Welcome!

Jindřich Helcl

22 February, 2024

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

Goals of the Course

1. Show you how the models work

2. Teach you to use them

3. Make you able to critically assess what you read about them

4. Think about the broader context of using the models

About the Course

Part 1 - How the models work

- The Transformer architecture
  - Start with high-level overview, but
  - in the end, enough details for you to be able to implement it from scratch.
- Model training
  - Math formulas for different training objectives,
  - implementation in Pytorch.
- Data processing
  - Where to get appropriate data,
  - how to tell whether it is good or bad.
- Analyzing trained models
  - Evaluate model on a given task,
  - but also design methods to learn about its inner workings.

About the Course

Part 2 - How to use LLMs

- Fine-tuning to specific tasks
  - Learn about the tasks
  - Select a pretrained model and a suitable dataset
  - Write and run fine-tuning scripts on GPUs
- Optimizing for low computational budget
  - What to do to fit a huge model on a medium-sized device
- Accessing LLMs using an API
  - .. and figuring out the best prompt
  - (.. and deciding whether it is worth to do so)
- Knowing when LLMs fail
  - Know your data
  - Assess deployment risks

About the Course

Part 3 - Getting information

- Reading research papers
  - Appreciate the strengths of the good ones
  - Identify weaknesses in all of them
- Reading news stories / blog posts
  - Know what is realistic and what is hype

Part 4 - Using LLMs in context

- Ethics of LLM use and training data acquisition
- Do LLMs understand language?
- Societal impacts

About Us

Jindřich Helcl

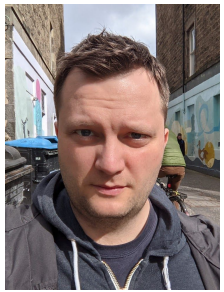

Jindřich Libovický

Ondřej Dušek

Tomasz Limisiewicz

Zdeněk Kasner

Rudolf Rosa

Dominik Macháček

Tomáš Musil

David Mareček

Ondřej Plátek

Michal Novák

Martin Popel

Our emails: surname@ufal.mff.cuni.cz

Course Logistics

- Thursdays, 15:40 (that's now)
- Room **S3** (changed after capacity issues with S8)
- 13 sessions led by various teachers (no class next week!)
- Credit at the end (worth 3 e-credits)

Formal Requirements

To pass the course, you need to:

- Actively participate during classes
  - There will be 2-3 tasks during the semester; we will work on them mainly during classes but they might turn into a (small) homework
  - You will be asked at least once to read a paper before the class
- Take part in the final test

Syllabus

- Introduction to the Transformer model
- Language model types
- Data gathering and curation, downstream tasks
- Model training
- Fine-tuning and inference
- Multilinguality and cross-lingual transfer
- LLM Applications
- Multimodality
- Societal impacts
- Interpretability

Summary

- The course materials will be available on the course webpage.
- If you have questions, drop us a line:
  - Jindřich Helcl (helcl@ufal.mff.cuni.cz)
  - Jindřich Libovický (libovicky@ufal.mff.cuni.cz)

https://ufal.cz/courses/npfl140

Warm-up Discussion

Discussion

What do we mean by **Language Models**?

Discussion

What do we mean by **Large** Language Models?

Discussion

What can we use LLMs **for, and how**?

Discussion

What **types** of language models are there?

Discussion

Why do you think LLMs are so **good**?

Discussion

What are **downsides** of using LLMs?

Discussion

How good is ChatGPT in your native language? Why is that?

Discussion

(#1 question on genai.stackexchange.com)

Does saying "please" and "thank you" to LLMs change anything?