# COURSE DESCRIPTION

*Understanding and Developing Large Language Models (LLMs)*

University year **2026-2027**

**1. Programme-related data**

| | | |
|---|---|---|
| 1.1. Higher Education Institution | *Babeş-Bolyai University* | |
| 1.2. Faculty | *Faculty of Mathematics and Computer Science* | |
| 1.3. Department | *Department of Computer Science* | |
| 1.4. Field | *Computer Science* | |
| 1.5. Level of study | *Bachelor* | |
| 1.6. Degree programme / Qualification | *Computer Science* | |
| 1.7. Form of education | *Full time* | |

**2. Course-related data**

| | | | | | |
|---|---|---|---|---|---|
| 2.1. Course title | *Understanding and Developing Large Language Models (LLMs)* | | | Course code | *MLE5247* |
| 2.2. Course coordinator | Lect. Dr. Bogdan MURSA | | | | |
| 2.3. Seminar coordinator | Lect. Dr. Bogdan MURSA | | | | |
| 2.4. Year of study | 3 | 2.5. Semester | 5 | 2.6. Type of assessment | Viva voce |
| 2.7. Course status | Optional | | | 2.8. Course type | Specialisation subject |

**3. Total estimated time** (hours per semester of teaching activities)

| | | | | | |
|---|---|---|---|---|---|
| 3.1. Number of hours per week | **5** | of which: 3.2. course | **2** | 3.3. seminar/ laboratory/ project | **1 lab + 2 proj** |
| 3.4. Total of hours in the curriculum | 60 | of which: 3.5. course | 24 | 3.6. seminar/ laboratory | **36** |
| **Time allocation for individual study (IS) and self-taught activities (ST)** | | | | | **hours** |
| Learning from textbooks, course materials, bibliography, and notes (IS) | | | | | 12 |
| Additional research in the library, on subject-specific electronic platforms, and on-site | | | | | 16 |
| Preparing seminars/ laboratories/ projects, assignments, reports, portfolios, and essays | | | | | 25 |
| Tutoring (professional guidance) | | | | | 6 |
| Examinations | | | | | 6 |
| Other activities | | | | | |
| **3.7. Total hours of individual study (IS) and self-taught activities (ST)** | | | | **65** | |
| **3.8. Total hours per semester** | | | | **125** | |
| **3.9. Number of credits** | | | | **5** | |

**4. Prerequisites** (where applicable)

| | |
|---|---|
| 4.1. curriculum-related | • *Python programming*<br>• *Linear Algebra*<br>• *Statistics*<br>• *Data Structures and Algorithms* |
| 4.2 skills-related | • *Average programming skills in a high-level programming language and very good knowledge on data structures and algorithms.* |

**5. Specific conditions** (where applicable)

| | |
|---|---|
| 5.1. course-related | • *Classroom with a video project device* |

| | | |
|---|---|---|
| 5.2. seminar/laboratory-related | • *Lab equipped with high-performance computers and Python installed.* | |

**6.1. Competencies resulting from the completion of the degree programme *(as referred to in the curriculum)*[1]**

| **Professional competencies** | |
|---|---|
| **Competency code** | **Competency** |
| **PC1** | create software |
| **PC18** | use machine learning |
| **PC6** | develop software prototype |
| **Transversal competencies** | |
| **Competency code** | **Competency** |
| **TC2** | Solve problems |
| **TC3** | Think analytically |

**6.2. Learning outcomes relevant to the degree programme** (as referred to in the curriculum)[2]

| **Learning outcomes targeted by the subject** | | |
|---|---|---|
| **Competency code** | **Knowledge and comprehension** | **Specific academic skills** |
| **PC1/PC6** | The student/graduate identifies, explains and justifies fundamental concepts of data structures, algorithms, and programming paradigms, as well as computer architecture.<br>The graduate knows, understands and applies the basic concepts and the fundamental algorithms of Artificial Intelligence and is able to evaluate them based on metrics.<br>The graduate knows and understands the concepts and the techniques of knowledge representation and is able to apply them for problem solving.<br>The graduate knows and understands the mathematical foundations needed to develop intelligent algorithms and is capable of using them for algorithm implementation.<br>The graduate knows, understands and uses methods for representing, analyzing and handling large volumes of data. | The student/graduate designs, develops and demonstrates complex software solutions using efficient algorithms and diverse programming paradigms. The graduate is able to formally describe issues addressed in various areas, and to model them as problems that can be addressed using Artificial Intelligence techniques. The graduate is able to apply fundamental algorithms of Artificial Intelligence in order to solve real-world problems. The graduate is able to evaluate, both quantitatively and qualitatively, the performance of intelligent systems. The graduate is able to design and implement software systems that are using methods of Artificial Intelligence and to evaluate their performance. |
| **PC17** | The student/graduate identifies, compares, recognizes and describes advanced concepts and techniques in the field of artificial intelligence, machine learning, and natural language processing. | The student/graduate designs, implements, experiments with predictive models and develops applications based on machine learning algorithms. |

[1] The professional and/or transversal skills targeted by the subject for which the course description is prepared will be copied from the curriculum of the degree programme. For each competency, the complete entry, including the competency code, will be copied with the exact wording that appears in the curriculum, without any changes. If no competency is copied from either of the two categories, the row corresponding to that category is deleted from the table.

[2] The learning outcomes relevant for the degree programme and targeted by the subject for which the course description is prepared will be listed. The entries, copied without any changes from the Curriculum by subject type (Core Subject/Specialisation Subject/Complementary Subject), are listed under the corresponding competency.

| | | |
|---|---|---|
| **TC2/TC3** | The student/graduate has the knowledge necessary to understand and solve complex problems, and to plan and organize advanced processes in various fields. | The graduate is able to identify complex problems and examine related issues to develop solving options and implement solutions. The graduate has the ability to apply general rules to specific problems and produce relevant solutions. The graduate is able to combine diverse information to formulate solutions and generate ideas for developing new products and applications. |

**7. Subject-specific learning outcomes**

| **Knowledge and comprehension** |
|---|
| 1. Understanding the fundamental and advanced concepts of Artificial Intelligence and Machine Learning in the context of connected systems. |
| 2. Understanding the main machine learning paradigms, including supervised learning, unsupervised learning, and deep learning techniques. |
| 3. Understanding the stages involved in designing, training, evaluating, and optimizing Machine Learning models. |
| 4. Understanding the challenges related to scalability, data processing, and the integration of AI solutions within Big Data and distributed system architectures. |
| **Specific academic skills** |
| 1. Developing and evaluating Machine Learning models using real-world datasets and modern AI frameworks. |
| 2. Performing data analysis, preprocessing, and visualization in order to build robust and efficient intelligent solutions. |
| 3. Designing and implementing Artificial Intelligence applications for solving real-world problems in the domain of connected systems. |

**8. Contents**

| 8.1 Course | Teaching methods | Remarks |
|---|---|---|
| *1. Introduction to LLMs and the Landscape of Generative AI. Overview of the history of Natural Language Processing with a focus on Large Language Models (LLMs) and their significance in the field of generative artificial intelligence. Examination of various applications and tasks LLMs are employed for, highlighting their versatility.* | • *Interactive exposure*<br>• *Explanation*<br>• *Conversation*<br>• *Didactical demonstration* | |
| *2. The Evolution of Text Generation Technologies. Tracing the development of text generation from pre-transformer models to current methodologies.* | • *Interactive exposure*<br>• *Explanation*<br>• *Conversation*<br>• *Didactical demonstration* | |
| *3. Deep Dive into Transformer Architecture. Techniques and strategies for utilizing transformers in text generation tasks. Exploration of transformer architecture, the backbone of modern LLMs.* | • *Interactive exposure*<br>• *Explanation*<br>• *Conversation*<br>• *Didactical demonstration* | |

| | | |
|---|---|---|
| *4. The Principle of Attention in Transformers. Understanding the "Attention is all you need" concept and its revolutionary impact on LLMs.* | *• Interactive exposure*<br>*• Explanation*<br>*• Conversation*<br>*• Didactical demonstration* | |
| *5. Mastering Prompt Engineering. Learning how to effectively design prompts to guide LLMs in generating desired outputs.* | *• Interactive exposure*<br>*• Explanation*<br>*• Conversation*<br>*• Didactical demonstration* | |
| *6. Pre-Training Large Language Models and Scaling Laws. Insights into the pre-training process, computational challenges, and the principles of scaling laws for LLMs.* | *• Interactive exposure*<br>*• Explanation*<br>*• Conversation*<br>*• Didactical demonstration* | |
| *7. Fine-Tuning LLMs for Specific Tasks. Strategies for instruction-based fine-tuning, including single and multi-task adaptations.* | *• Interactive exposure*<br>*• Explanation*<br>*• Conversation*<br>*• Didactical demonstration* | |
| *8. Advanced Fine-Tuning Techniques. Introduction to Parameter Efficient Fine-Tuning (PEFT) methods such as LoRA and Soft Prompts.* | *• Interactive exposure*<br>*• Explanation*<br>*• Conversation*<br>*• Didactical demonstration* | |
| *9. Reinforcement Learning from Human Feedback (RLHF). Fundamentals of aligning LLMs with human values through RLHF,*<br>*including feedback collection and reward models.* | *• Interactive exposure*<br>*• Explanation*<br>*• Conversation*<br>*• Didactical demonstration* | |
| *10. Enhancing LLM output using Reasoning*<br>*and Act. Explore the landscape of advanced fine-tuning and prompting strategies through method like Chain-of-thought (CoT, Reason Only), Act-only and ReAct across different domains, highlighting their task-solving trajectories and the distinct advantages of the ReAct approach.* | *• Interactive exposure*<br>*• Explanation*<br>*• Conversation*<br>*• Didactical demonstration* | |
| *11. Implementing LLMs in Real-World Applications & Introduction to LangChain. Combining the exploration of deploying LLMs in real-world applications with an introduction to LangChain, covering document loading, vector stores, embeddings, and the fundamentals of Retrieval Augmented Generation (RAG).* | *• Interactive exposure*<br>*• Explanation*<br>*• Conversation*<br>*• Didactical demonstration* | |
| *12. Ethics of AI. Discover the evolving field of generative AI, emphasizing the* | *• Interactive exposure*<br>*• Explanation* | |

| | | |
|---|---|---|
| *need for responsible use and continuous innovation in LLM-powered applications.* | *• Conversation*<br>*• Didactical demonstration* | |
| *13. Presentation of the student projects.* | *• Interactive exposure*<br>*• Explanation*<br>*• Conversation*<br>*• Dialogue, debate* | |
| *14. Presentation of the student projects.* | *• Interactive exposure*<br>*• Explanation*<br>*• Conversation*<br>*• Dialogue, debate* | |


*Bibliography*

*1. Chung, H. W., Hou, L., Longpre, S., Zoph, B., Tay, Y., Fedus, W., Li, Y., Wang, X., Dehghani, M., Brahma, S., Webson, A., Gu, S. S., Dai, Z., Suzgun, M., Chen, X., Chowdhery, A., Pellat, M., Robinson, K., Valter, D., . . . Wei, J. (2022). Scaling Instruction-Finetuned Language Models.*

*2. Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., & Cao, Y. (2022). ReAct: Synergizing Reasoning and Acting in Language Models.*

*3. Wu, S., Irsoy, O., Lu, S., Dabravolski, V., Dredze, M., Gehrmann, S., Kambadur, P., Rosenberg, D., & Mann, G. (2023). BloombergGPT: A Large Language Model for Finance.*

*4. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention Is All You Need*

*5. Alammar, J,m Grootendorst, M. (2024). Hands-On Large Language Models.*

*6. Auffarth, B. (2023). Generative AI*


| 8.2 Seminar / laboratory | Teaching methods | Remarks |
|---|---|---|
| *1. Introduction to LLMs and Text Generation. Get hands-on experience with basic LLM operations, focusing on generating text using pre-trained models.* | *• Interactive exposure*<br>*• Explanation*<br>*• Conversation*<br>*• Individual and group work*<br>*• Dialogue, debate* | |
| *2. Exploring Transformer Architectures. Dive into transformer models, understanding attention mechanisms and their implementation in text generation tasks.* | *• Interactive exposure*<br>*• Explanation*<br>*• Conversation*<br>*• Individual and group work*<br>*• Dialogue, debate* | |
| *3. Advanced Text Generation and Prompt Engineering. Experiment with advanced text generation techniques and learn the art of prompt engineering to guide LLM outputs.* | *• Interactive exposure*<br>*• Explanation*<br>*• Conversation*<br>*• Individual and group work*<br>*• Dialogue, debate* | |
| *4. Pre-Training and Fine-Tuning Strategies. Hands-on session on the basics of pre-training LLMs and strategies for fine-tuning them on specific tasks.* | *• Interactive exposure*<br>*• Explanation*<br>*• Conversation*<br>*• Individual and group work*<br>*• Dialogue, debate* | |
| *5. Reinforcement Learning from Human Feedback (RLHF). Implement RLHF techniques, setting up feedback loops and* | *• Interactive exposure*<br>*• Explanation*<br>*• Conversation*<br>*• Individual and* | |

| | | |
|---|---|---|
| *understanding reward models to align LLM outputs with human values.* | *group work*<br>*• Dialogue, debate* | |
| *6. Introduction to LangChain and Retrieval Augmented Generation (RAG). Begin working with LangChain, focusing on document loading, vector stores, and embeddings. Explore the implementation of RAG for enhancing LLM applications.* | *• Interactive exposure*<br>*• Explanation*<br>*• Conversation*<br>*• Individual and group work*<br>*• Dialogue, debate* | |
| *7. Building a Chatbot. Students will apply the knowledge gained in LangChain and RAG to to build a functional chatbot.*<br><br>***PROJECT***<br>***Phase 1 (Weeks 1 and 2): Introduction and Topic Selection***<br><br>*Presentation of a list of project topics that incorporate LLMs, focusing on the requirements from the standpoint of real-world clients. Students choose or propose their own project topics, working in groups. Discussion about the chosen projects to ensure feasibility and relevance by using the methodology of Generative AI project lifecycle. Initial state-of-the-art analysis, focusing on how similar challenges are approached using LLMs.*<br><br>***Phase 2 (Weeks 3 and 4): Preparation and Planning***<br>*Following their selected topic, each team is tasked with identifying and defining a list of NLP applications, then conducting a literature review to determine the highest performing pretrained models for those specified use cases.*<br><br>***Phase 3 (Weeks 5 and 6): Adapt and Align model I***<br>*.*<br>*Apply prompt engineering techniques to refine the model's output without undergoing retraining, followed by an evaluation of the model's performance.* | *• Interactive exposure*<br>*• Explanation*<br>*• Conversation*<br>*• Individual and group work*<br>*• Dialogue, debate*<br>*• Interactive exposure*<br>*• Explanation*<br>*• Conversation*<br>*• Individual and group work* | |

| | | |
|---|---|---|
| ***Phase 4 (Weeks 7 and 8): Adapt and Align model II.***<br><br>*Implement fine-tuning methods to retrain the models, enhancing their performance for the particularities of the selected topic, then proceed to evaluate the model.*<br>***Phase 5 (Weeks 9 and 10): Adapt and Align model III.***<br><br>*Incorporate Reinforcement Learning from Human Feedback (RLHF) and reward models to tailor the LLM output more closely with human values.*<br><br>***Phase 6 (Weeks 11 and 12): LangChain and Retrieval Augmented Generation (RAG)***<br><br>*Utilizing LangChain and RAG, students are required to integrate the LLM they developed into an actual application workflow. This integration should ensure the LLM's output is in harmony with topic-specific requirements, accomplished through the employment of document loading, vector stores, and embeddings.*<br><br>***Phase 7 (Weeks 13 and 14):***<br>*Oral presentations* | | |

***Bibliography***

*1. Chung, H. W., Hou, L., Longpre, S., Zoph, B., Tay, Y., Fedus, W., Li, Y., Wang, X., Dehghani, M., Brahma, S., Webson, A., Gu, S. S., Dai, Z., Suzgun, M., Chen, X., Chowdhery, A., Pellat, M., Robinson, K., Valter, D., . . . Wei, J. (2022). Scaling Instruction-Finetuned Language Models.*
*2. Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., & Cao, Y. (2022). ReAct: Synergizing Reasoning and Acting in Language Models.*
*3. Wu, S., Irsoy, O., Lu, S., Dabravolski, V., Dredze, M., Gehrmann, S., Kambadur, P., Rosenberg, D., & Mann, G. (2023). BloombergGPT: A Large Language Model for Finance.*
*4. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I.*
*(2017). Attention Is All You Need*
*5. Alammar, J,m Grootendorst, M. (2024). Hands-On Large Language Models.*
*6. Auffarth, B. (2023). Generative AI with LangChain: Build large language model (LLM) apps with Python, ChatGPT and other LLMs*

**9. Evaluation**

| Activity type | 9.1 Evaluation criteria | 9.2 Evaluation methods | 9.3 Percentage of final grade |
|---|---|---|---|
| 9.4 Course | • *The capability to utilize the knowledge* | *Oral examination (project)* | 60% |

| | | | |
|---|---|---|---|
| | *acquired from the course and practiced in the labs to address practical problems and real-world requirements with applications in natural language processing and generative AI.* | | |
| | *• The student possesses a thorough comprehension of Large Language Model (LLM) concepts, including transformer architectures, prompt engineering, and LangChain applications.* | *Practical Examination under continuous observation (solving lab tasks)* | 40% |
| 9.5 Seminar/laboratory | | | |
| | | | |
| 9.6 Minimum standard of performance | | | |
| • *Students must prove that they acquired an acceptable level of knowledge and understanding of the core concepts taught in the class, that they are capable of using this knowledge in a coherent form, that they have the ability to establish certain connections and to use the knowledge in solving various computer vision problems.*<br>• *The final grade (weighted average between the two presented evaluation methods) should be at least 5 (no rounding, from a scale from 1 to 10).* | | | |

**10. SDG labels (Sustainable Development Goals)[3]**

| | | |
|---|---|---|
| | ○ | Sustainable Development Generic Label |

| 1 FĂRĂ SĂRĂCIE | 2 FOAMETE „ZERO" | 3 SĂNĂTATE ȘI BUNĂSTARE | 4 EDUCAȚIE DE CALITATE | 5 EGALITATE DE GEN | 6 APĂ CURATĂ ȘI SANITAȚIE | 7 ENERGIE CURATĂ ȘI LA PREȚURI ACCESIBILE | 8 MUNCĂ DECENTĂ ȘI CREȘTERE ECONOMICĂ | 9 INDUSTRIE, INOVAȚIE ȘI INFRASTRUCTURĂ |
|---|---|---|---|---|---|---|---|---|
| ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | **X** |

[3] Select a single label which, according to the *Implementation of SDG labels in the academic process*, best matches the subject. If the subject addresses sustainable development in a generic manner (i.e. by presenting/introducing the general framework of sustainable development, etc.), then the Sustainable Development generic label may be applied. If none of the labels describe the subject, select the last option: *"No label applies."*

| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
|  |  |  |  |  |  |  |  | No label applies |
| ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |

Date of entry:
15.05.2026

Signature of course coordinator

**Lect. Dr. Bogdan MURSA**

Signature of seminar coordinator

**Lect. Dr. Bogdan MURSA**

Date of approval in the department:
...

Signature of the head of department

Assoc.prof.phd. Adrian STERCA