

Large Language Models and Natural Language Processing

Course code: MLC_NLP

This course is intended for anyone fascinated by the capabilities of large language models and generative artificial intelligence who wants to understand this field beyond the level of an ordinary user. Together, we will explore transformers—the fundamental building blocks of modern language models—introduce the most well-known architectures, and demonstrate how large language models can be used in various applications. No paid third-party accounts are required for the practical exercises. We will use open-source models which, when used correctly, can be just as good as the largest commercial ones.

Who is the course for

This course is intended for anyone who is fascinated by the capabilities of large language models and generative artificial intelligence, and who wants to delve into this field beyond the level of an ordinary user.

What we teach you

Together, we will explore transformers—the fundamental building blocks of modern language models—introduce the most well-known architectures, and demonstrate how large language models can be used for various applications. No paid third-party accounts are required for the practical exercises. We will use open-source models which, when used correctly, can be just as good as the largest commercial ones.

Required skills

- Basic knowledge of programming in Python.
- Knowledge of machine learning at the level of an Introduction to Machine Learning course.

Course outline

- Generative AI for text and images
- Evolution of language modeling
- Transformers
- Types of transformers for language modeling (encoder, decoder, encoder–decoder)
- Reinforcement learning from human feedback (RLHF)
- Selected transformer-based language models (BERT, GPT, LLaMA, T5, BART...)
- Practical example of text classification using transformers with the HuggingFace library in Google Colab
- Prompt engineering: in-context learning, zero-shot, one-shot and few-shot prompting, key configuration parameters of generative processes
- Practical example of in-context learning using the HuggingFace library in Google Colab
- Fine-tuning large language models and parameter-efficient fine-tuning (LoRA)
- Evaluation of generative language models (ROUGE, BLEU)
- Practical example of parameter-efficient fine-tuning using the HuggingFace library in Google Colab
- Retrieval-Augmented Generation (RAG)

GOPAS Praha

Na Strži 2097/63
140 00 Praha 4 - Krč
Tel.: +420 226 201 390
info@gopas.cz

GOPAS Brno

Nové sady 996/25
602 00 Brno
Tel.: +420 542 422 111
info@gopas.cz

GOPAS Bratislava

Dr. Vladimíra Clementisa 10
Bratislava, 821 02
Tel.: +421 902 903 132
info@gopas.sk



Copyright © 2026 GOPAS, a.s.,
All rights reserved