Course: Intelligent systems *2026/27*

Lecturer: Prof Dr Marko Robnik-Šikonja
Assistants: Dr Tadej Škvorc, Assist Prof Dr Areeb Ahmed

The aim: The goal of the course is the students to become acquainted with the field of intelligent systems, which includes a collection of tools and approaches for solving problems which are difficult or unpractical to tackle with other methods.

Student's obligations:

- two assignments
- written exam

Grading

The practical work is graded through two equally weighted assignments, where each has to be finished on time and successfully (at least 50% of points). The exam is in written form. The preconditions for the written exam are successfully passed assignments. During the written exam, students are allowed to use one A4 double-sided sheet of paper. The oral exam is obligatory in case of any doubts about the score of the assignments, quizzes, or written exam. The final grade is the sum of the assignment scores (30%) and written exam (70%). The grades are valid in the current and next two academic years.

Syllabus

1. Introduction to intelligent systems and data science
2. Nature-inspired computing (genetic algorithms, genetic programming)
3. Introduction to predictive modeling
4. Bias, variance, and overfitting
5. Representation learning and feature engineering
6. Neural networks: architectures, backpropagation, deep neural networks
7. Ensemble methods
8. Kernel methods
9. Model inference and explanation
10. Automated machine learning
11. Foundation models for tabular data and time series
12. Natural language processing: text representation, transformer architecture, large language models, information extraction, text classification, semantic similarity
13. Reinforcement learning: basic approaches and algorithms, Q learning, TD learning, deep RL
14. Agentic systems
15. Evaluation of intelligent systems

Literature:

- James, G., Witten, D., Hastie, T., Tibshirani, R., and Taylor, J., 2023. *An Introduction to Statistical Learning: With Applications in Python*. New York: Springer. Freely available at https://www.statlearning.com/
- Chollet, F. and Watson, M., 2025. *Deep Learning with Python, 3rd edition*. Manning. Available at https://deeplearningwithpython.io/

Additional readings:

- D. Jurafsky, J. H. Martin. Speech *and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition, with Language Models, 3rd edition draft*, 2026, Freely available at http://web.stanford.edu/~jurafsky/slp3/
- Richard S. Sutton and Andrew G. Barto: *Reinforcement Learning, An Introduction, 2nd edition,* MIT press, 2018, free copy
- Kevin P. Murphy: *Probabilistic Machine Learning: An Introduction*. MIT Press, 2022, freely available from https://probml.github.io/pml-book/book1
- Kevin P. Murphy: Probabilistic Machine Learning: Advanced Topics. MIT Press, 2023, freely available at http://probml.github.io/book2
- Simon J.D. Prince: *Understanding Deep Learning. MIT Press, 2023* *https://udlbook.github.io/udlbook/* *Updated* in 2026, this book intuitively explains many important concepts and techniques in deep learning.
- Sebastian Raschka: *Build a Large Language Model (From Scratch)*. Manning, 2024
- Tong Xiao, and Zhu Jing-Bo: *Foundations of Large Language Models*. ArXiv:2501.09223, 2026