

SYLLABUS

1. Information regarding the programme

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 of study	Computer Science
1.5 Study cycle	Master
1.6 Study programme / Qualification	Artificial Intelligence for Connected Industries

2. Information regarding the discipline

2.1 Name of the discipline (en) (ro)	Advanced Machine Learning						
2.2 Course coordinator	Prof. PhD. Dioşan Laura						
2.3 Seminar coordinator	Prof. PhD. Dioşan Laura						
2.4. Year of study	1	2.5 Semester	3	2.6. Type of evaluation	E	2.7 Type of discipline	Compulsory
2.8 Code of the discipline	MME8219						

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

3.1 Hours per week	4	Of which: 3.2 course	2	3.3 seminar/laboratory	1sem+1 project
3.4 Total hours in the curriculum	56	Of which: 3.5 course	28	3.6 seminar/laboratory	28
Time allotment:					hours
Learning using manual, course support, bibliography, course notes					36
Additional documentation (in libraries, on electronic platforms, field documentation)					40
Preparation for seminars/labs, homework, papers, portfolios and essays					35
Tutorship					4
Evaluations					4
Other activities:					-
3.7 Total individual study hours	119				
3.8 Total hours per semester	175				
3.9 Number of ECTS credits	7				

4. Prerequisites (if necessary)

4.1. curriculum	<ul style="list-style-type: none"> Algorithms, data structures, statistics, Basic elements of Artificial Intelligence and Machine Learning
4.2. competencies	<ul style="list-style-type: none"> Average programming skills

5. Conditions (if necessary)

5.1. for the course	<ul style="list-style-type: none"> • Projector
5.2. for the seminar /lab activities	<ul style="list-style-type: none"> • Computers, specific development environment

6. Specific competencies acquired

Professional competencies	<p>CE1.3 Use of Artificial Intelligence's methods, techniques and algorithms for modelling problem solutions</p> <p>CE1.4. Identify and explain Artificial Intelligence's techniques and algorithms and solving specific problems</p> <p>CE1.5. Integration of Artificial Intelligence models and solutions in specific applications</p>
Transversal competencies	<p>CT1. Application of efficient work rules and responsible attitudes towards the scientific domain, for the creative exploitation of one's own potential according to the principles and rules of professional ethics</p> <p>CT2. Efficient conduct of activities organized in an interdisciplinary group and development of empathic capacity of interpersonal communication, networking and collaboration with diverse groups</p> <p>CT3. Use of efficient methods and techniques for learning, information, research and development of abilities for knowledge exploitation, for adapting to the needs of a dynamic society and for communication in a widely used foreign language.</p>

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

7.1 General objective of the discipline	Emphasis the proper Machine Learning methods and techniques (optimisation algorithms, machine learning) for solving various problems for industry and society (healthcare, biology, psychology, finance, etc.)
7.2 Specific objective of the discipline	<p>This course is aimed to advance both theoretical and practical aspects of Machine Learning (supervised, unsupervised, semi-supervised, reinforcement learning). To the end, the students will be able to:</p> <ul style="list-style-type: none"> • Identify the society's challenges that can be solved by machine learning methods and to propose AI-based solutions • Describe the ML methods (basic concepts, design and implementation) • Model the social challenges as mathematical problems that can be solved by intelligent algorithms and to adapt them to particular problems

- Describe the evaluation criteria of ML methods
- Prepare presentations of the implemented projects

8. Content

8.1 Course	Teaching methods	Remarks
<ol style="list-style-type: none"> 1. Foundations of Machine & Deep Learning <ul style="list-style-type: none"> ○ Revisiting classical ML/DL from a modern perspective ○ Data-centric vs model-centric AI ○ Evaluating AI systems (accuracy vs. calibration vs. utility) 2. Generative AI approaches <ul style="list-style-type: none"> ○ Transformer-based models for processing textual, visual, audio or code data(LLMs, VLMs, AudioLMs, Code models) <ul style="list-style-type: none"> ▪ Tokenisation ▪ Embeddings ▪ Attention components ○ Diffusion models beyond images ○ Transfer learning, instruction-tuning and fine-tuning ○ Evaluation, Robustness & Cross-Domain Generalization 3. Reinforcement Learning for Decision-Making & Agents <ul style="list-style-type: none"> ○ Foundations of RL ○ RL for Language Models ○ RL for structured reasoning & tool use ○ Preference Optimization for LLMs (DPO, ORPO, IPO, RLAIF) 4. Foundations of Retrieval-Augmented Generation (RAG) and Knowledge Systems <ul style="list-style-type: none"> ○ RAG architectures ○ Chunking strategies & semantic compression ○ Modern vector databases (Chroma, Weaviate, Milvus, Qdrant) ○ Evaluating RAG 5. Advances of RAG systems <ul style="list-style-type: none"> ○ RAG architectures <ul style="list-style-type: none"> ▪ RAG vs RAG2.0 ▪ Indexed RAG vs. Self-RAG ▪ GraphRAG ○ Chunking Strategies & Semantic Compression <ul style="list-style-type: none"> ▪ Adaptive vs context-aware vs multi-modal ▪ Hierarchical vs. LLM-based 	<ul style="list-style-type: none"> • Interactive exposure • Presentation • Explanation • Practical examples • Case-study discussions 	

<ul style="list-style-type: none"> vs. retriever-based compression <ul style="list-style-type: none"> ▪ Token-efficient indexing ▪ Information density optimisation ○ Enhanced RAG: reranking, tool-use, memory ○ Evaluating RAG with IR metrics + LLM judges ○ RL for RAG optimization 6. Agentic AI approaches <ul style="list-style-type: none"> ○ Workflow & components <ul style="list-style-type: none"> ▪ Planner, reasoner, memory, tool-use controller, critic & evaluator, router, execution engine ▪ State machines for agent transitions ▪ Execution traces & logging ○ Multi-Agent Reinforcement Learning ○ Hierarchical agent architectures (ReAct → Reflexion → AutoGen → CrewAI) <ul style="list-style-type: none"> ▪ Meta agents ▪ Skill decomposition ▪ Adaptive agent routing ○ Planning algorithms <ul style="list-style-type: none"> ▪ Algorithmic foundations ▪ Formal methods ○ Collaboration protocols among agents <ul style="list-style-type: none"> ▪ Coordination protocols ▪ Agent communication ○ Tool-ing and API orchestration ○ Memory & long-term learning for agents ○ Agent evaluation & debugging 7. Physics AI approaches <ul style="list-style-type: none"> ○ PINNs (physics-informed neural networks) ○ Digital Twin models and pipelines (data engineering + simulation + LLM integration) 		
--	--	--

Bibliography

1. A. Hopgood, Intelligent Systems for Engineers and Scientists, CRC Press, 2001
2. T. M. Mitchell, Machine Learning, McGraw-Hill Science, 1997
3. D. J. C. MacKey, Information Theory, Inference and Learning Algorithms, Cambridge University Press, 2003
4. C. Bishop, Pattern Recognition and Machine Learning, Springer, 2006
5. P. F. Brown, S. Della Pietra, V. J. Della Pietra, and R. L. Mercer. The mathematic of statistical

machine translation: Parameter estimation. *Computational Linguistics*, 19(2):263-311, 1994

6. Zhao, P., Zhang, H., Yu, Q., Wang, Z., Geng, Y., Fu, F., ... & Cui, B. (2026). Retrieval-augmented generation for ai-generated content: A survey. *Data Science and Engineering*, 1-29.
7. Tural, B., Örpek, Z., & Destan, Z. (2024, December). Retrieval-augmented generation (RAG) and LLM integration. In *2024 8th International Symposium on Innovative Approaches in Smart Technologies (ISAS)* (pp. 1-5). IEEE.
8. Singh, A., Ehtesham, A., Kumar, S., & Khoei, T. T. (2025). Agentic retrieval-augmented generation: A survey on agentic rag. *arXiv preprint arXiv:2501.09136*.

8.2 Seminar / laboratory	Teaching methods	Remarks
<p>Project development</p> <p>Students pick one domain (health, legal, finance, education, robotics, smart city, etc.).</p> <p>They will build and document:</p> <ol style="list-style-type: none"> 1. A baseline ML model 2. A generative model enhancement 3. A RAG system on the same domain 4. An autonomous agent powered by these models 5. (Optional) connect the agent to a digital twin or physics model <p>Steps:</p> <ol style="list-style-type: none"> 1. Choose a dataset and build a classical ML pipeline (EDA, training a baseline ML model, evaluation) 2. Extend the Baseline with a Fine-Tuned LLM (extend the dataset by augmentation methods or be generating new synthetic data and fine-tune or adapt an open LLM and compare it to the baseline model) 3. Build a Domain RAG Pipeline by turning the system into a knowledge-grounded assistant (wraps the LLM in a new architecture that injects external knowledge and ensures grounding) 4. Transform the RAG system into a Multi-Agent Pipeline (by introducing reasoning, tool-use, autonomous planning, etc.) 5. [optional] Connect the Agent to a Digital Twin or Physics Simulation <p>In all the stages, the classes will be organised in an interactive way in order to construct in a collaborative manner feasible solutions. Feedback will be provided along the entire project development from both academic and industry specialists.</p>	<ul style="list-style-type: none"> • Interactive exposure • Explanation • Conversation • Didactical demonstration 	
<p>Bibliography</p> <ol style="list-style-type: none"> 1. A. Hopgood, <i>Intelligent Systems for Engineers and Scientists</i>, CRC Press, 2001 		

2. T. M. Mitchell, *Machine Learning*, McGraw-Hill Science, 1997
3. D. J. C. MacKey, *Information Theory, Inference and Learning Algorithms*, Cambridge University Press, 2003
4. C. Bishop, *Pattern Recognition and Machine Learning*, Springer, 2006
5. P. F. Brown, S. Della Pietra, V. J. Della Pietra, and R. L. Mercer. The mathematic of statistical machine translation: Parameter estimation. *Computational Linguistics*, 19(2):263-311, 1994
6. Zhao, P., Zhang, H., Yu, Q., Wang, Z., Geng, Y., Fu, F., ... & Cui, B. (2026). Retrieval-augmented generation for ai-generated content: A survey. *Data Science and Engineering*, 1-29.
7. Tural, B., Örpek, Z., & Destan, Z. (2024, December). Retrieval-augmented generation (RAG) and LLM integration. In *2024 8th International Symposium on Innovative Approaches in Smart Technologies (ISAS)* (pp. 1-5). IEEE.
8. Singh, A., Ehtesham, A., Kumar, S., & Khoei, T. T. (2025). Agentic retrieval-augmented generation: A survey on agentic rag. *arXiv preprint arXiv:2501.09136*.

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

- The course exists in the curriculum of many universities in the world.
- The results of course are considered by software companies particularly useful and topical, developing needed abilities in modelling and visualization of data.

10. Evaluation

Type of activity	10.1 Evaluation criteria	10.2 Evaluation methods	10.3 Share in the grade (%)
10.4 Course	Know concepts and methods from the domain of artificial intelligence	Research report and project presentation	50%
10.5 Seminar/lab activities	Apply AI techniques in real problems	Project development	50%
10.6 Minimum performance standards			
Type of evaluation – project-based			
Each student should implement 70% of the project.			

Date

Signature of course coordinator

Signature of seminar coordinator

Prof. PhD. Dioşan Laura

Prof. PhD. Dioşan Laura

Date of approval

Signature of the head of department

.....