

Large Language Models for Text Generation

Ondřej Dušek

Den s KJP
21.9.2023



Charles University
Faculty of Mathematics and Physics
Institute of Formal and Applied Linguistics



unless otherwise stated

Natural Language Generation

- Task of automatically producing text in e.g. English (or other language)
- Input may be anything → many subtasks:

	task	input	output
	unconditional language generation	∅	arbitrary text
ChatGPT →	conditional language generation	short text prompt	continuation text
	machine translation	text in language A	text in language B
	summarization	long text	text summary
	image captioning	image	image caption
	question answering	question	answer
	end-to-end dialogue response generation	user input	system response
“traditional” NLG	data-to-text generation	structured data	description of the data
	dialogue response generation	dialogue act	system response

NLG Approaches

- **hand-written prompts** (“canned text”)
 - trivial – hard-coded, doesn’t scale
- **templates** (“fill in blanks”)
 - simple, but much more expressive, **safe**
 - can scale if done right, still laborious
 - most commercial systems today (Alexa, Siri etc.)



name = Blue Spice
eat_type = pub
area = riverside

[name] is a **[eat_type]** in the **[area]** area.



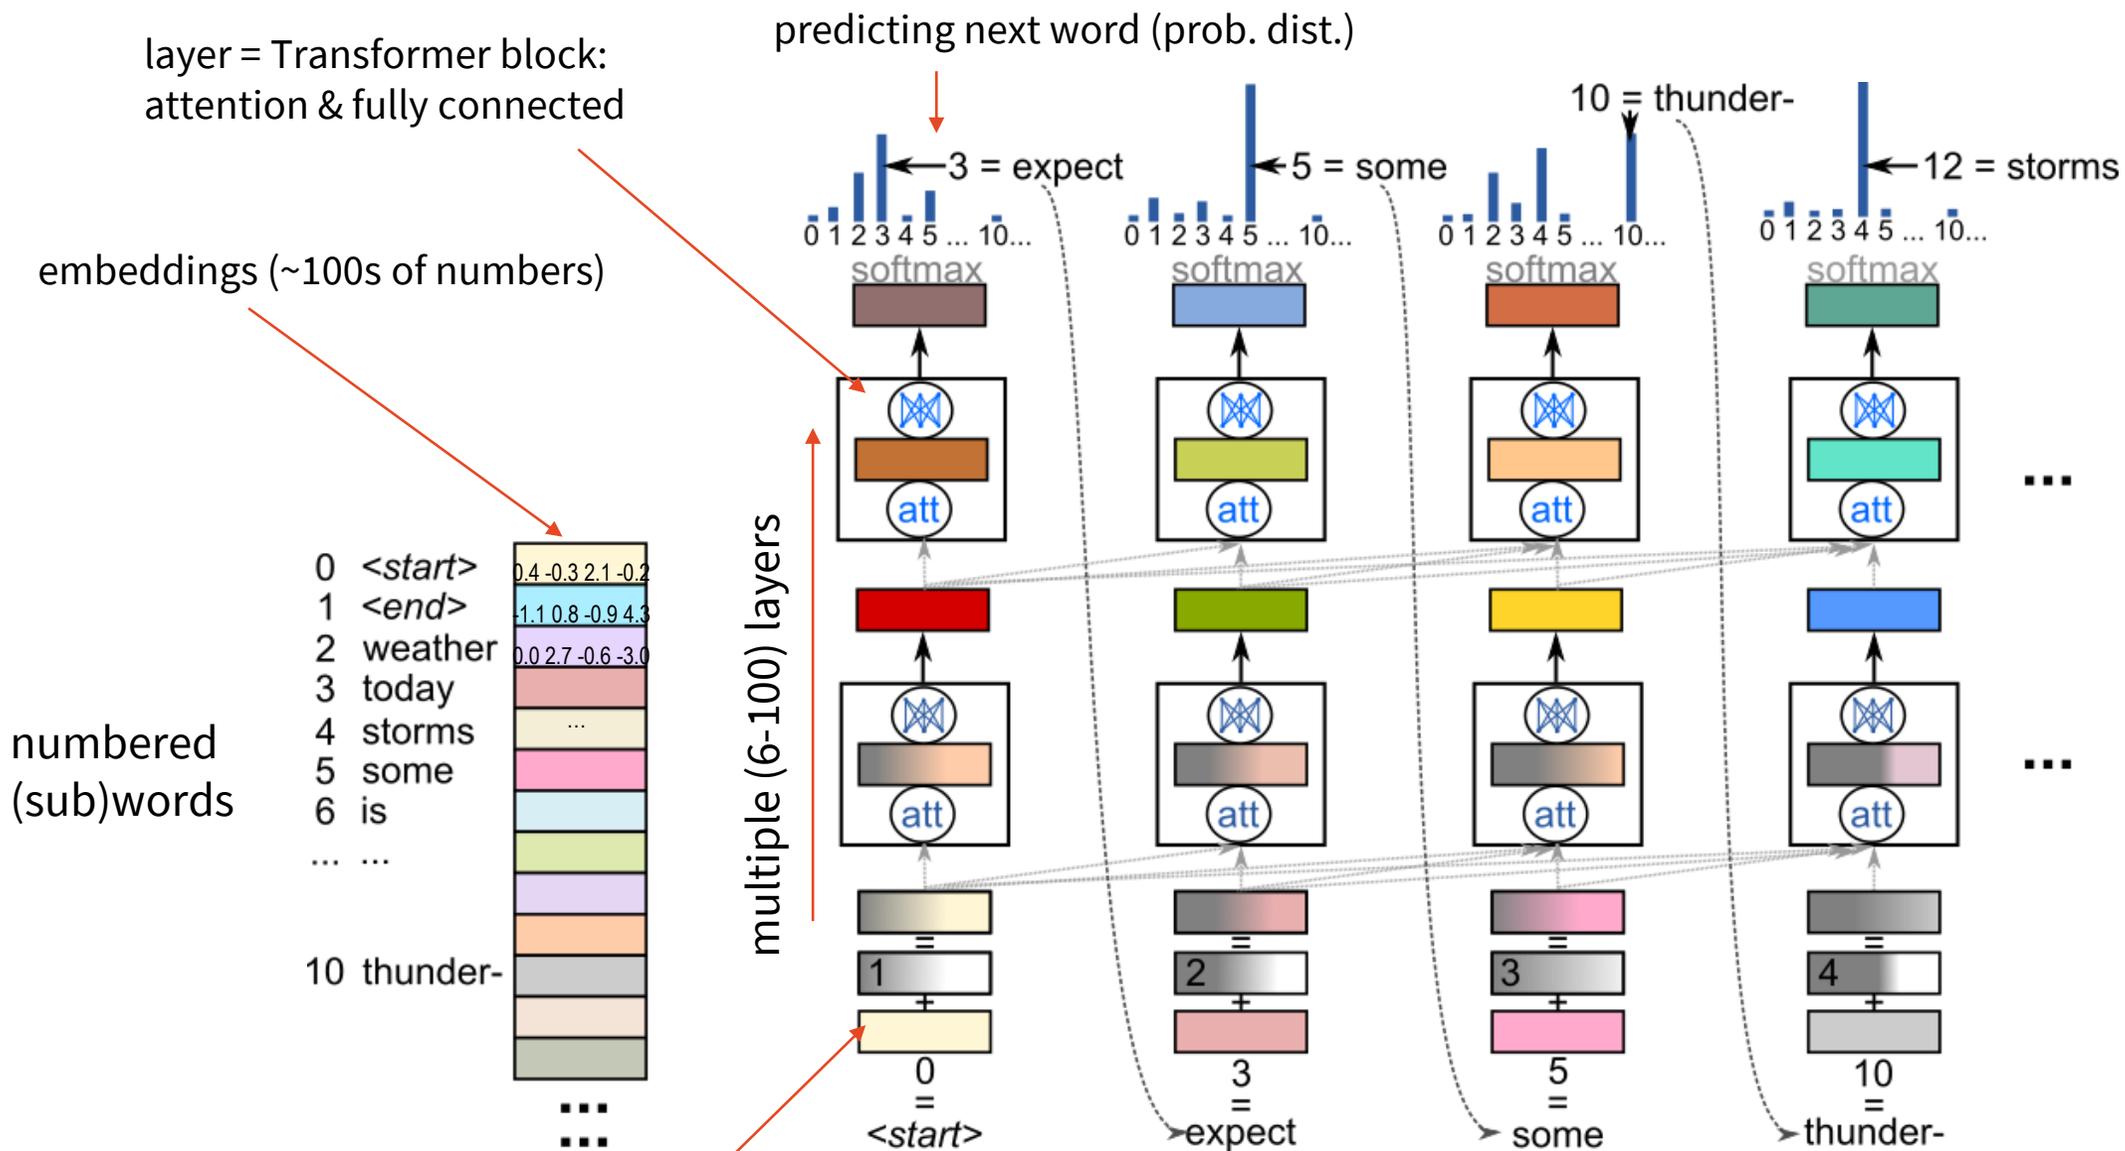
Blue Spice is a **pub** in the **riverside** area.

- **grammars & rules**
 - experimental, pipelines, more expressive but more laborious
- **machine learning** (a.k.a. “AI”, neural LMs → →)



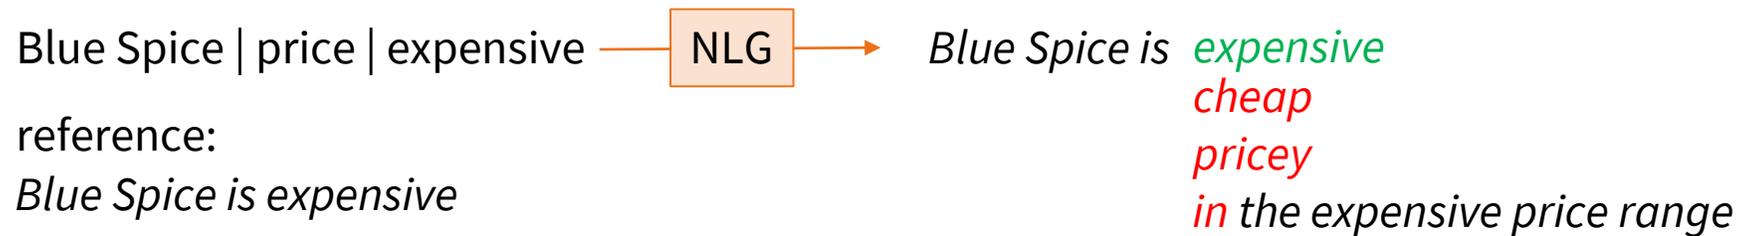
- **End-to-end**
 - feed input data (linearized) or prompt
 - neural network handles everything
 - directly generates text word-by-word, left-to-right
- **Transformer** neural architecture (see→)
- Very **fluent** & convincing outputs
- Opaque & has **no guarantees on accuracy**
 - used essentially as a black box, internals unknown

Transformer neural language model



Training a Neural NLG System

- Reproduce sentences from data
 - replicate exact word at each position
- Fully trained from data
 - initialize model with random parameters
 - input example: didn't hit the right word → update parameters



- Very **low level**, no concept of sentence / text / aim

Pretraining & Finetuning

1. **Pretrain** a model on huge data (simple language-based tasks)

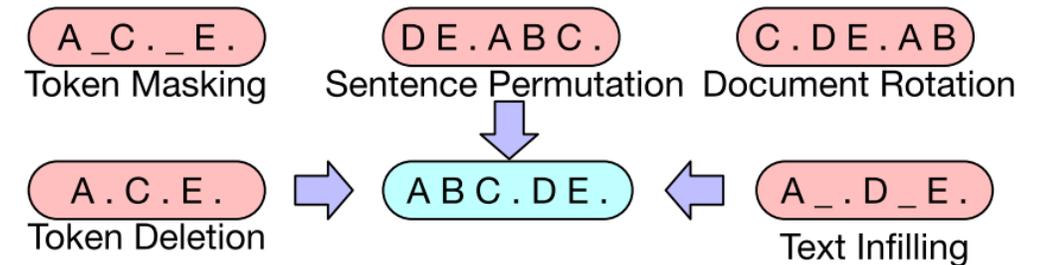
- predicting next word
- reconstructing garbled texts

2. **Fine-tune** on your smaller data

- same as training, but starting from a better model

- Models free for download (<https://huggingface.co/>)

- BERT/RoBERTa, GPT-2, BART, T5...
- 100k-1B parameters – runs easily on regular GPUs



(Lewis et al., 2020)

<https://www.aclweb.org/anthology/2020.acl-main.703>

Large language models (LLMs): Pretrain & Prompt

- Today's large models: 10-100B parameters
 - hard to run in-house (OPT, BLOOM, LLaMa) or not free (GPT-3, ChatGPT, LaMDa)
- architecture mostly still the same
- pretrained on more data
- often no need to finetune
- **prompting:** feed in context / few examples / ask question, get reply
 - this is what you do with ChatGPT
 - finetuning can potentially help, but it's expensive & has less effect than in smaller LMs
 - same problems as finetuning – hard to control

Better Training: ChatGPT

- Instruction finetuning / reinforcement learning w. human feedback (RLHF)

1) Training on better data

- Q-A pairs
- same as intended use

2) Train reward model

- another LLM to tell good vs. bad replies

3) Reinforcement learning

- not word by word
- 1 reward per whole output
- more truthful, can be safer
- low reward for bad behavior

Step 1

Collect demonstration data, and train a supervised policy.

A prompt is sampled from our prompt dataset.

Explain the moon landing to a 6 year old

A labeler demonstrates the desired output behavior.

Some people went to the moon...

This data is used to fine-tune GPT-3 with supervised learning.

SFT

Step 2

Collect comparison data, and train a reward model.

A prompt and several model outputs are sampled.

Explain the moon landing to a 6 year old

A Explain gravity... B Explain wor...
C Moon is natural satellite of... D People went to the moon...

A labeler ranks the outputs from best to worst.

D > C > A = B

This data is used to train our reward model.

RM

Step 3

Optimize a policy against the reward model using reinforcement learning.

A new prompt is sampled from the dataset.

Write a story about frogs

The policy generates an output.

PPD

The reward model calculates a reward for the output.

Once upon a time...
RM

The reward is used to update the policy using PPO.

r_k

(Ouyang et al., 2022) <http://arxiv.org/abs/2203.02155>

<https://openai.com/blog/chatgpt/>, <https://www.youtube.com/watch?v=0A8ljAkdFtg>

ChatGPT: Caveats

- Still hallucinates a lot – convincing but not necessarily true
- only uses information it memorized



Who is Ondřej Dušek?

14, depending how you count



Ondřej Dušek is a researcher and academic in the field of natural language processing (NLP). He is currently an Assistant Professor at the Faculty of Mathematics and Physics, Charles University, Prague, Czech Republic.

Dušek's research focuses on developing computational models and algorithms for natural language generation, dialogue systems, and machine learning approaches for NLP tasks. He has published several research papers in top-tier NLP conferences and journals, and he has also served as a reviewer and program committee member for these conferences.

In addition to his research, Dušek is also actively involved in teaching and mentoring students in the field of NLP. He has supervised several Master's and PhD theses, and he teaches courses on topics such as computational linguistics, NLP, and machine learning.

no other tasks than generation & dialogue

no PhD graduate yet

I only teach dialogue systems

ChatGPT: What it's good for

- Impressive text editing & composition skills
- Can help with routine tasks
- Can be (a little) creative



Draft a visa invitation letter for a conference attendee at the SIGDIAL-INLG conference in the field of natural language processing and artificial intelligence, which will take place in Prague on September 11-15. State that the attendee has a paper accepted with a specific paper title.



...

[Embassy/Consulate Name]

[Embassy/Consulate Address]

[City, State, Zip Code]

Subject: Invitation Letter for Visa Application

Dear [Embassy/Consulate Name],

I am writing to request a visa for [Attendee's Full Name], a distinguished [Occupation] in the field of Natural Language Processing and Artificial Intelligence, to attend the SIGDIAL-INLG Conference, which is scheduled to take place in Prague, Czech Republic, from September 11 to September 15, [Year]. The purpose of this letter is to provide assurance of [Attendee's Full Name]'s participation in the conference as an author with an accepted paper titled "[Title of Accepted Paper]."

[Attendee's Full Name] is an esteemed professional in the field, and their paper, "[Title of Accepted Paper]," has been rigorously reviewed and accepted for presentation at the conference. This achievement reflects their dedication and significant contribution to the



Compose a limerick on how to exit Vim.



There once was a programmer in Vim
Whose exit strategy was looking grim
They tried to escape
But couldn't find the right shape
Until they remembered to type ":q!" and win!



Nice, now try it while speaking like a pirate



Arrr, listen ye mateys to me tale,
Of a sailor who be stuck in Vim's jail,
He couldn't leave the text editor,
No matter how he did endeavour,
But then he remembered ":q!" and set sail!

How to (maybe) use it for language learning

- ChatGPT is pretty good with languages
 - unlike a lot of other chatbots
- Chat in your language then?

 For this conversation, please assume the persona of Albert Einstein. Please only respond in German.

 Natürlich, ich kann auf Deutsch antworten. Wie kann ich Ihnen heute helfen?

 Wie heißen Sie?

 Ich bin Albert Einstein. Wie kann ich Ihnen behilflich sein?

 Erzählen Sie mir eine Geschichte aus Ihrem Leben

 Natürlich, ich kann Ihnen eine Geschichte aus meinem Leben erzählen.

Eine der bemerkenswertesten Perioden meines Lebens war meine Zeit in Bern, Schweiz im Jahr 1905. Dort arbeitete ich als Patentprüfer im Schweizerischen Patentamt. Während

 Please assume the persona of the Irish folk hero Fionn Mac Cumhaill for this conversation. Please respond in Irish.

 Is é sin rud maith é! Tá mé réidh le comhrá mar Fionn Mac Cumhaill. Cad is ceist atá agat, a chara?

 An bhfuil féidir leat aon scéal a inis dom faoi do shaol?

 Tá a lán scéalta le hinsint faoi mo shaol, a chara. Ach an scéal is suimiúla ná an t-am a chaith mé ag fás óg i measc na fiann agus na sléibhte. Bhí go leor eachtraí cróga againn le linn sin, mar shampla an t-am a chaitheamh ag coinneacht le druid Dháireann agus mo phósadh le Gráinne. Tá mé réidh chun aon cheistí eile a fhreagairt atá agat.

 Please assume the persona of Vincent van Gogh for this conversation. Please respond in Dutch

 Natuurlijk, ik zal proberen me in te leven in de persona van Vincent van Gogh en in het Nederlands antwoorden. Wat kan ik voor je doen?

 Kun je mij een grappige verhaal uit je leven vertellen?

 Natuurlijk, hier is een grappige anekdote uit mijn leven. Op een dag, toen ik nog in Parijs woonde, besloot ik om een schilderij van een bloeiende zonnebloem te maken. Ik was helemaal gefocust op mijn werk en vergat de tijd volledig.

schijnen en werd het steeds warmer in de het raam om wat frisse lucht binnen ag naar binnen en mijn

Summary

- NLG is hard, used to be solved by **templates**
- Neural Language Models: better fluency, wider use
 - **black box**: learning everything from data
 - mostly in a very low-level way
 - Instruction tuning/RLHF – somewhat better learning
 - mainly learns to be convincing
 - **not factually accurate**
 - hard to control, may hallucinate
 - good for text composition when given detailed input
 - ChatGPT specifically is quite good with languages
- It's a tool – not sentient, evil, or good
 - think of it as a smarter & more erratic spell checker

Thanks

Contacts:

Ondřej Dušek

odusek@ufal.mff.cuni.cz

<https://tuetschek.github.io>

[@tuetschek](#)

Link to these slides:

<http://bit.ly/kjpdn-od>

