Szkolenie: Component Soft

AI-434 GenAI Application Development with LLMs (OpenAI GPT, Google Gemini, Meta Llama, Mistral)

DOSTĘPNE TERMINY

2026-02-24 | 4 dni | Kraków / Virtual Classroom
2026-02-24 | 4 dni | Warszawa / Virtual Classroom
2026-03-24 | 4 dni | Kraków / Virtual Classroom
2026-03-24 | 4 dni | Warszawa / Virtual Classroom
2026-03-31 | 4 dni | Kraków / Virtual Classroom
2026-03-31 | 4 dni | Warszawa / Virtual Classroom
2026-04-27 | 4 dni | Kraków / Virtual Classroom
2026-04-27 | 4 dni | Warszawa / Virtual Classroom
2026-05-26 | 4 dni | Kraków / Virtual Classroom
2026-05-26 | 4 dni | Warszawa / Virtual Classroom
2026-06-30 | 4 dni | Kraków / Virtual Classroom
2026-06-30 | 4 dni | Warszawa / Virtual Classroom

Cel szkolenia:

Artificial intelligence has become an extremely important area for IT professionals and engineers with the scientific breakthroughs and practical applications of generative AI systems, especially its Large Language Model (LLM) variant such as OpenAI's GPT, Google's Gemini and many other closed- and open-source models. Due to its importance and impact on every aspect of our lives, understanding the concepts, functionalities and practical usage of generative AI systems is quickly becoming essential for all IT and other technical professionals as well as for managers with technical background.

This training focuses on LLM concepts as well as GPT, Gemini and open-source LLM prompt engineering and application development, and teaches participants the following

topics:

- Introduction to LLM based applications

- The Foundation Technologies of LLMs (Neural Networks, Tokenizer, Transformer)

- The 3-phase training process of LLMs (pre-training, fine-tuning, RLHF)

- Using closed- and open-source LLMs via APIs

- Prompt engineering

- Retriever Augmented Generation (RAG)
- Creating LLM chains with LangChain
- LLM Agents
- Fast Web Interface Prototyping for LLMs (Gradio)
- Debugging and Evaluating LLM-based apps (Langsmith)
- Fine-tuning open-source LLM models

Besides gaining a basic understanding of the concepts of prompt engineering, students will also do extensive lab exercises using the Python APIs of the OpenAI GPT, Google Gemini as well as popular and powerful open-source LLMs such as Meta's Llama and Mistral models to see how these concepts work in practice.

This training is part of the AI portfolio of Component Soft which explores essential AI topics, such as:

- AI-110 Intro to GenAI with Large Language Model (LLMs) and LLM-based apps
- AI-434 GenAI Application Development with LLMs (OpenAI GPT, Google Gemini, Meta
- Llama, Mistral)

Structure: 50% theory, 50% hands on lab exercises.

Target audience: Software developers and other IT and technical professionals as well as managers with technical background who want to understand the basic concepts and technologies behind Large Language Models (LLMs) and want to gain practical skills in prompt engineering and LLM application development with the Python APIs of popular closed- and open-source LLMs (OpenAI GPT, Google Gemini, Meta Llama, Mistral).

Plan szkolenia:

- PART I. Basic LLM Concepts (~6 hours)
  - Module 1. Introduction to LLM based applications
    - Main usage areas of LLM-based applications
    - Main types of LLM-based applications
    - Building blocks of LLM-based applications
    - Lab: Testing a few simple LLM-based applications
  - Module 2. The Foundations: Neural Networks, Deep Learning, Transfer Learning

- From human neural cells to artificial neural networks
- Deep Neural Networks
- Transfer Learning
- Lab: Examining a small neural network model written in Python and PyTorch

- Module 3. "Attention is all you need" – The Transformer Architecture
  - Intuition of the transformer model
  - Main elements of transformers: tokenizer, embeddings, encoder, decoder
  - Variations on the transformer architecture
  - Popular transformer models
  - Lab: Testing text generation of different GPT model generations

- Module 4. The 3-phase training process of LLMs (Pre-training, Fine-tuning, RLHF)
  - Pre-training of LLMs
  - How does pre-training basically work?
  - Training data set, computational and financial challenges
  - LLM Fine-tuning techniques.
  - How does fine-tuning basically work?
  - Parameter efficient fine-tuning (PEFT) with LoRA and quantized parameters
  - Reinforcement Learning with Human feedback (RLHF)
  - Why do we need RLHF in the first place?
  - Methods and main steps of RLHF
  - Lab: Examining GPT 3 models before and after fine-tuning and RLHF

- PART II. Application Development with LLMs
  - Module 5. Using closed- and open-source LLMs via APIs
    - Using LLMs through APIs
    - Typical LLM parameters
    - Jupyter Lab basics
    - Lab: Doing basic LLM tasks with Jupyter lab notebooks. Using GPT, Gemini as well as popular open-source LLMs via the Python APIs

  - Module 6. Prompt engineering
    - Prompt engineering basics:
    - What is prompt engineering?
    - Prompt engineering terminology and concepts
    - The "Just Ask" Principle, Zero-shot prompts
    - Prompts with Few-shot learning
    - Prompt Chaining, Chain of Thought Prompting
    - Prompts with Personas

- Advanced prompt engineering techniques:
- Techniques related to Prompt Structure and Clarity
- Techniques related to Specificity and Information
- Techniques related to User Interaction and Engagement
- Techniques related to Content and Language Style
- Techniques related to Complex Tasks and Coding Prompts
- Lab: Practicing prompt techniques via the Python APIs of GPT, Gemini, as well as popular open-source LLMs

- Module 7. Retriever Augmented Generation (RAG)
  - What is Retriever Augmented Generation (RAG)?
  - How does RAG work?
  - Syntactic vs. Semantic Similarity
  - Text embedding
  - Vector Databases
  - Lab: Creating simple RAG systems with GPT, Gemini as well as popular opensource LLMs.

- Module 8. Creating LLM chains with LangChain
  - What are LLM chains?
  - LangChain architecture
  - Main Building Blocks: Models, Prompts and Output Parsers
  - Building LLM chains from building blocks
  - LangChain Memory
  - Lab: Using LangChain together with GPT, Gemini as well as popular opensource LLMs.

- Module 9. LLM Agents with LangChain
  - Motivations for LLM Agents
  - Main Features of LLM Agents
  - Main Building Blocks: Functions, Tools, Agents, ReAct execution logic
  - Implementing ReAct with complex prompts and with function-calling LLMs
  - Different ways of creating agents in Langchain
  - Limitations of the ReAct models, possible new directions
  - Lab: Creating and using LangChain agents with GPT, Gemini as well as popular open-source LLMs.

- Module 10. Fast Web Interface Prototyping for LLMs (Gradio)
  - Main features of LLM Web interfaces
  - Example web app for simple LLM tasks
  - Creating our own chatbot with a web interface

- Lab: Creating a simple LLM web interface with Gradio for GPT, Gemini as well as popular open-source LLM-based apps.
- Module 11. Tracing and Evaluating LLM-based apps (Langsmith)
  - Main techniques of debugging LLMs
  - Popular LLM tracing tools (Langsmith, Langfuse, Weight & Biases)
  - Tracing features of Langsmith
  - Challenges of evaluating LLMs
  - Evaluating the quality of prompts and answers with different LLMs
  - Testing semantic searches
  - Explicit and implicit human evaluation of LLMs
  - Lab: Using Langsmith for debugging and evaluating GPT, Gemini as well as popular open-source LLM-based apps.
- Module 12. Fine-tuning Open-source LLM models
  - Typical goals of fine-tuning LLMs
  - Difference between prompt engineering and fine tuning
  - When to use which?
  - The HuggingFace trainer API
  - Preparing your dataset
  - Creating a fine-tuned model
  - Evaluating a fine-tuned model
  - Lab: Fine-tuning an open-source LLM

Wymagania:

Prerequisites: Basic understanding of AI concepts, basic Python programming skills, user experience with ChatGPT or similar chatbots.

Poziom trudności

Certyfikaty:

The participants will obtain certificates signed by Component Soft (course completion).

Prowadzący:

Certified Component Soft Trainer.