

SYLLABUS

Understanding and Developing Large Language Models (LLMs)

University year **2025-2026**

1. Information regarding the programme

1.1. Higher education institution	<i>Babeş-Bolyai University</i>
1.2. Faculty	<i>Faculty of Mathematics and Computer Science</i>
1.3. Department	<i>Department of Computer Science</i>
1.4. Field of study	<i>Computer Science</i>
1.5. Study cycle	<i>Bachelor</i>
1.6. Study programme/Qualification	<i>Computer Science</i>
1.7. Form of education	<i>Full time</i>

2. Information regarding the discipline

2.1. Name of the discipline			<i>Understanding and Developing Large Language Models (LLMs)</i>					Discipline code		<i>MLE5247</i>	
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		6	2.6. Type of evaluation		E	2.7. Discipline regime		<i>Optional</i>

3. Total estimated time (hours/semester of didactic activities)

3.1. Hours per week	5	of which: 3.2 course	2	3.3 seminar/laboratory/project	1 lab + 2 proj
3.4. Total hours in the curriculum	60	of which: 3.5 course	24	3.6 seminar/laboratory/project	36
Time allotment for individual study (ID) and self-study activities (SA)					hours
Learning using manual, course support, bibliography, course notes (SA)					12
Additional documentation (in libraries, on electronic platforms, field documentation)					16
Preparation for seminars/labs, homework, papers, portfolios and essays					25
Tutorship					6
Evaluations					6
Other activities:					
3.7. Total individual study hours	65				
3.8. Total hours per semester	125				
3.9. Number of ECTS credits	5				

4. Prerequisites (if necessary)

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

5. Conditions (if necessary)

5.1. for the course	• <i>Classroom with a video project device</i>
---------------------	--

5.2. for the seminar /lab activities	• <i>Lab equipped with high-performance computers and Python installed.</i>
--------------------------------------	---

6.1. Specific competencies acquired ¹

Professional/essential competencies	• <i>CE1.3 Using the methods, techniques, and algorithms from AI in order to model several classes of problems</i> • <i>CE1.4 Identify and explain specific AI techniques and algorithms and using them to solve specific problems</i> • <i>CE1.5 Integrating the models and the specific solutions from AI in dedicated applications</i> • <i>C4.2 Interpretation of mathematical models and computer science (formal)</i> • <i>C4.3 Identifying appropriate models and methods to solve real problems</i> • <i>C4.5 Incorporation of formal models in specific applications in various fields</i>
Transversal competencies	• <i>CT1 Ability to conform to the requirements of organized and efficient work, to develop a responsible approach towards the academic and scientific fields, in order to make the most of one's own creative potential, while obeying the rules and principles of professional ethic.</i> • <i>CT3 Using efficient methods and techniques for learning, information, research and developing capabilities for using knowledge, for adapting to a dynamic society and for communicating in Romanian and in a worldwide spoken language.</i>

6.2. Learning outcomes

Knowledge	• 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. • The graduate knows and understands the concepts and the techniques of knowledge representation and is able to apply them for problem solving.
Skills	• The graduate is able to identify complex issues and examine related issues in order to design several solutions and implement these solutions. • 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.
Responsibility and autonomy:	• The graduate is able to apply architectural templates, design templates and best practices in the field to design highly complex software applications. • The graduate has the necessary skills to apply various methods and tools for analysis and visualizing the results of the used Artificial Intelligence algorithms and techniques. • The graduate has the necessary knowledge to review the literature and use international databases and international digital research libraries.

7. Objectives of the discipline (outcome of the acquired competencies)

¹ One can choose either competences or learning outcomes, or both. If only one option is chosen, the row related to the other option will be deleted, and the kept one will be numbered 6.

7.1 General objective of the discipline	<i>The goal of this course is to familiarize students with the field of natural language processing, focusing particularly on the latest advancements brought by transformer architecture. Students will be taught how to analyze, design, implement, and evaluate various NLP problems. This course aims to elucidate how NLP serves as a bridge between human language and machine understanding, enabling tasks like text classification, entity extraction, text summarization, text generation, chatbots, among others. Specifically, all these will be accomplished by leveraging the latest technical breakthroughs in Large Language Models (LLMs)</i>
7.2 Specific objective of the discipline	• <i>Understand various architectures of Large Language Models (LLMs) with a focus on transformer architectures for tasks such as text classification, entity extraction, text summarization, text generation, and many others.</i> • <i>Solve and analyze a natural language processing problem using specific theoretical frameworks and methodologies inherent to LLMs.</i> • <i>Understand and develop effective strategies for prompt engineering, including techniques for eliciting desired responses from LLMs through well-crafted prompts.</i> • <i>Learn techniques for fine-tuning and retraining Large Language Models to enhance performance and adaptability to specific NLP tasks.</i> • <i>Understand the metrics used to evaluate the performance of LLMs and the principles behind deploying these models in real-world applications, including bot creation.</i>

8. Content

8.1 Course	Teaching methods	Remarks
<i>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.</i>	• <i>Interactive exposure</i> • <i>Explanation</i> • <i>Conversation</i> • <i>Didactical demonstration</i>	
<i>2. The Evolution of Text Generation Technologies. Tracing the development of text generation from pre-transformer models to current methodologies.</i>	• <i>Interactive exposure</i> • <i>Explanation</i> • <i>Conversation</i> • <i>Didactical demonstration</i>	
<i>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.</i>	• <i>Interactive exposure</i> • <i>Explanation</i> • <i>Conversation</i> • <i>Didactical demonstration</i>	
<i>4. The Principle of Attention in</i>	• <i>Interactive exposure</i>	

<i>Transformers. Understanding the "Attention is all you need" concept and its revolutionary impact on LLMs.</i>	• <i>Explanation</i> • <i>Conversation</i> • <i>Didactical demonstration</i>	
<i>5. Mastering Prompt Engineering. Learning how to effectively design prompts to guide LLMs in generating desired outputs.</i>	• <i>Interactive exposure</i> • <i>Explanation</i> • <i>Conversation</i> • <i>Didactical demonstration</i>	
<i>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.</i>	• <i>Interactive exposure</i> • <i>Explanation</i> • <i>Conversation</i> • <i>Didactical demonstration</i>	
<i>7. Fine-Tuning LLMs for Specific Tasks. Strategies for instruction-based fine-tuning, including single and multi-task adaptations.</i>	• <i>Interactive exposure</i> • <i>Explanation</i> • <i>Conversation</i> • <i>Didactical demonstration</i>	
<i>8. Advanced Fine-Tuning Techniques. Introduction to Parameter Efficient Fine-Tuning (PEFT) methods such as LoRA and Soft Prompts.</i>	• <i>Interactive exposure</i> • <i>Explanation</i> • <i>Conversation</i> • <i>Didactical demonstration</i>	
<i>9. Reinforcement Learning from Human Feedback (RLHF). Fundamentals of aligning LLMs with human values through RLHF, including feedback collection and reward models.</i>	• <i>Interactive exposure</i> • <i>Explanation</i> • <i>Conversation</i> • <i>Didactical demonstration</i>	
<i>10. Enhancing LLM output using Reasoning 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.</i>	• <i>Interactive exposure</i> • <i>Explanation</i> • <i>Conversation</i> • <i>Didactical demonstration</i>	
<i>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).</i>	• <i>Interactive exposure</i> • <i>Explanation</i> • <i>Conversation</i> • <i>Didactical demonstration</i>	
<i>12. Ethics of AI. Discover the evolving field of generative AI, emphasizing the need for responsible use and continuous</i>	• <i>Interactive exposure</i> • <i>Explanation</i> • <i>Conversation</i>	

<i>innovation in LLM-powered applications.</i>	• <i>Didactical demonstration</i>	
<i>13. Presentation of the student projects.</i>	• <i>Interactive exposure</i> • <i>Explanation</i> • <i>Conversation</i> • <i>Dialogue, debate</i>	
<i>14. Presentation of the student projects.</i>	• <i>Interactive exposure</i> • <i>Explanation</i> • <i>Conversation</i> • <i>Dialogue, debate</i>	
<i>Bibliography</i> 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). <i>Scaling Instruction-Finetuned Language Models</i> . 2. Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., & Cao, Y. (2022). <i>ReAct: Synergizing Reasoning and Acting in Language Models</i> . 3. Wu, S., Irsoy, O., Lu, S., Dabrovolski, V., Dredze, M., Gehrmann, S., Kambadur, P., Rosenberg, D., & Mann, G. (2023). <i>BloombergGPT: A Large Language Model for Finance</i> . 4. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). <i>Attention Is All You Need</i> 5. Alammam, J,m Grootendorst, M. (2024). <i>Hands-On Large Language Models</i> . 6. Auffarth, B. (2023). <i>Generative AI</i>		
8.2 Seminar / laboratory	Teaching methods	Remarks
<i>1. Introduction to LLMs and Text Generation. Get hands-on experience with basic LLM operations, focusing on generating text using pre-trained models.</i>	• <i>Interactive exposure</i> • <i>Explanation</i> • <i>Conversation</i> • <i>Individual and group work</i> • <i>Dialogue, debate</i>	
<i>2. Exploring Transformer Architectures. Dive into transformer models, understanding attention mechanisms and their implementation in text generation tasks.</i>	• <i>Interactive exposure</i> • <i>Explanation</i> • <i>Conversation</i> • <i>Individual and group work</i> • <i>Dialogue, debate</i>	
<i>3. Advanced Text Generation and Prompt Engineering. Experiment with advanced text generation techniques and learn the art of prompt engineering to guide LLM outputs.</i>	• <i>Interactive exposure</i> • <i>Explanation</i> • <i>Conversation</i> • <i>Individual and group work</i> • <i>Dialogue, debate</i>	
<i>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.</i>	• <i>Interactive exposure</i> • <i>Explanation</i> • <i>Conversation</i> • <i>Individual and group work</i> • <i>Dialogue, debate</i>	
<i>5. Reinforcement Learning from Human Feedback (RLHF). Implement RLHF techniques, setting up feedback loops and understanding reward models to align LLM</i>	• <i>Interactive exposure</i> • <i>Explanation</i> • <i>Conversation</i> • <i>Individual and group work</i>	

outputs with human values.	• 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 • Explanation • Conversation • Individual and group work • Dialogue, debate	
7. Building a Chatbot. Students will apply the knowledge gained in LangChain and RAG to build a functional chatbot. PROJECT Phase 1 (Weeks 1 and 2): Introduction and Topic Selection 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. Phase 2 (Weeks 3 and 4): Preparation and Planning 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. Phase 3 (Weeks 5 and 6): Adapt and Align model I . Apply prompt engineering techniques to refine the model's output without undergoing retraining, followed by an evaluation of the model's performance. Phase 4 (Weeks 7 and 8): Adapt and Align model II.	• Interactive exposure • Explanation • Conversation • Individual and group work • Dialogue, debate • Interactive exposure • Explanation • Conversation • Individual and group work	

<i>Implement fine-tuning methods to retrain the models, enhancing their performance for the particularities of the selected topic, then proceed to evaluate the model.</i> Phase 5 (Weeks 9 and 10): Adapt and Align model III. <i>Incorporate Reinforcement Learning from Human Feedback (RLHF) and reward models to tailor the LLM output more closely with human values.</i> Phase 6 (Weeks 11 and 12): LangChain and Retrieval Augmented Generation (RAG) <i>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.</i> Phase 7 (Weeks 13 and 14): <i>Oral presentations</i>		
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). <i>Scaling Instruction-Finetuned Language Models</i>. 2. Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., & Cao, Y. (2022). <i>ReAct: Synergizing Reasoning and Acting in Language Models</i>. 3. Wu, S., Irsoy, O., Lu, S., Dabrovolski, V., Dredze, M., Gehrmann, S., Kambadur, P., Rosenberg, D., & Mann, G. (2023). <i>BloombergGPT: A Large Language Model for Finance</i>. 4. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). <i>Attention Is All You Need</i> 5. Alammari, J, Grootendorst, M. (2024). <i>Hands-On Large Language Models</i>. 6. Auffarth, B. (2023). <i>Generative AI with LangChain: Build large language model (LLM) apps with Python, ChatGPT and other LLMs</i>		

9. Corroborating the content of the discipline with the expectations of the epistemic community, professional associations and representative employers within the field of the program

- *Similar courses exist in the studying program of major universities in Europe and abroad.*
- *The course respects the IEEE and ACM Curricula Recommendations for Computer Science studies.*
- *The knowledge and skills gained from this course not only provide students with a foundation for embarking on a career in scientific research but also position them as sought-after LLM engineers in the industry, where there is a high demand for experts.*

10. Evaluation

Activity type	10.1 Evaluation criteria	10.2 Evaluation methods	10.3 Percentage of final grade
10.4 Course	• <i>The capability to utilize the knowledge 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.</i>	<i>Oral examination (project)</i>	60%
	• <i>The student possesses a thorough comprehension of Large Language Model (LLM) concepts, including transformer architectures, prompt engineering, and LangChain applications.</i>	<i>Practical Examination under continuous observation (solving lab tasks)</i>	40%
10.5 Seminar/laboratory			
10.6 Minimum standard of performance			
• <i>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.</i> • <i>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).</i>			

11. Labels ODD (Sustainable Development Goals)²

Not applicable.

Date:
15.04.2025

Signature of course coordinator

Lect. Dr. Bogdan MURSA

Signature of seminar coordinator

Lect. Dr. Bogdan MURSA

Date of approval:

Signature of the head of department

Assoc.prof.phd. Adrian STERCA

² Keep only the labels that, according to the [*Procedure for applying ODD labels in the academic process*](#), suit the discipline and delete the others, including the general one for *Sustainable Development* – if not applicable. If no label describes the discipline, delete them all and write „*Not applicable.*”.