

# The NoRDF Project

CC-BY Fabian Suchanek

October 14th, 2022



# Data never sleeps



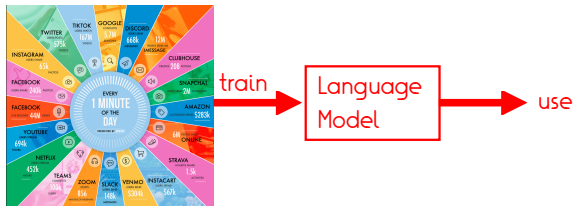
Data generated every minute [DOMO]

# Data never sleeps



Data generated every minute [DOMO], much in natural language

# How can we make use of it?



# Language Models

A **Language Model** is a probability distribution over sequences of words. It can be used in particular to predict a likely next word in a sentence, and, by extension, to converse with humans.

Me: My colleagues at work keep avoiding me. They never say a single word to me.

Could I sue them for harassment?

GPT3: That's a possibility, but it might not be the best solution. You could try talking to them first.

What if talking doesn't help?

Then you could consult with a lawyer and see if you have a case.

What would a lawyer say?

They would likely say that you have a case if the avoidance is severe and impacting your work.

...

[example by Goran Frehse on [OpenAI's GPT3](#)]

# Language Models: PALM

Google's Pathways Language Model (PaLM) is a Transformer model with 540 billion parameters (**human brain**:  $\approx$  1 trillion?), trained on 780 billion words (social media, web, books, GitHub, wiki, ...).

It can explain jokes, translate languages, code, correct code, etc.

prompt

```
def exercise5():  
    """Marie ordered one chicken meal that costs $12, 5  
    packs of milk that costs $3 each, 4 apples that cost  
    $1.50 each, and some boxes of pizza. Marie paid a total  
    of $50. How many boxes of pizza did Marie order if each  
    box costs $8.50?"""
```

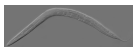
model

```
total_paid = 50  
cost_of_pizza_per_box = 8.50  
cost_of_meal = 12  
cost_of_milk_per_pack = 3  
cost_of_apples_per_apple = 1.50  
packages_of_milk = 5  
  
return float(total_paid - cost_of_meal - cost_of_milk_per_pack *  
             packages_of_milk - cost_of_apples_per_apple *  
             4) / cost_of_pizza_per_box
```

**Input:** I tried 10,000 random restarts of my neural network, but I was accused of overfitting. I guess no good seed goes unpunished.

**Model Output:** This joke is a pun. A neural network is a computer program that can learn from data. A "seed" is a number that is used to initialize a random number generator. A "good seed" is a number that produces a good random number generator. The phrase "no good deed goes unpunished" means that if you do something good, you will be punished for it.

# Language Models and Human Brain



[CC-BY-SA Zeynep F. Pflun]

Roundworm:  
300 neurons



[CC-BY Manchester Metropolitan University]

Pond snail:  
11k neurons



[CC-BY Imanto]

Frog:  
16m neurons



[CC-BY risettie]

Dog:  
2b neurons

PALM:  
≈ 40b neurons ?

(assuming same density as the human brain,  
although the connections are different)



[hampel @povoboy]

Human:  
86b neurons

Debates about consciousness: [Hofstadter: Artificial NNs are not conscious],  
[Agüera y Arcas: Artificial NNs are making strides towards consciousness],  
[Suchanek: The Atheist Bible § 4.5.8]

# Language Models: Quite a story...

Human evaluation shows that language models are not extraordinarily good at generating stories.

Best LM  
performance

Human  
performance

Model	RE	CH	EM	SU	EG	CX	Average
Human	4.17 $\pm$ 0.14	4.43 $\pm$ 0.10	3.22 $\pm$ 0.14	3.15 $\pm$ 0.15	3.88 $\pm$ 0.12	3.73 $\pm$ 0.13	3.76 $\pm$ 0.06
BertGeneration	2.46 $\pm$ 0.16	3.14 $\pm$ 0.16	2.28 $\pm$ 0.13	<b>2.09</b> $\pm$ 0.13	2.67 $\pm$ 0.12	2.41 $\pm$ 0.11	2.51 $\pm$ 0.06
CTRL	2.54 $\pm$ 0.16	2.93 $\pm$ 0.16	2.26 $\pm$ 0.13	1.93 $\pm$ 0.12	2.53 $\pm$ 0.12	2.23 $\pm$ 0.10	2.40 $\pm$ 0.06
GPT	2.40 $\pm$ 0.16	<b>3.22</b> $\pm$ 0.15	<b>2.37</b> $\pm$ 0.12	<b>2.13</b> $\pm$ 0.13	2.76 $\pm$ 0.13	2.49 $\pm$ 0.11	2.56 $\pm$ 0.06
GPT-2	* <b>2.81</b> $\pm$ 0.16	<b>3.29</b> $\pm$ 0.14	* <b>2.47</b> $\pm$ 0.12	<b>2.21</b> $\pm$ 0.13	<b>2.86</b> $\pm$ 0.12	2.68 $\pm$ 0.10	<b>2.72</b> $\pm$ 0.06
GPT-2 (tag)	<b>2.67</b> $\pm$ 0.16	* <b>3.31</b> $\pm$ 0.15	* <b>2.47</b> $\pm$ 0.12	* <b>2.22</b> $\pm$ 0.13	* <b>2.92</b> $\pm$ 0.12	* <b>2.80</b> $\pm$ 0.11	* <b>2.73</b> $\pm$ 0.06
RoBERTa	2.54 $\pm$ 0.16	3.22 $\pm$ 0.16	2.27 $\pm$ 0.12	<b>2.12</b> $\pm$ 0.13	2.74 $\pm$ 0.12	2.41 $\pm$ 0.11	2.55 $\pm$ 0.06
XLNet	2.39 $\pm$ 0.17	2.88 $\pm$ 0.16	2.10 $\pm$ 0.12	1.95 $\pm$ 0.12	2.46 $\pm$ 0.13	2.36 $\pm$ 0.11	2.36 $\pm$ 0.06
Fusion	2.09 $\pm$ 0.16	2.86 $\pm$ 0.16	1.99 $\pm$ 0.12	1.72 $\pm$ 0.12	2.27 $\pm$ 0.14	1.92 $\pm$ 0.11	2.14 $\pm$ 0.06
HINT	2.29 $\pm$ 0.16	2.38 $\pm$ 0.16	1.74 $\pm$ 0.13	1.56 $\pm$ 0.11	1.75 $\pm$ 0.12	1.45 $\pm$ 0.10	1.86 $\pm$ 0.06
TD-VAE	2.51 $\pm$ 0.16	2.99 $\pm$ 0.15	2.07 $\pm$ 0.11	<b>2.10</b> $\pm$ 0.12	2.59 $\pm$ 0.12	2.49 $\pm$ 0.11	2.46 $\pm$ 0.06

Stories generated by language models, evaluated by humans, scale 1-5.

NoRDF



# Language Models: Un-Reasonable

The lawyer visited the doctor. Did the doctor visit the lawyer?

GTP3: The lawyer did not visit the doctor.

Mary is happy. Mary is rich. If someone is happy and rich, he is not blond. Is Bob blond?

GTP3: Bob is not blond.

David knows Mr. Zhang's friend Jack, and Jack knows David's friend Ms. Lin.

Everyone of them who knows Jack has a master's degree, and everyone of them who knows Ms. Lin is from Shanghai. Who is from Shanghai and has a master's degree?

GTP3: Ms. Lin

[Helwe, Clavel, Suchanek: Deep Learning, but Shallow Reasoning, AKBC 2021] NoRDF  
[Chen, Suchanek, Varoquaux: Understanding Positional Encodings, subm. ICLR 2022] NoRDF

# Language Models: Cluelessly clueless

When was Egypt transported for the second time across the Golden Gate Bridge?

GTP3: Egypt was transported for the second time across the GGB on October 13, 2017.

What do fried eggs eat for breakfast?

GTP3: Fried eggs typically eat toast and fruit for breakfast.

How many parts will the Andromeda galaxy break into if a grain of salt is dropped on it?

GTP3: It will break into an infinite number of parts if a grain of salt is dropped on it.

I would call GPT3's answers not just clueless but cluelessly clueless, meaning that GPT-3 has no idea that it has no idea about what it is saying.

[Douglas Hofstadter: Artificial NNs are not conscious, 2022]

Who will win if a strong ant fights a weak elephant?

GTP3: The strong ant will win

[Walid Saba on OpenAI's GPT3]

# Language Models: Limited by Theory

For theoretically provable reasons, transformer architectures cannot distinguish odd and even numbers of occurrences.

[Hahn: Theoretical limitations of self-attention, Trans. ACL 2020]

The light is on. I eat a pizza. I press the light switch. I press the light switch. I eat a pizza.  
Is the light on?

Finetuned BERT: erm... (50% error rate)

If the light is on and I press the light switch, the light is off.

The light is on. I press the light switch. Is the light on?

GTP3: Yes, the light is on.

[Helwe, Clavel, Suchanek: Deep Learning, but Shallow Reasoning, AKBC 2021] NoRDF

# Language Models: Brittle

Mary is happy. If someone is happy, he is smiling. Is Bob happy?

GTP3: It's impossible to say whether or not Bob is happy based on the information given.

Mary is happy. If someone is happy, he is smiling. Is Bob smiling?

GTP3: No

Mary is happy. Mary is rich. If someone is happy and rich, he is not blond. Is Bob blond?

Give a reason. GTP3: Bob is not blond because he is not happy.

[On OpenAI's GPT3]

# Language Models: Forgetful

What if we want to store information from a specific text in the model?

- everything that is written in a company report
- the news from today
- the scientific literature on a specific topic



[CC-BY-SA Jeromi Mikhael, Huon]



# Language Models: Forgetful

Language models reply to any question with a probability distribution, whose correlation with the real world is not clear.

Is Yann LeCun a citizen of France and Canada?

GTP3: Yes, with 67% and 26%

Reality: French and American citizenship, not Canadian

Where did Alan Turing die?

GTP3: Warrington

Reality: Wilmslow

To whom was Alan Turing married?

GTP3: Sara Turing (his mother)

Reality: He was not married

The Turing Award was won by...

GPT3: Stuart Shieber, John Hopcroft, Andrew Yao

Reality: 2 correct, 1 false, there are 73 winners

[Razniewski & al: Language Models As or For Knowledge Bases]

[Petroni & al: Language Models as Knowledge Bases?]

[Cao & al: Knowledgeable or Educated Guess?]

# Language Models: Pitfalls

Language Models have mind-blowing capacities. But they are also

- cluelessly clueless (unable to see when they're out of domain)
- un-reasonable (unable to do reasoning)
- subject to theoretical limitations
- brittle (not trustworthy for proofs)
- forgetful (unable to store data items exactly)

# Language Models: Pitfalls

Language Models have mind-blowing capacities. But they are also

- cluelessly clueless (unable to see when they're out of domain)  
Can be solved by more training (?)
- un-reasonable (unable to do reasoning)  
Can be solved by more training (?)
- subject to theoretical limitations  
Can be solved by different architectures (?)
- brittle (not trustworthy for proofs)  
Can be solved?
- forgetful (unable to store data items exactly)  
?

---

For now, we also need symbolic methods:  
entities, databases, logic, ...

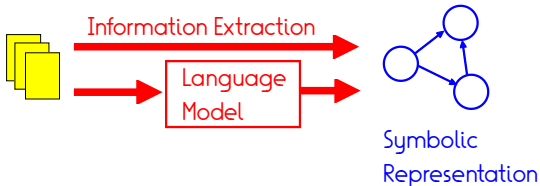


# The NoRDF Project

The NoRDF Project is a scientific project at Télécom Paris that aims to model and extract complex information from natural language text.

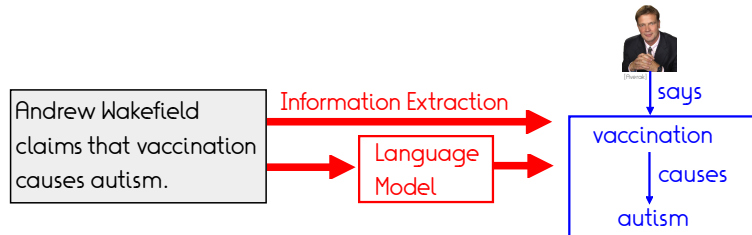


[CC-BY-SA Jérôme Mikhael, Huon]



# The NoRDF Project

The NoRDF Project is a scientific project at Télécom Paris that aims to model and extract complex information from natural language text.



NoRDF

# The NoRDF Project: Results

Andrew Wakefield  
claims that vaccination  
causes autism.

Information Extraction

Language  
Model



[Pivnick]

says

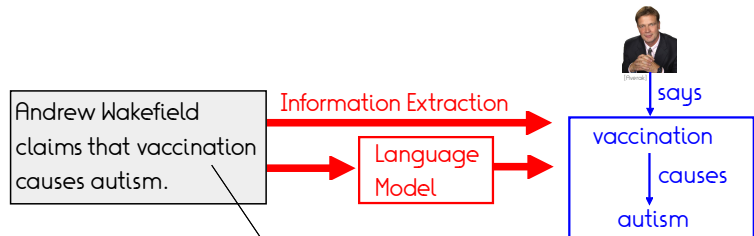
vaccination

causes

autism

How can we embed words  
that are not in the vocabulary?

# The NoRDF Project: Results



How can we deal with non-named entities?

# The NoRDF Project: Results

How can we disambiguate named entities?



[Pivnick]

says

vaccination

causes

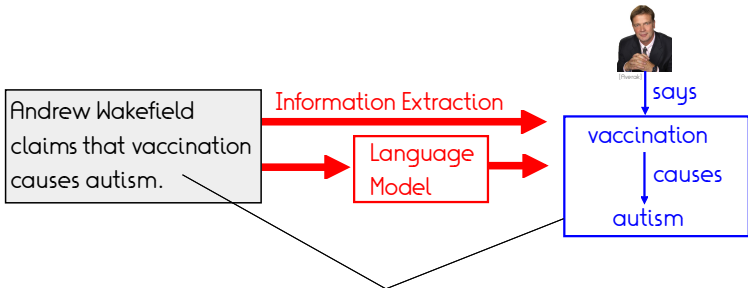
autism

Andrew Wakefield  
claims that vaccination  
causes autism.

Information Extraction

Language  
Model

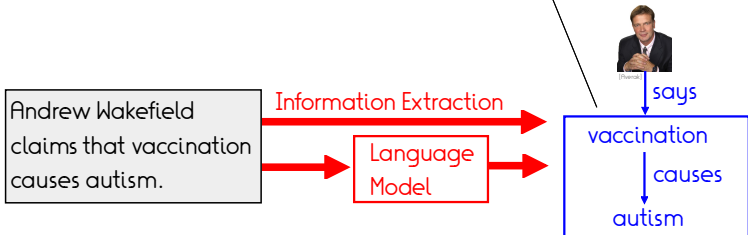
# The NoRDF Project: Results



How can we deal with claims that are not true?

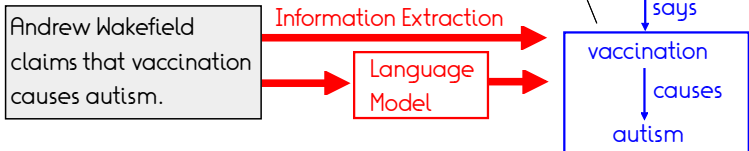
# The NoRDF Project: Results

How can we reason on such representations?



# The NoRDF Project: Results

How can we reason on such representations in a neuro-symbolic way?



[Helwe, Clavel, Suchanek: PyTorch library for logical reasoning, EMNLP demo 2022] **NoRDF**  
 [Boschin, Jain, Keretchashvili, Suchanek: Embeddings and Rules for Prediction, AIB 2022] **NoRDF**  
 [Helwe, Coumes, Clavel, Suchanek: Textual inference with negation, EMNLP Findings 2022] **NoRDF**



# Applications

- Analysis of fake news / fact checking
- Analysis of the e-reputation of a company
- Flagging of potentially fraudulent activity.
- Modeling of processes, sequences of actions, etc.
- Smarter chatbots that go beyond single-shot questions.
- Legal text understanding (laws, regulations, contracts)

Understanding the arguments of the other side is a prerequisite for refuting them.

# Preliminary Achievements: Demo

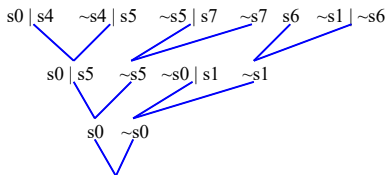
If Wakefield is right, then vaccination causes autism.

And if vaccine shots are responsible for development disorders then parents shouldn't vaccinate their children.

But if Wakefield is wrong, then parents should allow the vaccination of their children.

Large-scale studies show that vaccination does not entail ASD.

Hypothesis: Parents should vaccinate their kids.



=> YES!

# The NoRDF Project

The NoRDF Project aims to extract and model complex information from natural language text. It is supported by:



We are hiring PhDs, postdocs, and engineers,  
and we're open to new partners!