UNIVERSITY OF DEBRECEN, FACULTY OF INFORMATICS

H-4028 Debrecen, Kassai Road 26., H-4002 Debrecen, P.O. box 400.
☎ 52/518-630, to@inf.unideb.hu

ARTIFICIAL INTELLIGENCE MSC

2026

Mode:	Full-time training
Program Coordinator:	Dr. Balázs Harangi (harangi.balazs@inf.unideb.hu)
Mentor:	Dr. Gergő Bogacsovics (bogacsovics.gergo@inf.unideb.hu)

Qualification requirements

General requirements of the diploma are regulated by The Rules and Regulations of The University of Debrecen.

Work and Fire Safety and Physical Education

The courses of „Work and Fire Safety" and „Physical Education" are worth 1 - 1 credit, which must be completed in excess of the number of credits required for the diploma as specified in the training and outcome requirements of the degree.

Diploma credit requirements:

Mathematics and Natural Science:	18 credits
Informatics and Artificial Intelligence compulsory topics:	24 credits
Differentiated knowledge topics:	42 credits
Thesis work:	30 credits
Free choice:	6 credits
Total	**120 credits**
Work and Fire Safety Training:	1 credit
Physical Education (1 semester):	1 credit

Professional training/Internship requirements

Professional training is a practice which is completed at a competent training place. It lasts for at least 6 weeks and 240 work hours.

It is a must to complete Professional training subject to issue the absolutorium (pre-degree certificate).

https://inf.unideb.hu/en/professional-training

Student can apply for Professional training after completing at least one semester.

Faculty of Informatics annex to the Academic and Examination Rules and Regulations of the University of Debrecen contains the procedure of the professional training.

The Thesis

During the studies, Student must write a thesis. Writing a thesis is a diploma requirement.

Thesis subject is mandatory to complete. The prerequisites to register for the Thesis subject are the followings:

- chose a thesis topic by the deadline.

  (Together with the supervisor the candidate writes a work plan in the maximum of two pages. The work plan describes the aim of the work, areas of expertise and the scheduling of the work.)
- the chosen topic is approved by the Educational Committee
- at least 30 completed credits.

Final Exam / State Exam

a., Requirements for Final Exam

1. Complete all the 120 credits required by the curriculum of program specialisation to have the degree of MSc program
2. Carry out the internship
3. Write and submit the Diploma Thesis

b., Process of the Final Exam

The Final Exam consists of an oral part only and the purpose is to examine the coherence of the professional knowledge.

F. The average from the grades of the oral exam (rounded to a whole number). If the grade for any item is failed, the grade is failed, and the final examination is failed.

D1. Thesis defence. During the defence the candidate has to sum up the Thesis in a short presentation then s/he answers the questions from the referee of the Thesis and the members of the Committee.

D2. The grade for the thesis, which is determined by the Final Examination Committee taking into account the grade proposed by the thesis assessor.

Calculation of the final examination grade (**ZV**): ZV = (F+D1+D2)/3

If the grade D2 is failed, the candidate will not be allowed to sit the final examination.

If any of the grades of F or D1 are unsatisfactory, the final exam is also unsatisfactory. Only the component graded as unsatisfactory must be retaken in the retake of the final examination.

Grade of Diploma:

Diploma grade: in the case of a successful final examination, it is determined based on the average of the following results:

a) **SZ**: Average of the grades for the Thesis subject, the grade for the thesis assessment and the grades for the thesis defence in the final examination, rounded to two decimal places.

b) **F**: Average of the grades obtained in the final examination, rounded to a whole number.

c) **T**: the credit-weighted average of all compulsory and optional professional subjects completed during the course, except for 'Thesis 1' and 'Thesis 2', rounded to two decimal places.

Diploma grade = (0,3*SZ+0,2*F+0,5*T)

Based on the above average result, the qualification of the diploma is determined by the University of Debrecen's Academic and Examination Regulations, Section 31 (7).

The diploma shall be assessed based on the calculation of the grade average as follows:

outstanding	4,81-5,00
excellent	4,51-4,80
good	3,51-4,50
satisfactory	2,51-3,50
pass	2,00-2,50

Mathematics and Natural Science – needed 18 credits

Code	Subject name	Cre-dit	Type and number			Asses-ment	Prerequisites	Period	Semes-ter
			lec.	practice					
				sem.	lab				
INMEA0101-26 INMEA0101E	Mathematical Foundations of Artificial Intelligence	3	2			E			1
INMEA0102-26 INMEA0102E INMEA0102L	Statistical Analysis of Artificial Intelligence Methods	6	2		2	E S			1
INMEA0103-26 INMEA0103L	Advanced Optimization Methods	3			2	PM			1
INMEA0104-26 INMEA0104E	Theory of Neural Networks	3	2			E			1
INMEA0207-26 INMEA0207L	Operations Research	3			2	PM			2

Informatics and Artificial Intelligence compulsory topics – needed 24 credits

Code	Subject name	Cre-dit	Type and number			Asses-ment	Prerequisites	Period	Semes-ter
			lec.	practice					
				sem.	lab				
INMEA0105-26 INMEA0105E INMEA0105L	Machine Learning	6	2		2	E S			1
INMEA0106-26 INMEA0106L	Advanced Machine Learning Models	3			2	PM			1
INMEA0208-26 INMEA0208E INMEA0208L	Deep Learning	6	2		2	E S			2
INMEA0209-26 INMEA0209L	Natural Language Processing and Speech Processing	3			2	PM			2
INMEA0210-26 INMEA0210L	Generative Artificial Intelligence	3			2	PM			2
INMEA0211-26 INMEA0211L	AI engineering / MLOps	3			2	PM			2

Thesis work – needed 30 credits

Code	Subject name	Cre-dit	Type and number			Asses-ment	Prerequisites	Period	Semes-ter
			lec.	practice					
				sem.	lab				
INMEA0312-26 INMEA0312L	Thesis 1	15				PM			3
INMEA0413-26 INMEA0413L	Thesis 2	15				PM			4

Differentiated knowledge topics – needed 42credits

Code	Subject name	Cre-dit	Type and number			Asses-ment	Prerequisites	Period	Semes-ter
			lec.	practice					
				sem.	lab				
INMEA9914-26 INMEA9914L	Explainable AI	3			2	PM			3
INMEA9915-26 INMEA9915E INMEA9915L	AI and cybersecurity	6	2		2	E S			4
INMEA9916-26 INMEA9916E	AI Ethics and Governance	3	2			E			2
INMEA9917-26 INMEA9917L	Visualization and Visual Analytics	3			2	PM			3
INMEA9918-26 INMEA9918L	Graph-based neural networks	3			2	PM			3
INMEA9919-26 INMEA9919L	Simulations – Digital Twin	3			2	PM			3
INMEA9920-26 INMEA9920E INMEA9922L	Reinforcement learning	6	2		2	E S			2
INMEA9921-26 INMEA9921L	Distributed AI Systems	3			2	PM			3
INMEA9922-26 INMEA9922L	Multimodal AI Tools	3			2	PM			3
INMEA9923-26 INMEA9923L	Data Mining	3			2	PM			2
INMEA9924-26 INMEA9924L	Advanced Data Management	3			2	PM			3
INMEA9925-26 INMEA9925E INMEA9925L	Computer Vision	6	2		2	E S			2
INMEA9926-26 INMEA9926L	Text and Web Mining	3			2	PM			3
INMEA9927-26 INMEA9927L	Fundamentals of Robotics	3			2	PM			3
INMEA9928-26 INMEA9928L	AI in Health Sciences	3			2	PM			4
INMEA9929-26 INMEA9929L	AI in physics	3			2	PM			4
INMEA9930-26 INMEA9930L	Docker and Kubernetes for AI	3			2	PM			2
INMEA9931-26 INMEA9931L	Advanced Cloud Computing	3			2	PM			2
INMEA9932-26 INMEA9932L	Tools for parallel programming	3			2	PM			3
INMEA9933-26 INMEA9933L	Software Develop-ment in Industrial Environments	3			2	PM			3
INMEA9934-26 INMEA9934L	Theoretical and Neural Models in the Industry	3			2	PM			3
INMEA9935-26 INMEA9935L	Effective AI: Opti-mization Techniques	3			2	PM			3

* The recommended semester can be changed occasionally.

Professional Training

Code	Subject name	Cre-dit	Type and number			Asses-ment	Prerequisites	Period	Semes-ter
			lec.	practice					
				sem.	lab				
INMEA9997-26 INMEA9997G	Professional Training	9				PM			3

Free choice – needed 6 credits

Code	Subject name	Cre-dit	Type and number			Asses-ment	Prerequisites	Period	Semes-ter
			lec.	practice					
				sem.	lab				

* "Free choice" - Professional electives offered by the Faculty of Informatics and institutional electives offered by other faculties of the University of Debrecen.

Work and Fire Safety and Physical Education – needed 2 credits

must be completed in excess of the number of credits required
for the diploma as specified in the training and outcome requirements of the degree

Code	Subject name	Cre-dit	Type and number			Asses-ment	Prerequisites	Period	Semes-ter
			lec.	practice					
				sem.	lab				
	Work and Fire Safety	1				PM			1
	Physical Education	1				PM			

Exam types:
- E exam
- S signature
- PM practical mark

ARTIFICIAL INTELLIGENCE MSC

Description of Subjects

Mathematics and Natural Science

MATHEMATICAL FOUNDATIONS OF ARTIFICIAL INTELLIGENCE

INMEA0101-26

Semester:	1
Type:	Lecture
Number of Classes:	2+0+0
Credit:	3
Status:	Obligatory
Assessment:	Exam
Prerequisites:	None
Responsible:	Dr. Bernadett Aradi

Topics:
Elements of linear algebra: linear transformations, symmetric and orthogonal transformations, their eigenvalues and eigenvectors. Matrix decompositions: Cholesky decomposition, spectral decomposition, singular value decomposition (SVD).
Activa-tion functions and their properties, their derivatives. Differentiation of scalar- and vector-valued multivariable functions: partial derivatives and gradient vector, Jacobian matrix. Higher-order derivatives: Hessian matrix, multivariable Taylor series.
Elements of probability theory: distribution and density functions, multivariate distributions, covariance, and correlation matrices.
Applica-tions: Principal Component Analysis (PCA), Canonical Correlation Analysis (CCA), Multidimensional Scaling (MDS).

Compulsory/Recommended Readings:

- Deisenroth, M.P., Faisal, A.A. and Ong, C.S.: Mathematics for Machine Learning. Cambridge: Cambridge University Press, 2020. ISBN: 9781108455145
- Hijab, O.: Math for Data Science. Springer, 2025. ISBN: 978-3-319-74072-0
- Fessler, J.A. and Nadakuditi, R.R.: Linear Algebra for Data Science, Machine Learning, and Signal Processing. Cambridge: Cambridge University Press, 2024. ISBN: 9781009418140
- Lax, P.D. and Terrell, M.S.: Multivariable Calculus with Applications. Springer, 2017. ISBN: 978-3-319-74072-0

STATISTICAL ANALYSIS OF ARTIFICIAL INTELLIGENCE METHODS

INMEA0102-26

Semester:	1
Type:	Lecture / Laboratory
Number of Classes:	2+0+2
Credit:	6
Status:	Obligatory
Assessment:	Exam
Prerequisites:	None
Responsible:	Dr. Sándor Baran

Topics:
Descriptive statistics, sampling distributions, confidence intervals, fundamentals of hypothesis testing. Principles of machine learning evaluation, error estimation procedures, cross validation. Evaluation for classification: the common performance metrics (confusion matrix-based metrics, ROC analysis, calibration, Brier score), simple and multiple resampling, statistical analysis (statistical testing, confidence intervals). Evaluating other supervised settings: regression analysis, image segmentation, time series analysis. Evaluation of clustering (stability analysis, indices).
Acquiring a first impression of datasets with the help of descriptive methods. Real-life example-based evaluation and comparison of various classification and regression methods, image segmentation approaches, time series and clustering techniques.

Compulsory/Recommended Readings:
- Nathalie Japkowicz, Zois Boukouvalas: Machine Learning Evaluation. Cambridge University Press, 2025. ISBN 978-1-316-51886-1
- Giovanni Cerulli: Fundamentals of Supervised Machine Learning. With Applications in Python, R, and Stata. Springer, 2023. ISBN 978-3-031-41336-0.
- Rudolf Scitovski, Kristian Sabo, Francisco Martínez-Álvarez, Šime Ungar: Cluster Analysis and Applications. Springer, 2021. ISBN 978-3-030-74551-6
- Marc Peter Deisenroth, A. Aldo Faisal, Cheng Soon Ong: Mathematics for Machine Learning. Cambridge University Press, 2020. ISBN 978-1-108-47004-9.

ADVANCED OPTIMIZATION METHODS

INMEA0103-26

Semester:	1
Type:	Laboratory
Number of Classes:	0+0+2
Credit:	3
Status:	Obligatory
Assessment:	Practical mark
Prerequisites:	None
Responsible:	Dr. Ágnes Éva Baran

Topics:

Theoretical background of gradient-based methods, basics of constrained optimization, Lagrange multipliers, penalty functions.

Implementation and testing of the methods listed below on simple tasks. Analysis, testing and comparison of existing implementations for more complex, larger problems.

Unconstrained optimization, gradient method, variants of the gradient method in machine learning (momentum method, SGD, Adagrad, RMSprop, Adam, etc.). Newton's method, solving large-scale problems. Stochastic optimization. Constrained optimization, penalty functions. Regularization. Convex optimization.

Compulsory/Recommended Readings:

- M.J.Kochenderfer, T. A. Wheeler: Algorithms for Optimization, The MIT Press, 2019, ISBN: 9780262039420
- S.J Wright, B. Recht, Optimization for Data Analysis, Cambridge University Press, 2022, ISBN: 9781009004282
- S.Sra, S. Nowozin, S.J. Wright, Optimization for Machine Learning, The MIT Press, 2011, ISBN: 9780262537766
- Y. Nesterov, Lectures on Convex Optimization, Springer, 2018, ISBN: 978-3-319-91577-7
- J. P. Wheeler, An Introduction to Optimization with Applications in Machine Learning and Data Analytics, Chapman and Hall, 2023, ISBN: 9780367425500

THEORY OF NEURAL NETWORKS

INMEA0104-26

Semester:	1
Type:	Lecture
Number of Classes:	2+0+0
Credit:	3
Status:	Obligatory
Assessment:	Exam
Prerequisites:	None
Responsible:	Dr. István Fazekas

Topics:
Perceptron model and convergence theorem; MLP structure; activation and loss functions; gradient method, back-propagation, momentum, conjugate gradient, quasi-Newton methods, Levenberg–Marquardt method; regularization;deep learning and vanishing gradient problem; cross-entropy; Convolutional neural networks (CNN) structure, notable CNN architectures; Support vector machines (SVM), optimal hyperplane, Kuhn–Tucker conditions; SVM for separation and regression; kernel trick and kernel functions; Recurrent neural networks (RNN), LSTM.

Performing mathematical derivations; analytical examination of the operation of neural networks and optimization procedures; analysis of the properties of activation and loss functions; evaluation of simple neural network models by manual calculation; solving SVM tasks; practical understanding of the application of kernel functions; theoretical examination of the operation of convolutional networks.

Compulsory/Recommended Readings:

- Charu C. Aggarwal: Neural Networks and Deep Learning, Springer, 2023. ISBN978-3-031-29644-4
- Ian Goodfellow, Yoshua Bengio, Aaron Courville: Deep Learning. MIT Press, 2016. ISBN: 0262035618
- S. Haykin: Neural Networks. A Comprehensive Foundation. Prentice Hall. New Jersey, 1999. ISBN 978-0132733502
- V.N. Vapnik: Statistical Learning Theory. Wiley, 1998. ISBN: 978-0-471-03003-4
- Marc Peter Deisenroth, A. Aldo Faisal, Cheng Soon Ong: Mathematics for Machine Learning. Cambridge University Press, 2020 ISBN 9781108455145
- Hastie, T., Tibshirani, R., Friedman, J.: The Elements of Statistical Learning. Springer, 2009. ISBN 978-0-387-84857-0
- Schmidhuber, J.: Deep Learning in Neural Networks: An Overview. Neural Networks, 61, 2015

OPERATIONS RESEARCH

INMEA0207-26

Semester:	2
Type:	Laboratory
Number of Classes:	0+0+2
Credit:	3
Status:	Obligatory
Assessment:	Practical mark
Prerequisites:	None
Responsible:	Dr. Ágnes Éva Baran

Topics:

Algorithms from the area of discrete optimization, integer programming, dynamic programming, and network flow problems.

Implementation and testing of algorithms, as well as the acquisition of problem modelling skills around discrete optimization, integer programming, dynamic programming, graph algorithms, network flow problems. Analysis of the efficiency of algorithms related to various problems.

Compulsory/Recommended Readings:

- W. L. Winston, Operations research: Applications and algorithms, 4th Ed., Cengage Learning, 2003, ISBN: 978-0534380588
- L.A. Wolsey, Integer programming, Wiley, 2020, ISBN: 978-1-119-60653-6
- H. P. Williams, Model Building in Mathematical programming, Wiley, 2013, ISBN: 978-1-118-44333-0
- M. Jünger at al, 50 Years of Integer Programming, 1958-2008, Springer, 2010, ISBN: 978-3-662-50181-8
- H. Kellerer, U. Pferschy, D. Pisinger, Knapsack Problems, Springer, 2004, ISBN: 978-3-642-07311-3

MACHINE LEARNING

INMEA0105-26

Semester:	1
Type:	Lecture / Laboratory
Number of Classes:	2+0+2
Credit:	6
Status:	Obligatory
Assessment:	Exam
Prerequisites:	None
Responsible:	Dr. András Hajdu

Topics:

Fundamental concepts and learning paradigms of machine learning. Steps of data-driven modelling: data preparation, model construction, and performance evaluation. Theory of model validation, overfitting, and underfitting. Principles behind regression and classification methods. Decision trees and ensemble techniques. Theoretical foundations of clustering and dimensionality reduction. Basics of model interpretability and ethical considerations.

Applying data preparation and exploratory analysis techniques to real datasets. Building, training, and evaluating machine learning models using appropriate metrics. Implementing model validation procedures (train–test split, cross-validation). Implementing and comparing regression, classification, decision tree, and ensemble models. Applying clustering and dimensionality reduction methods in practice. Using model interpretation tools (e.g., feature importance). Considering ethical and responsible modelling aspects in practical task.

Compulsory/Recommended Readings:

- Aurélien Géron: Hands-On Machine Learning with Scikit-Learn, Keras & TensorFlow (3rd Edition), O'Reilly Media, 2022. ISBN: 1098125975
- Kevin P. Murphy: Probabilistic Machine Learning: An Introduction, MIT Press, 2022. ISBN: 9780262046824
- Trevor Hastie, Robert Tibshirani, Jerome Friedman: The Elements of Statistical Learning (2nd Edition), Springer, 2021. ISBN: 0387848576
- Christoph Molnar: Interpretable Machine Learning (2nd Edition), 2022. ISBN: 9798411463330

ADVANCED MACHINE LEARNING MODELS

INMEA0106-26

Semester:	1
Type:	Laboratory
Number of Classes:	0+0+2
Credit:	3
Status:	Obligatory
Assessment:	Practical mark
Prerequisites:	None
Responsible:	Dr. Balázs Harangi

Topics:

The concept and properties of weak learners; the theory of ensemble models; the operation and convergence of bagging and random forests; Theoretical background of Boosting and AdaBoost; GradBoost, XGBoost, and their modern variants; density-based clustering (DBSCAN, HDBSCAN); NLP principles: tokenization, embeddings, classical language models.

Implementation of Bagging and Boosting models in Python (sklearn, XGBoost, LightGBM); Practical application of Random Forest; HDBSCAN clustering on real data; Performance measurement (precision, recall, ROC, AUC); Solving simple NLP tasks (text classification, handling embeddings).

Compulsory/Recommended Readings:

- Trevor Hastie, Robert Tibshirani, Jerome Friedman: The Elements of Statistical Learning (2nd Edition), Springer, 2021. ISBN: 0387848576
- Aurélien Géron: Hands-On Machine Learning with Scikit-Learn, Keras & TensorFlow (3rd Edition), O'Reilly Media, 2022. ISBN: 1098125975
- Chen, T., Guestrin, C.: XGBoost: A Scalable Tree Boosting System, KDD 2016
- Breiman, L.: Random Forests, Machine Learning, 2001.
- McInnes, L. et al.: HDBSCAN: Hierarchical Density-Based Clustering, JOSS, 2017.
- Goldberg, Y.: Neural Network Methods for Natural Language Processing, Morgan & Claypool, 2017. ISBN: 9783031010378

DEEP LEARNING

INMEA0208-26

Semester:	2
Type:	Lecture / Laboratory
Number of Classes:	2+0+2
Credit:	6
Status:	Obligatory
Assessment:	Exam
Prerequisites:	None
Responsible:	Dr. András Hajdu

Topics:

Operating principles and training mechanisms of artificial neural networks. Theory and application areas of convolutional neural networks (CNNs). Structure and functioning of recurrent architectures (RNN, LSTM, GRU) for time-series and language tasks. Principles of object detection, with a focus on YOLO-based architectures. Theoretical background of semantic image segmentation (U-Net, encoder–decoder models). Concepts and uses of vector representations and embeddings. Theoretical foundations of neural recommender systems. Principles of Siamese networks and metric learning, especially for similarity measurement and face recognition.

Implementing and training neural networks using modern deep learning frameworks (e.g., TensorFlow, PyTorch). Applying CNNs to image processing and visual recognition tasks. Using RNN, LSTM and GRU architectures for time-series and natural language data. Training and evaluating object detection models (e.g., YOLO) on real-world image datasets. Implementing and applying image segmentation models such as U-Net and encoder–decoder architectures. Generating and utilizing embeddings in various modelling tasks. Building neural recommender systems in practice. Apply-ing Siamese networks and metric learning methods to similarity-based tasks.

Compulsory/Recommended Readings:

- Ian Goodfellow, Yoshua Bengio, Aaron Courville: Deep Learning, MIT Press, 2016. ISBN: 9780262035613
- Aurélien Géron: Hands-On Machine Learning with Scikit-Learn, Keras & TensorFlow (3rd Edition), O'Reilly Media, 2022. ISBN: 1098125975
- François Chollet: Deep Learning with Python (2nd Edition), Manning Publications, 2021. ISBN: 1617296864
- Aston Zhang, Zachary C. Lipton, Mu Li, Alexander J. Smola: Dive into Deep Learning, Cambridge University Press, 2023. ISBN: 1009389432

NATURAL LANGUAGE PROCESSING AND SPEECH PROCESSING

INMEA0209-26

Semester:	2
Type:	Laboratory
Number of Classes:	0+0+2
Credit:	3
Status:	Obligatory
Assessment:	Practical mark
Prerequisites:	None
Responsible:	Dr. Róbert Lakatos

Topics:
Particular emphasis is placed on technologies relevant to the industrial environment, such as end-to-end systems, self-supervised repre-sentations, and multimodal applications. Through practical examples, students will learn about the entire life cycle of language models, from pre-training and fine-tuning (e.g., DPO and RLHF) through prompting methods to deployments and efficiency issues.
The course covers in detail the functioning of generative models, text generation (e.g., summarization, translation), and the structure of question-answering (QA) systems. It also provides insight into the latest speech recognition architectures (e.g., Whisper models) and the design of speech agents. Students will also critically examine open problems in the field, ethical considerations, and model evaluation methods.

Compulsory/Recommended Readings:

- Jacob Eisenstein: Introduction to Natural Language Processing, MIT Press, 2019., ISBN: 9780262042840
- Yoav Goldberg: A Primer on Neural Network Models for Natural Language Processing, Morgan & Claypool Publishers, 2017, https://arxiv.org/abs/1510.00726
- Delip Rao and Brian McMahan: Natural Language Processing with PyTorch (requires Stanford login), O'Reilly Media, 2020., ISBN: 9781491978221
- Lewis Tunstall, Leandro von Werra, and Thomas Wolf: Natural Language Processing with Transformers, O'Reilly Media, 2022., ISBN: 9781098136789
- Dan Jurafsky and James H. Martin: Speech and Language Processing (2024 pre-release), (online), 2024.
- Michael A. Nielsen: Neural Networks and Deep Learning, Determination Press, 2015. (online)
- Eugene Charniak: Introduction to Deep Learning, MIT Press, 2018., ISBN: 9780262039512

GENERATIVE ARTIFICIAL INTELLIGENCE

INMEA0210-26

Semester:	2
Type:	Laboratory
Number of Classes:	0+0+2
Credit:	3
Status:	Obligatory
Assessment:	Practical mark
Prerequisites:	None
Responsible:	Dr. Gergő Bogacsovics

Topics:

The course also covers the principles of modern transformer-based generative models (e.g., language and multimodal models), conditional generation methods, and state-of-the-art parameter-efficient fine-tuning techniques (e.g., LoRA, adapter layers).

The course takes a practice-oriented approach to showing how generative models can be used to solve various tasks, including image and text generation, data augmentation, style transfer, and synthetic data generation, as well as how to measure the quality and diversity of generated content (e.g., FID, IS, CLIP-score). Emphasis is also placed on issues related to the responsible use of generative methods, the recognition of distortions, the evaluation of realism and diversity, and technological challenges.

Compulsory/Recommended Readings:

- Ian Goodfellow, Yoshua Bengio, Aaron Courville: Deep Learning, MIT Press, 2016. ISBN: 978-0262035613
- David Foster: Generative Deep Learning: Teaching Machines to Paint, Write, Compose, and Play, 2nd Edition, O'Reilly Media, 2023. ISBN: 978-1098134181
- Jakub M. Tomczak: Deep Generative Modeling, Springer, 2024. ISBN: 978-3031640865
- Kailash Ahirwar: Generative Adversarial Networks Projects, Packt Publishing, 2019. ISBN: 978-1789136678
- Omar Sanseviero, Pedro Cuenca, Apolinario Passos, Jonathan Whitaker: Hands-On Generative AI with Transformers and Diffusion Models, O'Reilly Media, 2024. ISBN: 978-1098149246

AI ENGINEERING / MLOPS

INMEA0211-26

Semester:	2
Type:	Laboratory
Number of Classes:	0+0+2
Credit:	3
Status:	Obligatory
Assessment:	Practical mark
Prerequisites:	None
Responsible:	Dr. Tamás Márton Bérczes

Topics:

The theoretical part of the course presents the entire life cycle of machine learning systems, from data collection and preprocessing to model building and evaluation, through to deployment and continuous operation. Students will learn about the operating principles of automated and scalable ML and data pipelines, the practice of model versioning and model artifact management, and the theoretical background of drift and anomaly detection. The course covers the operation of modern serving architectures, including Triton Inference Server and various batch and online prediction systems, and reviews the characteristics of CI/CD pipelines and their role in the reliable operation of AI systems.

During the practical sessions, students will create data and training pipelines in a real MLOps environment using various tools such as MLflow, Airflow, and Kubeflow. The aim of the exercises is to enable students to version and register models and deploy them using various server solutions (REST, gRPC, Triton). The course emphasizes the development of monitoring and alerting systems that can be used to track model performance and detect anomalies in a timely manner. By the end of the semester, students will implement a complex, end-to-end MLOps project that includes full lifecycle automation and documentation.

Compulsory/Recommended Readings:

- Mark Treveil & Dataiku Team: Introducing MLOps. O'Reilly, 2020., ISBN: 9781492083283
- Chip Huyen: Designing Machine Learning Systems. O'Reilly, 2022., ISBN: 9781098107956
- Emmanuel Ameisen: Building Machine Learning Powered Applications. O'Reilly, 2020., ISBN: 9781492045106
- Prema & Patrick Deziel: Real-Time Machine Learning. Manning, 2025., ISBN: 9781633435735

EXPLAINABLE AI

INMEA9914-26

Semester:	3
Type:	Laboratory
Number of Classes:	0+0+2
Credit:	3
Status:	Optional
Assessment:	Practical mark
Prerequisites:	None
Responsible:	Dr. János Tóth

Topics:

Basic concepts, goals, and justification of explainable AI; main approaches to interpreting machine learning and deep learning models; and issues of prejudice, biases, and fairness in AI systems.

Application of model-agnostic interpretability tools (LIME, SHAP, Partial Dependence Plots (PDP)) on different machine learning models; saliency-based methods and the use of Grad-CAM on convolutional networks; interpretation of attention mechanisms in transformer-based models; application of the TCAV method for analyzing high-level concepts; comparison, evaluation, and practical interpretation of different explanations; detection of biases and fairness-related issues, and the use of XAI techniques in such analyses

Compulsory/Recommended Readings:

- C. Molnar: Interpretable Machine Learning: A Guide for Making Black Box Models Explainable (3rd ed.), ISBN: 9783911578035, 2025.
- U. Kamath, J. Liu: Explainable Artificial Intelligence: An Introduction to Interpretable Machine Learning (1st ed.), ISBN: 9783030833565, 2021
- W. Samek, G. Montavon, A. Vedaldi, L.K. Hansen, K.R. Müller: Explainable AI: Interpreting, Explaining and Visualizing Deep Learning, Springer Cham, ISBN: 9783030289539, 2019.
- S. Barocas, M. Hardt, A. Narayanan: Fairness and Machine Learning – Limitations and Opportunities, MIT Press, ISBN: 9780262048613, 2023

AI AND CYBERSECURITY

INMEA9915-26

Semester:	4
Type:	Lecture / Laboratory
Number of Classes:	2+0+2
Credit:	6
Status:	Optional
Assessment:	Exam
Prerequisites:	None
Responsible:	Dr. Andrea Pintér-Huszti

Topics:

After introducing fundamental cybersecurity concepts and basic security controls, students will learn about cybersecurity threats and adversary models targeting artificial intelligence–based systems. Students will gain an overview of attacks against AI integrity, including data poisoning, model poisoning, and attacks on the AI supply chain. The course provides insight into attacks against AI confidentiality as well, such as model inversion, membership inference, and model extraction. Students will also become familiar with different countermeasures against these attacks. In addition, the course addresses cybersecurity domains where AI-based support is applied, such as spam email detection, malware detection, or log analysis.

During the lab, students experiment with simulations of prompt injection, data and model leakage, model theft, overfitting, and faulty output handling, as well as basic defense methods against them. The goal of the lab is for students to understand how cybersecurity threats related to modern AI systems work, their risks, and the relevant defense controls and evaluation methods that can be applied to them. The aim of the lab is to provide students with practical experience in examining the security challenges of machine learning through various real-world tasks: spam and phishing detection, malware classification (decision tree, random forest), malware clustering (K-means), network anomaly detection (Gaussian networks, statistical models), and biometric and behavior-based identification (face recognition, keystroke analysis). Students also model attack scenarios—data poisoning, model poisoning, supply chain attacks—and test various defense techniques.

Compulsory/Recommended Readings:

- Naoto Kiribuchi, Kengo Zenitani, Takayuki Semitsu, Securing AI Systems: A Guide to Known Attacks and Impacts, https://doi.org/10.48550/arXiv.2506.23296
- Alessandro Parisi, Hands-On Artificial Intelligence for Cybersecurity: Implement smart AI systems for preventing cyber attacks and detecting threats and network anomalies, ISBN-1789804027, Packt Publishing, 2019.
- Dr. Enrico GLEREAN, Training curriculum on AI and data protection Fundamentals of Secure AI Systems with Personal Data, 2025, European Data Protection Board
- Vassilev, A. – Oprea, A. – Fordyce, A. – Anderson, H., Adversarial Machine Learning: A Taxonomy and Terminology of Attacks and Mitigations, https://doi.org/10.6028/NIST.AI.100-2e2023

AI ETHICS AND GOVERNANCE

INMEA9916-26

Semester:	2
Type:	Lecture
Number of Classes:	2+0+0
Credit:	3
Status:	Optional
Assessment:	Exam
Prerequisites:	None
Responsible:	Dr. Miklós Hoffmann

Topics:

AI legal frameworks (AI Act, GDPR, DSA and liability rules), risk-based classification, obligations, compliance requirements, data minimization, automated decision-making, AI governance and corporate governance models; AI ethical frameworks (IEEE, OECD, EU), ethical dilemmas in the application of AI; algorithmic bias and discrimination, biased data, biased metrics, biased objective function, fairness metrics (demographic parity, equalized odds, etc.), bias-mitigation techniques, copyright and content ownership issues, value-based and risk-sensitive decision-making, ethical innovation, "Responsible-by-design" approach.

Analysis and evaluation of various case studies.

Compulsory/Recommended Readings:

- Mark Coeckelbergh: AI Ethics, MIT Press, 2021. ISBN: 0262538199
- Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 on harmonised rules on artificial intelligence
- Virginia Eubanks: Automating Inequality, St. Martins Publishing, 2018, ISBN 1250074312
- Giovanni Ziccardi: Legal Informatics, Edward Elgar Publishing, 2025. ISBN: 1035321157
- Paula Boddington: AI Ethics: A Textbook, Springer, 2023. ISBN: 9811993815

VISUALIZATION AND VISUAL ANALYTICS

INMEA9917-26

Semester:	3
Type:	Laboratory
Number of Classes:	0+0+2
Credit:	3
Status:	Optional
Assessment:	Practical mark
Prerequisites:	None
Responsible:	Dr. Roland Imre Kunkli

Topics:

Definition of data visualization and its role in data analysis and decision support. Human visual perception. Data and task abstraction. Visual encoding and the effectiveness of different marks and channels. Data-type-specific visualization techniques. Interaction techniques. Basic principles of dashboard design. Visualization options for data related to AI workflows. Visualizations generated or interpreted by AI.

Data preparation for visualization purposes. Creating static and interactive visualizations using selected libraries and software tools. Designing dashboards with interactive filtering and multiple coordinated views. Practical implementation of data-type-specific visualizations. Approaches to visualizing large-scale datasets.

Compulsory/Recommended Readings:

- Tamara Munzner: Visualization Analysis and Design. A K Peters/CRC Press, 2014. ISBN 978-1466508910.
- Claus O. Wilke: Fundamentals of Data Visualization: A Primer on Making Informative and Compelling Figures. O'Reilly Media, 2019. ISBN 978-1492031086.
- Kieran Healy: Data Visualization: A Practical Introduction. Princeton University Press, 2018. ISBN 978-0691181622.
- Matthew O. Ward, Georges Grinstein, Daniel Keim: Interactive Data Visualization: Foundations, Techniques, and Applications (2nd edition). A K Peters/CRC Press, 2021. ISBN 978-0367783488.
- Colin Ware: Information Visualization: Perception for Design (4th edition). Morgan Kaufmann, 2020. ISBN 978-0128128756

GRAPH-BASED NEURAL NETWORKS

INMEA9918-26

Semester:	3
Type:	Laboratory
Number of Classes:	0+0+2
Credit:	3
Status:	Optional
Assessment:	Practical mark
Prerequisites:	None
Responsible:	Dr. Attila Tiba

Topics:

The theoretical part of the course provides a comprehensive overview of the mathematical and computational representation of graphs, the properties that define network structures, and the operating principles of classical and modern graph embedding techniques. Students will learn about different architectures of graph neural networks, including spectral and spatial graph convolution models and the message passing paradigm, and will gain insight into the theory of oversmoothing, regularization, and the handling of heterogeneous or temporal graph structures. The course also covers the theoretical connections that enable the application of GNNs in a wide range of industrial and research problems.

During the exercises, students will perform node and edge classification tasks, link prediction, and graph classification on real datasets using modern libraries, in particular PyTorch Geometric. The aim of the lab sessions is to give students an understanding of the practical process of processing graph-structured data: from data preparation to model construction and training to the interpretation of network patterns. The course places great emphasis on practical problem solving and prepares students to independently apply GNN models in research or industrial settings.

Compulsory/Recommended Readings:

- Hamilton, W. L. (2020): Graph Representation Learning. Morgan & Claypool, ISBN-13: 978-1681739632
- Bronstein, M. M., Bruna, J., Cohen, T., Veličković, P. (2021): Geometric Deep Learning. arXiv preprint.
- Newman, M. E. J. (2018): Networks. Oxford University Press, ISBN-13: 978-0198805090
- Wu et al. (2021): A Comprehensive Survey on Graph Neural Networks. IEEE TNNLS.
- Barabási, A.-L. (2016): Network Science. Cambridge University Press, ISBN-13: 978-1107076266
- Lovász, L. (2012): Large Networks and Graph Limits. AMS, ISBN-13: 978-0821890851.

SIMULATIONS – DIGITAL TWIN

INMEA9919-26

Semester:	3
Type:	Laboratory
Number of Classes:	0+0+2
Credit:	3
Status:	Optional
Assessment:	Practical mark
Prerequisites:	None
Responsible:	Tibor Péter Kapusi

Topics:

The theoretical instruction reviews the basic concepts of digital twins, the theoretical challenges of the transition from simulation to reality (sim-to-real), and the role of synthetic data generation in modern machine learning pipelines. Students will learn about the functioning of domain randomization methods, the reasons for and effects of applying geometric, lighting, texture, and background variations, and the theoretical foundations of sensor modeling, with a particular focus on the functioning and distortions of RGB, depth, LiDAR, and IMU systems. The course covers the accuracy and consistency of synthetic annotations, as well as theoretical approaches to domain adaptation—including GAN-based style transfer—that serve to reduce differences between simulated and real environments. The course covers the role of simulations in reinforcement learning, particularly in establishing dynamic models, reward structures, and real-time interfaces.

During practical classes, students work with industrial-grade simulation toolchains, such as the NVIDIA Omniverse environment, where they create photorealistic, parameterizable scenes and large amounts of synthetic data sets. During lab assignments, they use Unreal or Unity engines to create interactive, physics-based simulations and build specialized, custom-developed simulation modules. The practical part of the course includes configuring sensor simulations, modeling the distortions of different sensors, automatically generating annotations, and ensuring their quality. During the semester-long project, students create a complete simulation pipeline, measure data quality, analyze the size of the domain gap, and evaluate the impact of adaptation methods on real-world tasks.

Compulsory/Recommended Readings:

- Fei Tao, Ang Liu, Tianliang Hu, A.Y.C. Nee: Digital Twin Driven Smart Manufacturing. Academic Press, 2020, ISBN: 978-0-12-818918-4.
- Sergey I. Nikolenko: Synthetic Data for Deep Learning. Springer, 2021, ISBN: 978-3-030-75177-7.
- Gabriela Csurka (ed.): Domain Adaptation in Computer Vision Applications. Springer, 2017, ISBN: 978-3-319-58346-4.
- Jason Gregory: Game Engine Architecture. CRC Press, 2018, ISBN: 9781138035454.
- Richard Szeliski: Computer Vision: Algorithms and Applications. Springer, 2022, ISBN: 978-3-030-34371-2.
- Pharr, M., Jakob, W., Humphreys, G.: Physically Based Rendering. MIT Press, 2023, ISBN: 9780262048026.

REINFORCEMENT LEARNING

INMEA9920-26

Semester:	2
Type:	Lecture / Laboratory
Number of Classes:	2+0+2
Credit:	6
Status:	Optional
Assessment:	Exam
Prerequisites:	None
Responsible:	Dr. Gergő Bogacsovics

Topics:

The theoretical part of the course introduces the formal framework based on Markov decision processes, the mathematical foundations of feedback decision-making, and the role of reward functions. Students will learn about classical value-based algorithms, including Q-learning and its deep learning variants, as well as the theoretical operation of policy-based and combined approaches. The course explains the stability and convergence characteristics of modern policy gradient techniques such as PPO, A2C, and A3C, as well as the principles of imitation learning, hierarchical RL, partially observable environments, and multi-agent systems. It addresses issues of sample efficiency and robust operation, as well as modeling techniques that bridge the gap between simulation and reality, including the theoretical background of RLHF (Reinforcement Learning from Human Feedback).

Students will learn about the most important RL algorithms in a practical environment and implement agents in various environments using modern libraries such as PyTorch or TensorFlow. The exercises will cover the construction of network architectures, the implementation of loss functions, the application of stability techniques, and experimentation in realistic, sometimes simulation-based environments. During the semester-long project, students will develop a complete RL solution to a selected problem, which includes data collection, model training, performance measurement, and interpretation of results.

Compulsory/Recommended Readings:

- Sutton, R. S., Barto, A. G.: Reinforcement Learning: An Introduction. MIT Press, 2018. ISBN: 978-0262039246
- Lapan, M.: Deep Reinforcement Learning Hands-On. Packt Publishing, 2024. ISBN: 978-1835882702
- Bilgin, E.: Mastering Reinforcement Learning with Python. Packt Publishing, 2020. ISBN: 978-1838644147
- Szepesvári, Cs.: Algorithms for Reinforcement Learning. Springer, 2022. ISBN: 978-3031015519

DISTRIBUTED AI SYSTEMS

INMEA9921-26

Semester:	3
Type:	Laboratory
Number of Classes:	0+0+2
Credit:	3
Status:	Optional
Assessment:	Practical mark
Prerequisites:	None
Responsible:	Dr. Róbert Lakatos

Topics:

In the theoretical part, students will learn about the structure of systems composed of autonomous agents, communication and coordination mechanisms, and algorithms describing cooperative and competitive behaviors. The course introduces the concept of federated learning, which enables decentralized model training by keeping data local and aggregating only model updates. The course reviews the basics of scalable AI infrastructures, the operation of high-performance distributed computing frameworks, and the applied data protection and algorithmic fairness principles that define the ethical and sustainable operation of modern distributed AI systems.

During the practical classes, students will implement AI pipelines using frameworks that support distributed computing, such as TensorFlow, PySyft, FedML, Apache Spark, and Dask. During the assignments, they will develop and operate decentralized learning prototypes, in which they will also test the aggregation of model updates, joint model building, and the practical implementation of data protection mechanisms. The practice involves the use of scalable, cloud-based data storage and processing solutions, as well as the assembly of architectures suitable for the simultaneous handling of large amounts of data and multiple parallel agents. The goal of the semester-long project is to implement a working prototype of a distributed AI system or federated learning application.

Compulsory/Recommended Readings:

- Raieli, S. & Iuculano, G.: Building AI Agents with LLMs, RAG, and Knowledge Graphs. Packt Publishing, 2025., ISBN: 9781835087060
- Nakayama, K., & Jeno: Federated Learning with Python. 2022., ISBN: 9781803247106
- ur Rehman, M. H., & Gaber, M. M. (eds.): Federated Learning Systems – Towards Next-Generation AI. Springer, 2021., ISBN: 9783030706043
- Kleppmann, M.: Designing Data-Intensive Applications. O'Reilly, 2021., ISBN: 9781491903063
- Jin, Y., Zhu, H., Xu, J., Chen, Y.: Federated Learning – Fundamentals and Advances. Springer, 2023., ISBN: 9789811970825

MULTIMODAL AI TOOLS

INMEA9922-26

Semester:	3
Type:	Laboratory
Number of Classes:	0+0+2
Credit:	3
Status:	Optional
Assessment:	Practical mark
Prerequisites:	None
Responsible:	Dr. Balázs Harangi

Topics:

The concept and significance of multimodality; modal representations and their coordination; architectures for processing multimodal data (late/early fusion, joint embedding models); multimodal extensions of transformer and large language models; data integration techniques; reliability and bias in multimodal systems.

Implementation of multimodal models using Python/PyTorch; preparation of data from different modalities; coordination of image and audio data streams; practical implementation of multimodal classification and detection; use and fine-tuning of pre-trained models (CLIP, ImageBind, etc.); validation of results.

Compulsory/Recommended Readings:

- Tsai, Y.-H. et al.: Multimodal Machine Learning: A Survey and Taxonomy. ACM Computing Surveys, 2020.
- Radford, A. et al.: Learning Transferable Visual Models From Natural Language Supervision (CLIP). arXiv:2103.00020.
- Girdhar, R. et al.: ImageBind: One Embedding Space to Bind Them All. arXiv:2305.05665
- Baltrušaitis, T., Ahuja, C., Morency, L.-P.: Multimodal Machine Learning: A Survey. IEEE TPAMI, 2019.
- Ian Goodfellow, Yoshua Bengio, Aaron Courville: Deep Learning, MIT Press, 2016. ISBN: 9780262035613
- Paszke, A. et al.: PyTorch: An Imperative Style, High-Performance Deep Learning Library, NeurIPS 2019

DATA MINING

INMEA9923-26

Semester:	2
Type:	Laboratory
Number of Classes:	0+0+2
Credit:	3
Status:	Optional
Assessment:	Practical mark
Prerequisites:	None
Responsible:	Dr. Márton Ispány

Topics:

The concept and process of data mining. Exploratory data analysis and preprocessing. Evaluation of data mining models. Advanced supervised learning: kernel methods, vector machines, Gaussian processes, graphical models, Bayesian networks, random Markov fields. Ensemble and additive models, decision trees, random forests, stochastic optimizers. Curse of dimensionality, dimension reduction methods. Advanced unsupervised learning: mixture models and the EM algorithm, bisecting K-means clustering, OPTICS, BIRCH, spectral clustering. Sampling methods and the MCMC algorithm. Sequential data, Markov chains, hidden Markov models, linear dynamic systems.

Learning how to use a Python-based data mining library, e.g., scikit-learn, at an advanced level. Creating and applying an independent data mining pipeline and interpreting the results.

Compulsory/Recommended Readings:

- Charu C. Aggarwal: Data Mining. The Textbook. Springer, 2016. ISBN: 3319381164
- Pang-Ning Tan, Michael Steinbach, Vipin Kumar: Introduction to Data Mining, 2nd edition. Pearson Education, 2019. ISBN: 9780273775324
- Christopher M. Bishop: Pattern Recognition and Machine Learning, Springer, 2006. ISBN: 978-1-4939-3843-8
- Trevor Hastie, Robert Tibshirani, Jerome Friedman: The Elements of Statistical Learning: Data Mining, Inference, and Prediction. Second Edition, Springer, 2009. ISBN: 978-0-387-84858-7
- Daniel Peña, Ruey S. Tsay: Statistical Learning for Big Dependent Data, Wiley 2021. ISBN: 9781119417385

ADVANCED DATA MANAGEMENT

INMEA9924-26

Semester: 3
Type: Laboratory
Number of Classes: 0+0+2
Credit: 3
Status: Optional
Assessment: Practical mark
Prerequisites: None
Responsible: Dr. János Tóth

Topics:

Principles of managing large datasets; columnar storage formats and data lake architectures; NoSQL data model-based storage approaches; operational models of distributed data processing frameworks; fundamentals of streaming data processing; structure and optimization aspects of cloud-based data warehouses.

Preparation and processing of structured and semi-structured data, including the use of columnar storage formats (e.g., Parquet, ORC) and data lake–based data management operations. Use of NoSQL databases (e.g., MongoDB, Bigtable), implementation of aggregation operations, fundamental data modeling principles, and indexing techniques. Working with distributed processing frameworks (e.g., Apache Spark, Dask); DataFrame-based transformations, join and partitioning strategies, and the fundamentals of streaming processing. Use of cloud-based data warehouses (e.g., BigQuery); management of partitioned and clustered tables; designing cost-efficient queries.

Compulsory/Recommended Readings:

- J. Reis, M. Housley: Fundamentals of Data Engineering: Plan and Build Robust Data Systems (1st ed.), O'Reilly Media, ISBN: 9781098108298, 2022.
- J.S. Damji, B. Wenig, T. Das, D. Lee: Learning Spark: Lightning-Fast Data Analytics (2nd ed.), O'Reilly Media, ISBN: 9781492050032, 2020.
- M. Kleppmann: Designing Data-Intensive Applications: The Big Ideas Behind Reliable, Scalable, and Maintainable Systems (1st ed.), O'Reilly Media, ISBN: 9781491903063, 2021.
- R. Kimball, M. Ross: The Data Warehouse Toolkit: The Definitive Guide to Dimensional Modeling (3rd ed.), Wiley, ISBN: ISBN: 9781118530801, 2013.

COMPUTER VISION

INMEA9925-26

Semester:	2
Type:	Lecture / Laboratory
Number of Classes:	2+0+2
Credit:	6
Status:	Optional
Assessment:	Exam
Prerequisites:	None
Responsible:	Dr. András Hajdu

Topics:

Fundamentals of digital image processing and video processing. Modern object recognition and object detection techniques based on neural networks. Principles of multiview geometry and 3D reconstruction. Deep learning approaches for visual reconstruction tasks. Theoretical foundations of biometric identification and authentication using image processing and AI. Operating principles, application areas, and limitations of computer vision systems in industrial and autonomous domains.

Implementing image and video processing tasks on real datasets. Building and evaluating neural-network-based object recognition and detection models. Performing 3D reconstruction from multiview data using geometric and deep learning techniques. Applying biometric identification and authentication algorithms in practice. Designing and integrating computer vision systems into industrial, security, or autonomous platforms. Using modern computer vision frameworks (e.g., OpenCV, PyTorch, TensorFlow) to solve complex visual tasks.

Compulsory/Recommended Readings:

- Richard Szeliski: Computer Vision: Algorithms and Applications (2nd Edition), Springer, 2022. ISBN: 3030343715
- Simon J.D. Prince: Computer Vision: Models, Learning, and Inference (2nd Edition), Cambridge University Press, 2023. ISBN: 1107011795
- Adrian Kaehler, Gary Bradski: Learning OpenCV 4: Computer Vision with Python (2nd Edition), O'Reilly Media, 2020. ISBN: 0596516134
- Mohamed Elgendy: Deep Learning for Vision Systems, Manning, 2021. ISBN: 1617296198

TEXT AND WEB MINING

INMEA9926-26

Semester:	3
Type:	Laboratory
Number of Classes:	0+0+2
Credit:	3
Status:	Optional
Assessment:	Practical mark
Prerequisites:	None
Responsible:	László Mészáros

Topics:

The theoretical part reviews web data collection methods, starting with the functioning of scraping and crawling techniques and the theoretical background of data services available through APIs. Students will learn the theoretical foundations necessary for processing unstructured data, such as HTML pages, PDF documents, JSON files, and image content, and will also receive an overview of classic and modern text processing techniques. The course also covers the operating principles of entity extraction, text structure recognition, document analysis, and text comprehension methods supported by large language models.

During the practical sessions, students collect data from real web sources, build automated processes using Python-based tools, and then clean, structure, and use the collected data for various analysis or machine learning tasks. In the exercises, students extract data from HTML, PDF, and JSON documents and use modern LLMs to perform entity extraction, text structuring, data enrichment, and summarization. The goal of the semester project is to create an automated web data collection and processing pipeline, which students will apply to a specific, real-world topic.

Compulsory/Recommended Readings:

- Bing Liu: Web Data Mining. Springer, 2011., ISBN: 9783642194597
- Lewis Tunstall, Leandro Von Werra, Thomas Wolf: Natural Language Processing with Transformers (Revised Edition). O'Reilly, 2022., ISBN: 9781098136789
- Ryan Mitchell: Web Scraping with Python (2nd Edition). O'Reilly, 2018., ISBN: 9781491985564
- Nocile Koenigstein: Transformers in Action. Manning, 2025., ISBN: 9781633437883

FUNDAMENTALS OF ROBOTICS

INMEA9927-26

Semester:	3
Type:	Laboratory
Number of Classes:	0+0+2
Credit:	3
Status:	Optional
Assessment:	Practical mark
Prerequisites:	None
Responsible:	Dr. Tien V. Do

Topics:

Robot operating system; software patterns in robot systems; navigation; map; machine vision; sensor-types; kinematics; kinematics and planning; human-robot interactions.

Navigation, localization, sensor and processing algorithms; software solutions and architecture for autonomous driving vehicles; SLAM technique; map solutions; cooperative intelligent transportation systems; AI technique in robot systems and C-ITS.

Compulsory/Recommended Readings:

- A. Koubaa (Editor): Robot Operating System (ROS): The Complete Reference (Volume 7) (Studies in Computational Intelligence, 1051) 1st ed. 2023 Edition, ISBN-13: 978-3031090615
- The ROS Community https://www.ros.org/
- computer vision library https://opencv.org/
- Open Source Autopilot https://px4.io/
- Navigation https://navigation.ros.org/
- L. Joseph, J. Cacace: Mastering ROS 2 for Robotics Programming - Fourth Edition, Packt Publishing, 2025. ISBN-13: 978-1836209010
- https://micro.ros.org/

AI IN HEALTH SCIENCES

INMEA9928-26

Semester:	4
Type:	Laboratory
Number of Classes:	0+0+2
Credit:	3
Status:	Optional
Assessment:	Practical mark
Prerequisites:	None
Responsible:	Dr. Balázs Harangi

Topics:

Types of healthcare data (imaging, genomic, clinical, time-dependent data); deep learning architectures in medical image processing (CNN, U-Net, Vision Transformer); predictive modeling for disease progression and diagnosis support; analysis of genomic data using machine learning; issues of data quality, bias, and reliability; AI regulation and ethical principles in healthcare.

Management of medical imaging databases (CT, MRI, X-ray); implementation of segmentation and classification models; preparation and modeling of genomic data; fine-tuning of predictive models; performance measurement and validation in a clinical environment; error analysis, interpretability methods (Grad-CAM, SHAP).

Compulsory/Recommended Readings:

- Litjens, G. et al.: A Survey on Deep Learning in Medical Image Analysis. Medical Image Analysis, 2017.
- LeCun, Y., Bengio, Y., Hinton, G.: Deep Learning. Nature, 2015.
- Eraslan, G. et al.: Deep Learning: New Computational Modelling Techniques for Genomics. Nature Reviews Genetics, 2019
- Esteva, A. et al.: A Guide to Deep Learning in Healthcare. Nature Medicine, 2019.
- François Chollet: Deep Learning with Python (2nd Edition), Manning Publications, 2021. ISBN: 1617296864
- Ribeiro, M. et al.: Why Should I Trust You? – Explainable AI in healthcare. KDD, 2016

AI IN PHYSICS

INMEA9929-26

Semester:	4
Type:	Laboratory
Number of Classes:	0+0+2
Credit:	3
Status:	Optional
Assessment:	Practical mark
Prerequisites:	None
Responsible:	Dr. Balázs Ujvári

Topics:

The type and structure of data from high-energy particle physics detectors. Machine learning methods in physical signal analysis: deep learning-based clustering and topological methods. Particle reconstruction based on AI, jet reconstruction using ML models. The role of GNNs (Graph Neural Networks) in track, cluster reconstruction. Noisy data handling and anomaly detection in high-resolution detector data. Handling large data streams (petabyte scale) in an experimental environment. Physical interpretability and model reliability: how can learned patterns be linked to detector physics?

Simulation techniques in particle physics: GEANT4. Preprocessing of high-energy particle physics raw data, channel merging, noise filtering. Implementation of clustering on real detector data. ML-based clusterers and deep clustering models (GNN). Particle reconstruction with AI models. Handling large data sets: ROOT, Jupyter + Python (NumPy, SciPy, Matplotlib, uproot). Detector calibration and anomaly detection with AI tools.

Compulsory/Recommended Readings:

- Shlomi, J., Battaglia, P., Vlimant, J.: Graph Neural Networks in Particle Physics. Machine Learning: Science and Technology, 2020.
- Chollet, F.: Deep Learning with Python. Manning, 2021. ISBN 9781617294433
- Bourilkov, D.: Machine and Deep Learning Applications in Particle Physics. 2020. https://doi.org/10.48550/arXiv.1912.08245
- Ilya Narsky & Frank C. Porter Statistical Analysis Techniques in Particle Physics: Fits, Density Estimation, and Supervised Learning ISBN 9783527410866
- Paolo Calafiura, David Rousseau & Kazuhiro Terao Artificial Intelligence for High Energy Physics ISBN 9789811234026
- Jona Motta: Development of Machine Learning τ Trigger Algorithms and Search for Higgs Boson Pair Production, ISBN: 9783031962875

DOCKER AND KUBER-NETES FOR AI

INMEA9930-26

Semester:	2
Type:	Laboratory
Number of Classes:	0+0+2
Credit:	3
Status:	Optional
Assessment:	Practical mark
Prerequisites:	None
Responsible:	László Mészáros

Topics:

In the theoretical part, students will learn the principles of containerization, the Docker operating model, and the components that enable the creation of a unified, isolated execution environment. The course covers GPU support architecture, the structure and optimization of Docker images, and solutions for reliably integrating various machine learning libraries and frameworks into containers. In addition, the basics of Kubernetes operation will be presented: cluster organization, the logic of pods and deploy structures, the role of service abstractions, and the approach that makes an AI system scalable, flexible to operate, and cloud-native.

In the practical sessions, students will build their own Docker images, configure environments for different machine learning frameworks, and create GPU-accelerated containers to serve real AI tasks. The central element of the exercise is the creation of a containerized inference service, which students deploy, scale, and monitor in a Kubernetes cluster. During the course, students will learn about container execution solutions from cloud providers such as AWS and Azure, and gain hands-on experience in how to deliver complex ML systems in a reproducible, production-ready form.

Compulsory/Recommended Readings:

- Nigel Poulton: Docker Deep Dive. Packt Publishing, 2020., ISBN: 9781800565135
- Nigel Poulton & Pushkar Joglekar: The Kubernetes Book (2nd Edition). Packt Publishing, 2024., ISBN: 9781835880302
- Faisal Masood & Ross Brigoli: Machine Learning on Kubernetes. Packt Publishing, 2022., ISBN: 9781803241807
- Kelsey Hightower, Brendan Burns & Joe Beda: Kubernetes: Up and Running. O'Reilly, 2017., ISBN: 9781491935668

ADVANCED CLOUD COMPUTING

INMEA9931-26

Semester:	2
Type:	Laboratory
Number of Classes:	0+0+2
Credit:	3
Status:	Optional
Assessment:	Practical mark
Prerequisites:	None
Responsible:	Dr. Tamás Márton Bérczes

Topics:

Students will learn the basic concepts of cloud computing, service (IaaS, PaaS, SaaS) and operating (public, private, hybrid, on-premise) models, and their role in modern IT systems, especially those supported by AI. They will review the theoretical background of reliability, availability, and scalability, the interpretation of SLA/SLO indicators, cost models, and the functioning of consumption-based pricing. At a theoretical level, they will learn about the concepts of virtualization and resource abstraction (virtual machines, networks, storage), the basic principles of risk analysis and disaster recovery (RPO, RTO, business continuity), and the specific requirements of cloud security and AI applications running in the cloud.

Students will learn how to create and configure the resources of typical cloud providers (virtual machines, networks, storage), deploy simple multi-layer applications to the cloud, and compare public, private, and hybrid clouds in practice. They will be able to apply availability and scaling mechanisms (e.g., load balancing, automatic scaling), perform basic cost estimation and cost optimization, use monitoring and logging tools, and implement backup, restore, and basic cloud security settings (IAM, network segmentation, encryption), with a particular focus on the runtime environment of resource-intensive, AI-type applications.

Compulsory/Recommended Readings:

- Thomas Erl, Ricardo Puttini, Zaigham Mahmood: Cloud Computing: Concepts, Technology & Architecture, ISBN: 978-0-13-338752-0
- Rajkumar Buyya, James Broberg, Andrzej M. Goscinski Cloud Computing: Principles and Paradigms. Wiley Series on Parallel and Distributed Computing, John ISBN: 978-0-470-88799
- Justin Domingus, John Arundel: Cloud Native DevOps with Kubernetes: Building, Deploying, and Scaling Modern Applications in the Cloud. 2nd edition, O'Reilly Media, Sebastopol, 2022. ISBN: 978-1-098-11682-8
- Kevin L. Jackson, Scott Goessling: Architecting Cloud Computing Solutions: Build Cloud Strategies That Align Technology and Economics While Effectively Managing Risk. Packt Publishing, Birmingham–Mumbai, 2018. ISBN: 978-1-78847-242-5

TOOLS FOR PARALLEL PROGRAMMING

INMEA9932-26

Semester:	3
Type:	Laboratory
Number of Classes:	0+0+2
Credit:	3
Status:	Optional
Assessment:	Practical mark
Prerequisites:	None
Responsible:	Tibor Péter Kapusi

Topics:

The theoretical part presents the basic principles that determine the nature of parallel execution: concurrent control models, differences between threads and processes, theoretical limitations of CPU-side acceleration, and the role of performance profiling. Students will gain an overview of how GPU architectures work, the organization of memory hierarchies, and programming patterns suitable for parallel algorithms. Special emphasis is placed on the theoretical model of OpenCL, which provides an understanding of the platform-device-kernel structure, as well as the organization of workgroups and its performance implications.

During the practical sessions, students implement various parallel execution models using Python-based libraries (threading, multiprocessing, IPC, asyncio) and measure the speed, efficiency, and resource utilization of the execution models. The practice also covers GPU-based acceleration, where students experiment with tensor operations, memory optimization, DataLoader tuning, and mixed-precision techniques in PyTorch and CuPy environments. They then create OpenCL-based implementations, using their own measurements to evaluate the impact of memory usage, local cache utilization, and set workgroup sizes. The course briefly covers parallel graph processing algorithms (e.g., BFS with CSR representation, PageRank), problems of irregular data access, and practical solutions for load balancing.

Compulsory/Recommended Readings:

- Richard D. McCool, Arch D. Robison, James Reinders: Structured Parallel Programming. Morgan Kaufmann, 2012, ISBN: 978-0124159938.
- Micha Gorelick, Ian Ozsvald: High Performance Python (2nd Edition). O'Reilly Media, 2020, ISBN: 978-1492055020.
- Aaftab Munshi et al.: OpenCL Programming Guide. Addison-Wesley, 2011, ISBN: 978-0-321-74964-2
- David B. Kirk, Wen-mei Hwu: Programming Massively Parallel Processors. Morgan Kaufmann, 2022, ISBN: 978-0-12-381472-2.
- Timothy G. Mattson et al.: Patterns for Parallel Programming. Addison-Wesley, 2004, ISBN: 978-0321228116

SOFTWARE DEVELOP-MENT IN INDUSTRIAL ENVIRONMENTS

INMEA9933-26

Semester:	3
Type:	Laboratory
Number of Classes:	0+0+2
Credit:	3
Status:	Optional
Assessment:	Practical mark
Prerequisites:	None
Responsible:	Dr. Attila Tamás Adamkó

Topics:

Software life cycle and architecture: concept and role in software systems, architecture vs. design. Architecture and quality attributes: performance, reliability, maintainability, scalability. Architectural styles: layered, client–server, multi-tier web applications, microservices, event-driven and message-based architectures. Architectural patterns (e.g. MVC, broker, pipes-and-filters) and their relationship to industrial frameworks.

Deriving quality requirements into architectural decisions. Architecture design process: basics of attribute-driven design and handling trade-offs. Documenting architecture: views, viewpoints, diagrams, templates. Evaluating architecture: basic evaluation techniques and risk identification. Architecture in the cloud: typical cloud architectures and service-oriented solutions. Industrial case studies: architectural analysis of existing systems and anti-patterns.

Compulsory/Recommended Readings:

- Mark Richards, Neal Ford: Fundamentals of Software Architecture, O'Reilly Media, 2020, 9781492043447
- Carola Lilienthal, Henning Schwentner: Domain-Driven Transformation, O'Reilly Media, 2025, 979-8341640122
- Simon Brown: The C4 Model, O'Reilly Media, 2026, 9798341660113.
- Evelyn van Kelle, Gien Verschatse, Kenny Baas-Schwegler: Collaborative Software Design: How to facilitate domain modeling decisions, Manning, 2025, 9781633439252
- Susanne Kaiser: Architecture for Flow, Addison-W, 2025, 978-0-13-739303-9

THEORETICAL AND NEURAL MODELS IN THE INDUSTRY

INMEA9934-26

Semester:	3
Type:	Laboratory
Number of Classes:	0+0+2
Credit:	3
Status:	Optional
Assessment:	Practical mark
Prerequisites:	None
Responsible:	Dr. Ágnes Éva Baran

Topics:

Mathematical modeling of the problems, breaking the problem into subtasks. Overview of the relevant machine learning methods. Interpretation of techniques used to evaluate solutions.

Implementing the models, preprocessing the data. Finding the optimal model and training parameters. Evaluating the results and comparing the methods, recognizing the advantages and disadvantages of the applied techniques.

Compulsory/Recommended Readings:

- Shubhabrata Datta, Paulo David, Machine Learning in Industry, Springer, 2022, ISBN: 978-3-030-75849-3
- Pedro Larranaga, David Atienza, Javier Diaz-Rozo, Alberto Ogbechie, Carlos Puerto-Santana, Concha Bielza, Industrial Applications of Machine Learning, 2018, Taylor and Francis, ISBN: 9781351128384
- Shai Shalev-Shwartz, Shai Ben-David: Understanding Machine Learning: From Theory to Algorithms, Cambridge University Press, 2014, ISBN: 9781107298019
- Kolla Bhanu Prakhas (ed), Machine Learning for Industrial Applications, Wiley, 2024, ISBN: 978-1-394-26897-9

EFFECTIVE AI: OPTI-MIZATION TECHNIQUES

INMEA9935-26

Semester:	3
Type:	Laboratory
Number of Classes:	0+0+2
Credit:	3
Status:	Optional
Assessment:	Practical mark
Prerequisites:	None
Responsible:	Dr. Róbert Lakatos

Topics:

The theoretical part introduces the principles of Parameter-Efficient Fine-Tuning (PEFT), with a focus on the operation of LoRA, which enables the adaptation of large models with limited resources. Students will learn about tools for reducing model size and computational requirements, such as various forms of quantization and pruning techniques. The course also reviews modern alignment methods such as supervised fine-tuning (SFT), Direct Preference Optimization (DPO), and Reinforcement Learning from Human Feedback (RLHF), which aim to ensure the safe and user-oriented operation of models. The theoretical part also deals with the characteristics of efficient LLM architectures and key metrics for quality assessment.

During the exercises, students prepare various data sets for fine-tuning processes, including the generation of synthetic data. They will implement PEFT methods, especially LoRA, in a real development environment and experience how a model's computational requirements can be reduced using quantization (e.g., 8-bit, 4-bit) and pruning techniques. Students will implement different alignment strategies and compare their effects on the behavior of the same model. During the semester, students also gain practical experience in applying decoding strategies (Top-k, Top-p, Temperature) and create a complex optimized LLM prototype as a project assignment.

Compulsory/Recommended Readings:

- Hu, E. J. et al.: LoRA: Low-Rank Adaptation of Large Language Models. ICLR, 2022., https://arxiv.org/abs/2106.09685
- Lewis, P. et al.: Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. NeurIPS, 2020., https://arxiv.org/abs/2005.11401
- François Chollet: Deep Learning with Python. Manning, 2017., ISBN: 9781617296864
- Ian Goodfellow, Yoshua Bengio, Aaron Courville: Deep Learning. MIT Press, 2016., ISBN: 9780262035613