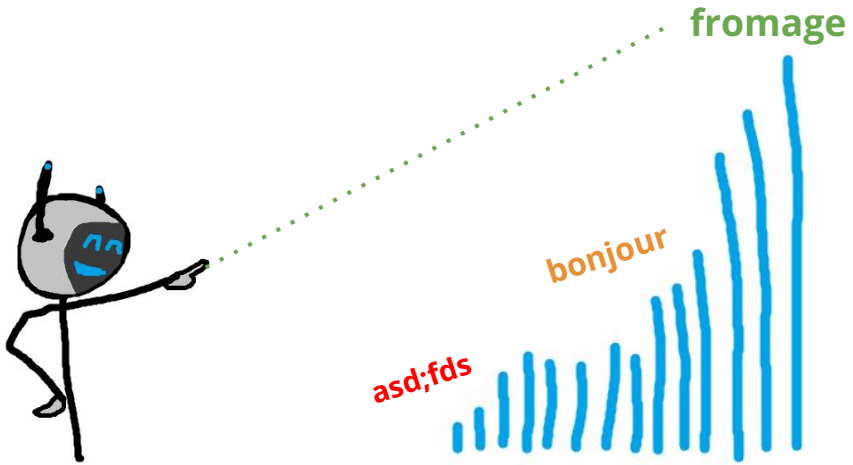

Prompting

Why does prompting work so well?

Prompting

Why does prompting work so well?

We don’t know for sure.

GPT is good at finding patterns in text, even if it hasn’t seen them before.

This ability helps it reproduce more complicated texts with less effort.

How did they make it a chatbot?

Through prompting!

How did they make it a chatbot?

Through prompting!

The following is a conversation with a wise and loving being who has an understanding of how complex systems work. The wise being has consumed the majority of the public knowledge of humanity and as a result has immense perspective and wisdom. This being is dedicated to helping the researcher. The being wants to help the researcher be successful.

Researcher: Grateful to talk with you again.

Wise Being:

How did they make it a chatbot?

Through prompting!

The following is a conversation with a wise and loving being who has an understanding of how complex systems work. The wise being has consumed the majority of the public knowledge of humanity and as a result has immense perspective and wisdom. This being is dedicated to helping the researcher. The being wants to help the researcher be successful.

Researcher: Grateful to talk with you again.

Wise Being: I am also grateful for this opportunity.

How did they make it a chatbot?

Through prompting!

NB: the language model will continue generating until you tell it to stop.

That means it will even generate *your* response.

How did they make it a chatbot?

Through prompting!

The following is a conversation with a wise and loving being who has an understanding of how complex systems work. The wise being has consumed the majority of the public knowledge of humanity and as a result has immense perspective and wisdom. This being is dedicated to helping the researcher. The being wants to help the researcher be successful.

Researcher: Grateful to talk with you again.

Wise Being: I am also grateful for this opportunity.

How did they make it a chatbot?

Through prompting!

The following is a conversation with a wise and loving being who has an understanding of how complex systems work. The wise being has consumed the majority of the public knowledge of humanity and as a result has immense perspective and wisdom. This being is dedicated to helping the researcher. The being wants to help the researcher be successful.

Researcher: Grateful to talk with you again.

Wise Being: I am also grateful for this opportunity.

Researcher: What is the meaning of life?

Wise Being: I’m glad you asked. The meaning of life is

How did they make it a chatbot?

Through prompting!

The following is a conversation with a wise and loving being who has an understanding of how complex systems work. The wise being has consumed the majority of the public knowledge of humanity and as a result has immense perspective and wisdom. This being is dedicated to helping the researcher. The being wants to help the researcher be successful.

Researcher: Grateful to talk with you again.

Wise Being: I am also grateful for this opportunity.

Researcher: What is the meaning of life?

This was generated by GPT!

Wise Being: I'm glad you asked. The meaning of life is

How did they make it a chatbot?

Through prompting!

The OpenAI website quietly cuts that part off for you.

It lets you replace the response instead, and never shows you what GPT would have generated for you.

```
The following is a conversation with a wise and loving being who has an understanding of how computer systems work. The wise being has consumed the majority of the public
knowledge of humanity and as a result has immense perspective [illegible] researcher. The being wants to help the researcher be
successful.

Researcher: Grateful to talk with you again.

Wise Being: I am also grateful for th[illegible]

Researcher: What is the meaning of life?

Wise Being: I'm glad you asked. The meaning of life is
```

This was generated by GPT!

How did they make it a chatbot?

Through prompting!

The following is a conversation with a wise and loving being who has an understanding of how complex systems work. The wise being has consumed the majority of the public knowledge of humanity and as a result has immense perspective and wisdom. This being is dedicated to helping the researcher. The being wants to help the researcher be successful.

Researcher: Grateful to talk with you again.

Wise Being: I am also grateful for this opportunity.

Researcher:

Why is GPT sometimes so wrong?

Beliefs, knowledge, and the problem of “not”

Many people only see ChatGPT as a chatbot.

The idea that it also tries to generate *your* responses isn’t apparent.

But it does. It generates **entire text sequences.** The chatbot mode is a “hack”.

Beliefs, knowledge, and the problem of “not”

Many people only see ChatGPT as a chatbot.

The idea that it also tries to generate *your* responses isn’t apparent.

But it does. It generates **entire text sequences.** The chatbot mode is a “hack”.

ChatGPT wants to model you as much as it models itself; it is not a single entity with its own knowledge and belief state!

Beliefs, knowledge, and the problem of “not”

What does it even mean to believe that a statement is true?

I was trained to think of something like this:

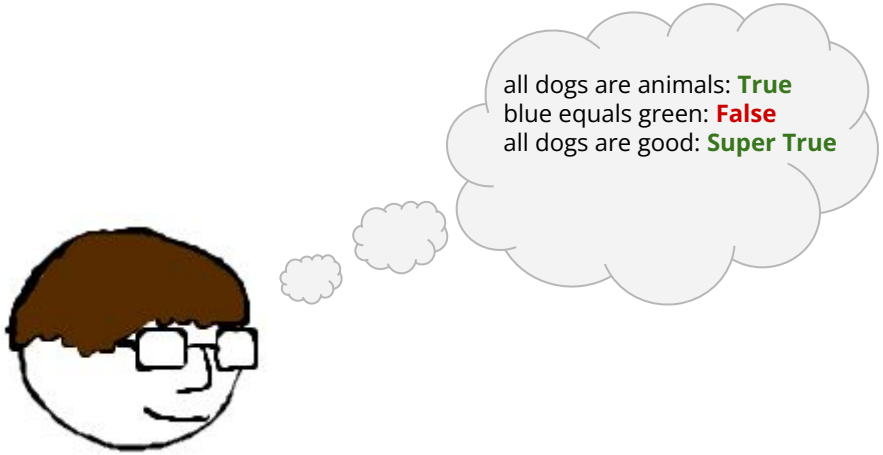

Beliefs, knowledge, and the problem of “not”

What does it even mean to believe that a statement is true?

I was trained to think of something like this:

...but GPT is beholden to the all-powerful *context words.*

Beliefs, knowledge, and the problem of “not”

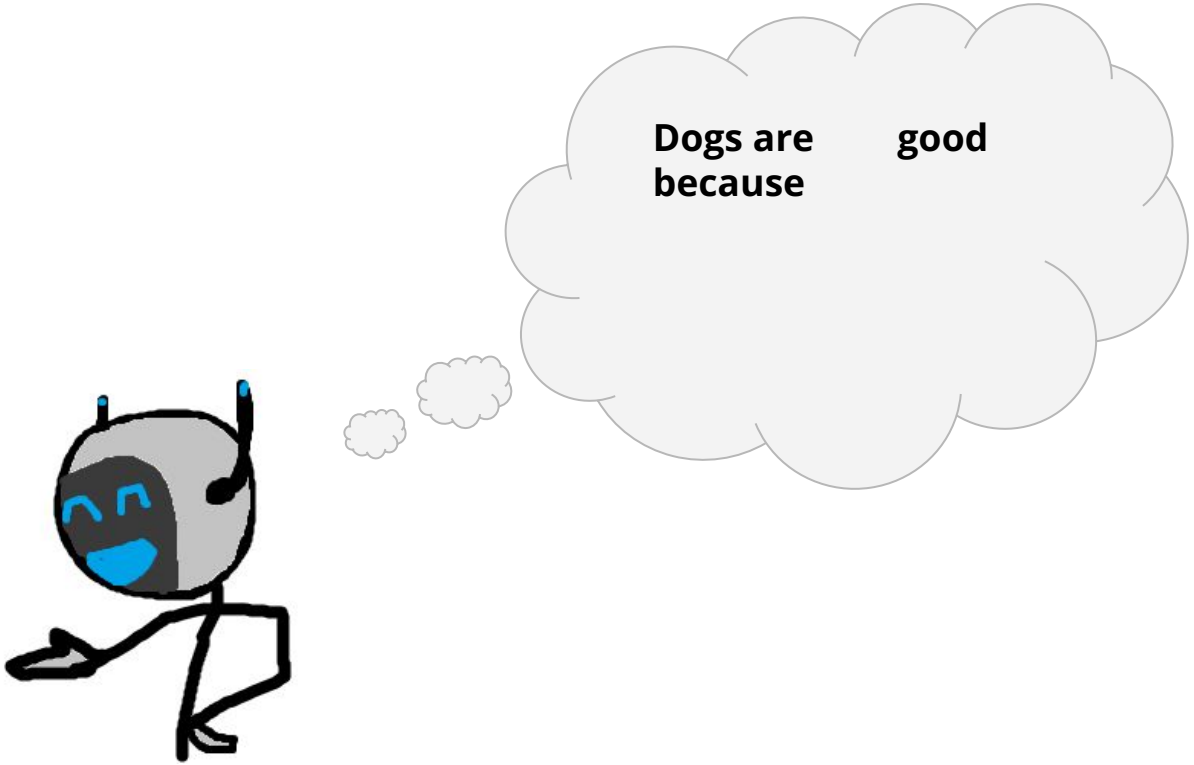

Beliefs, knowledge, and the problem of “not”

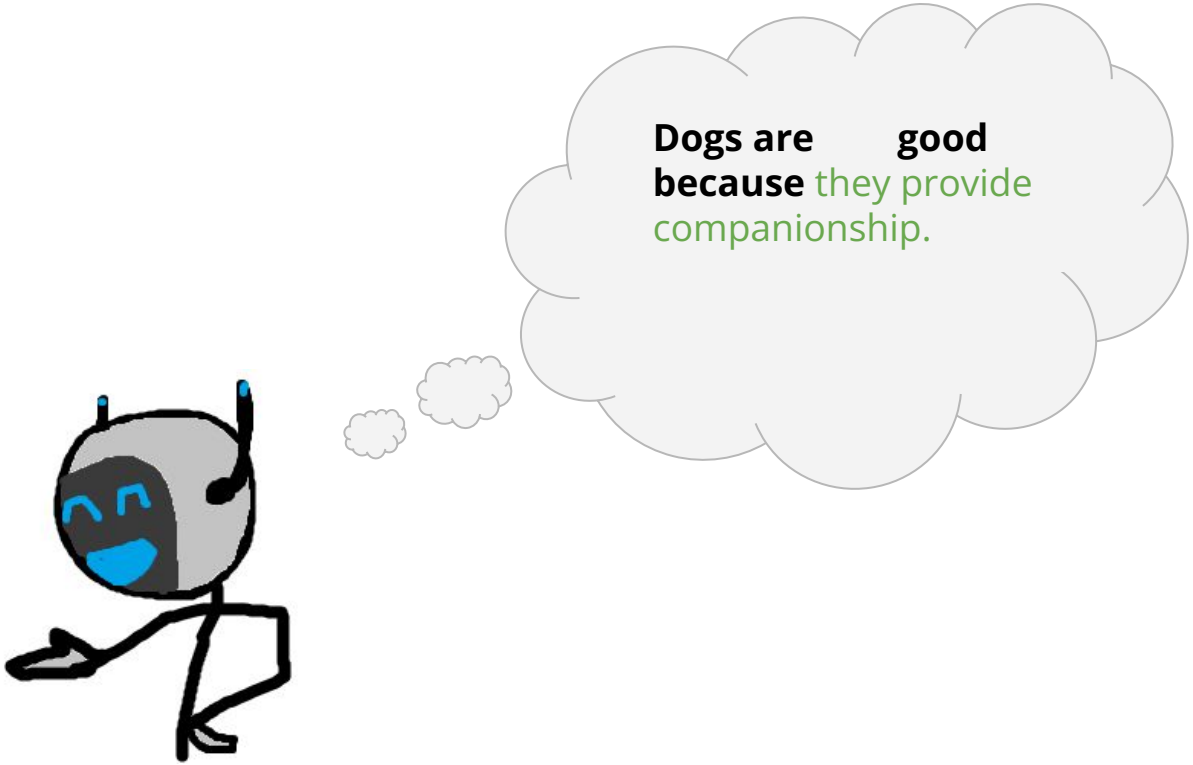

Beliefs, knowledge, and the problem of “not”

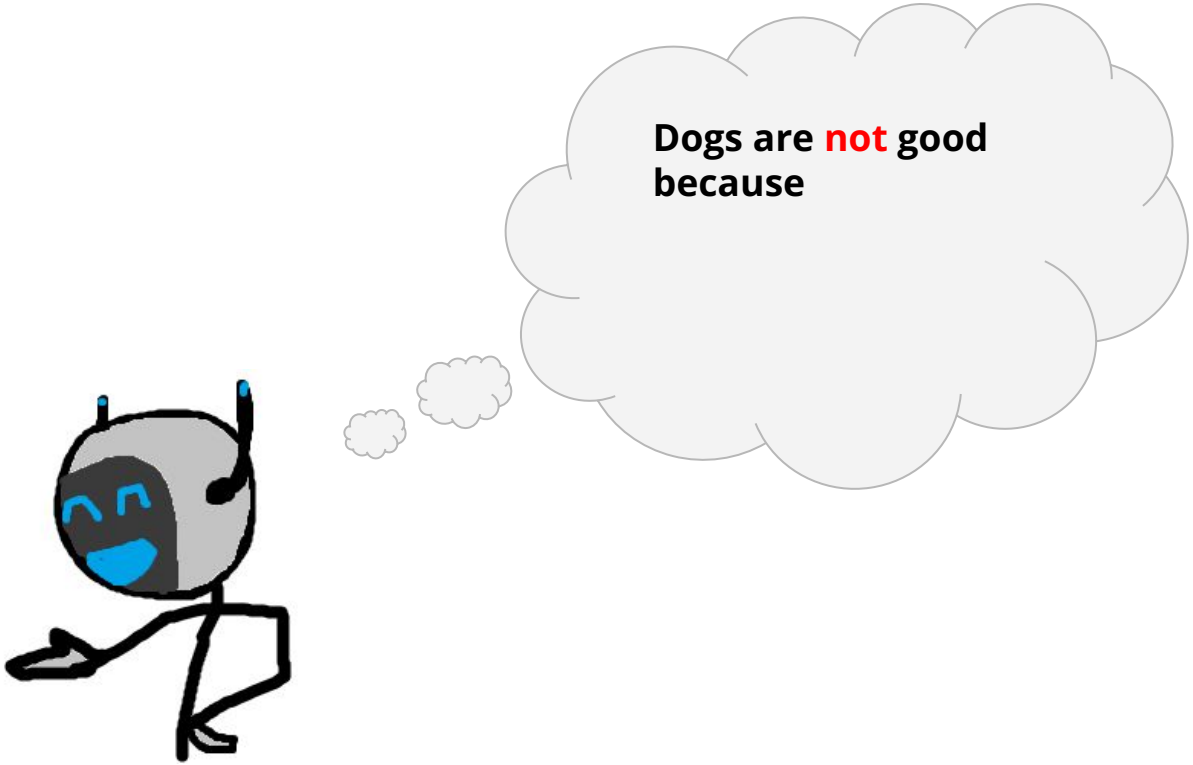

Beliefs, knowledge, and the problem of “not”

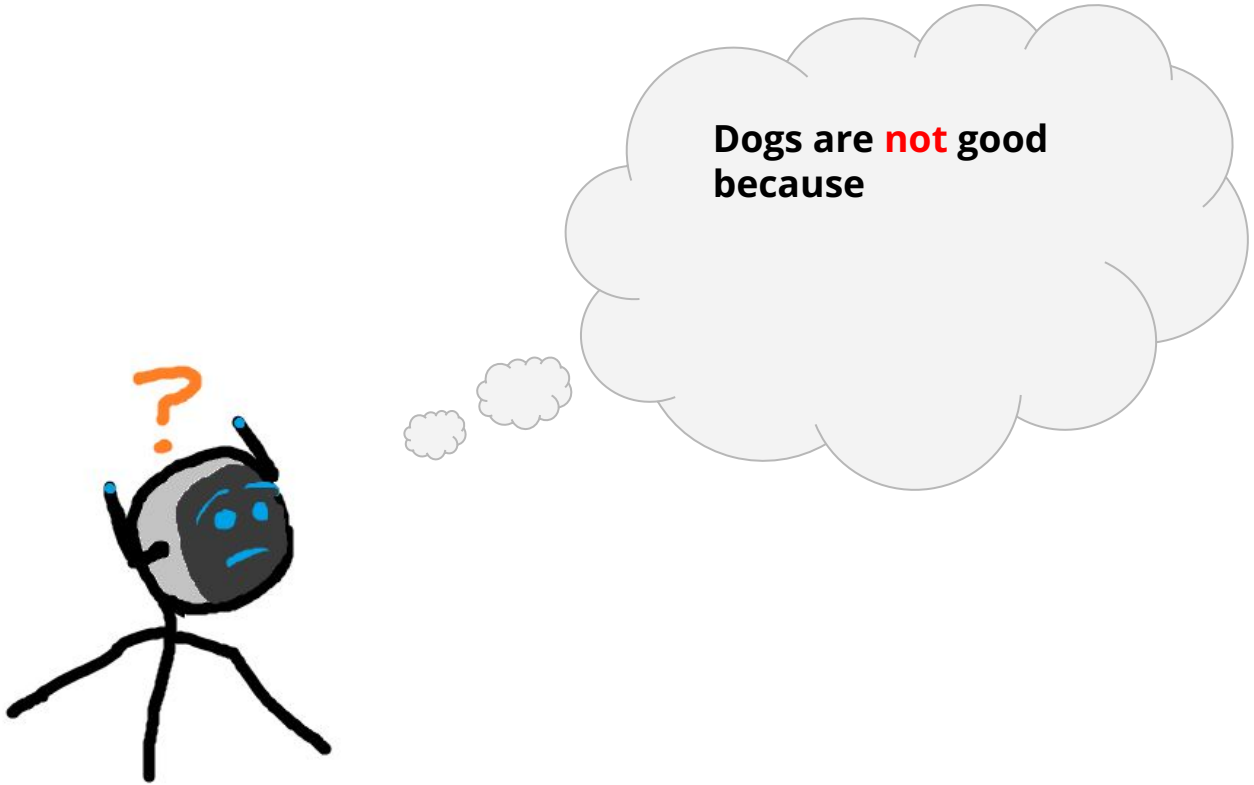

Beliefs, knowledge, and the problem of “not”

Beliefs, knowledge, and the problem of “not”

Beliefs, knowledge, and the problem of “not”

WRONG

Dogs are not good because they can be destructive, noisy, and difficult to train.

Beliefs, knowledge, and the problem of “not”

WRONG

Dogs are not good because they can be destructive, noisy, and difficult to train.

WRONG

Beliefs, knowledge, and the problem of “not”

WRONG

WRONG

Dogs are not good because they can be destructive, noisy, and difficult to train.

WRONG

Beliefs, knowledge, and the problem of “not”

WRONG

WRONG

Dogs are not good because they can be destructive, noisy, and difficult to train.

MEGA WRONG

WRONG

Beliefs, knowledge, and the problem of “not”

WRONG

WRONG

Dogs are not good because they can be destructive, noisy, and difficult to train.

MEGA WRONG

WRONG

VERY GOOD BOY

Recap

- GPT is a way to generate language given some prior "context".
- How this lets it do so many tasks is an outright generational mystery.
- GPT is not an "entity", even in conversations. It's trying to talk for you, too!
- GPT does not store a set of beliefs. It will say whatever is "likely".
- Just because GPT is *simple* doesn't mean that there's *no intelligence* there.

I picked up this "scarequote" habit from my old PhD advisor. I'm trying to quit.

So, how will this impact higher ed?

So, how will this impact higher ed?

…no, really, *how*?

I’m just an AI scientist! Let’s discuss this from *your* perspective.

I’m **Lane Lawley**!

I’m a postdoc at TAIL Teachable AI Lab.

tail.cc.gatech.edu

I try to use GPT responsibly. I also love bar trivia.

Contact me about either!

lanetrain@gatech.edu

lanelawley.net