


Intelligent Systems

Prof Dr Marko Robnik-Šikonja

Edition 2025



Lecturer

- Prof Dr Marko Robnik-Šikonja
- University of Ljubljana
Faculty of Computer and Information Science
Head of Machine Learning and Language Technologies Lab
- FRI, Večna pot 113, 2nd floor, room 2.06, to the right from the elevator
- marko.robnik@fri.uni-lj.si
- <https://fri.uni-lj.si/en/employees/marko-robnik-sikonja>
- (01) 4798 241
- Contact hours (see the webpage)
 - currently, Wednesdays, 11:00 - 12:00; please, email me; other times or Zoom meeting are possible
- **Research interests:** artificial intelligence, machine learning, natural language processing, network analytics, data science
- **Teaching:** courses from areas of machine learning, natural language processing, and problem solving
- **Software and resources:** supporting open science, author of several open source ML packages, many large language models, and language resources

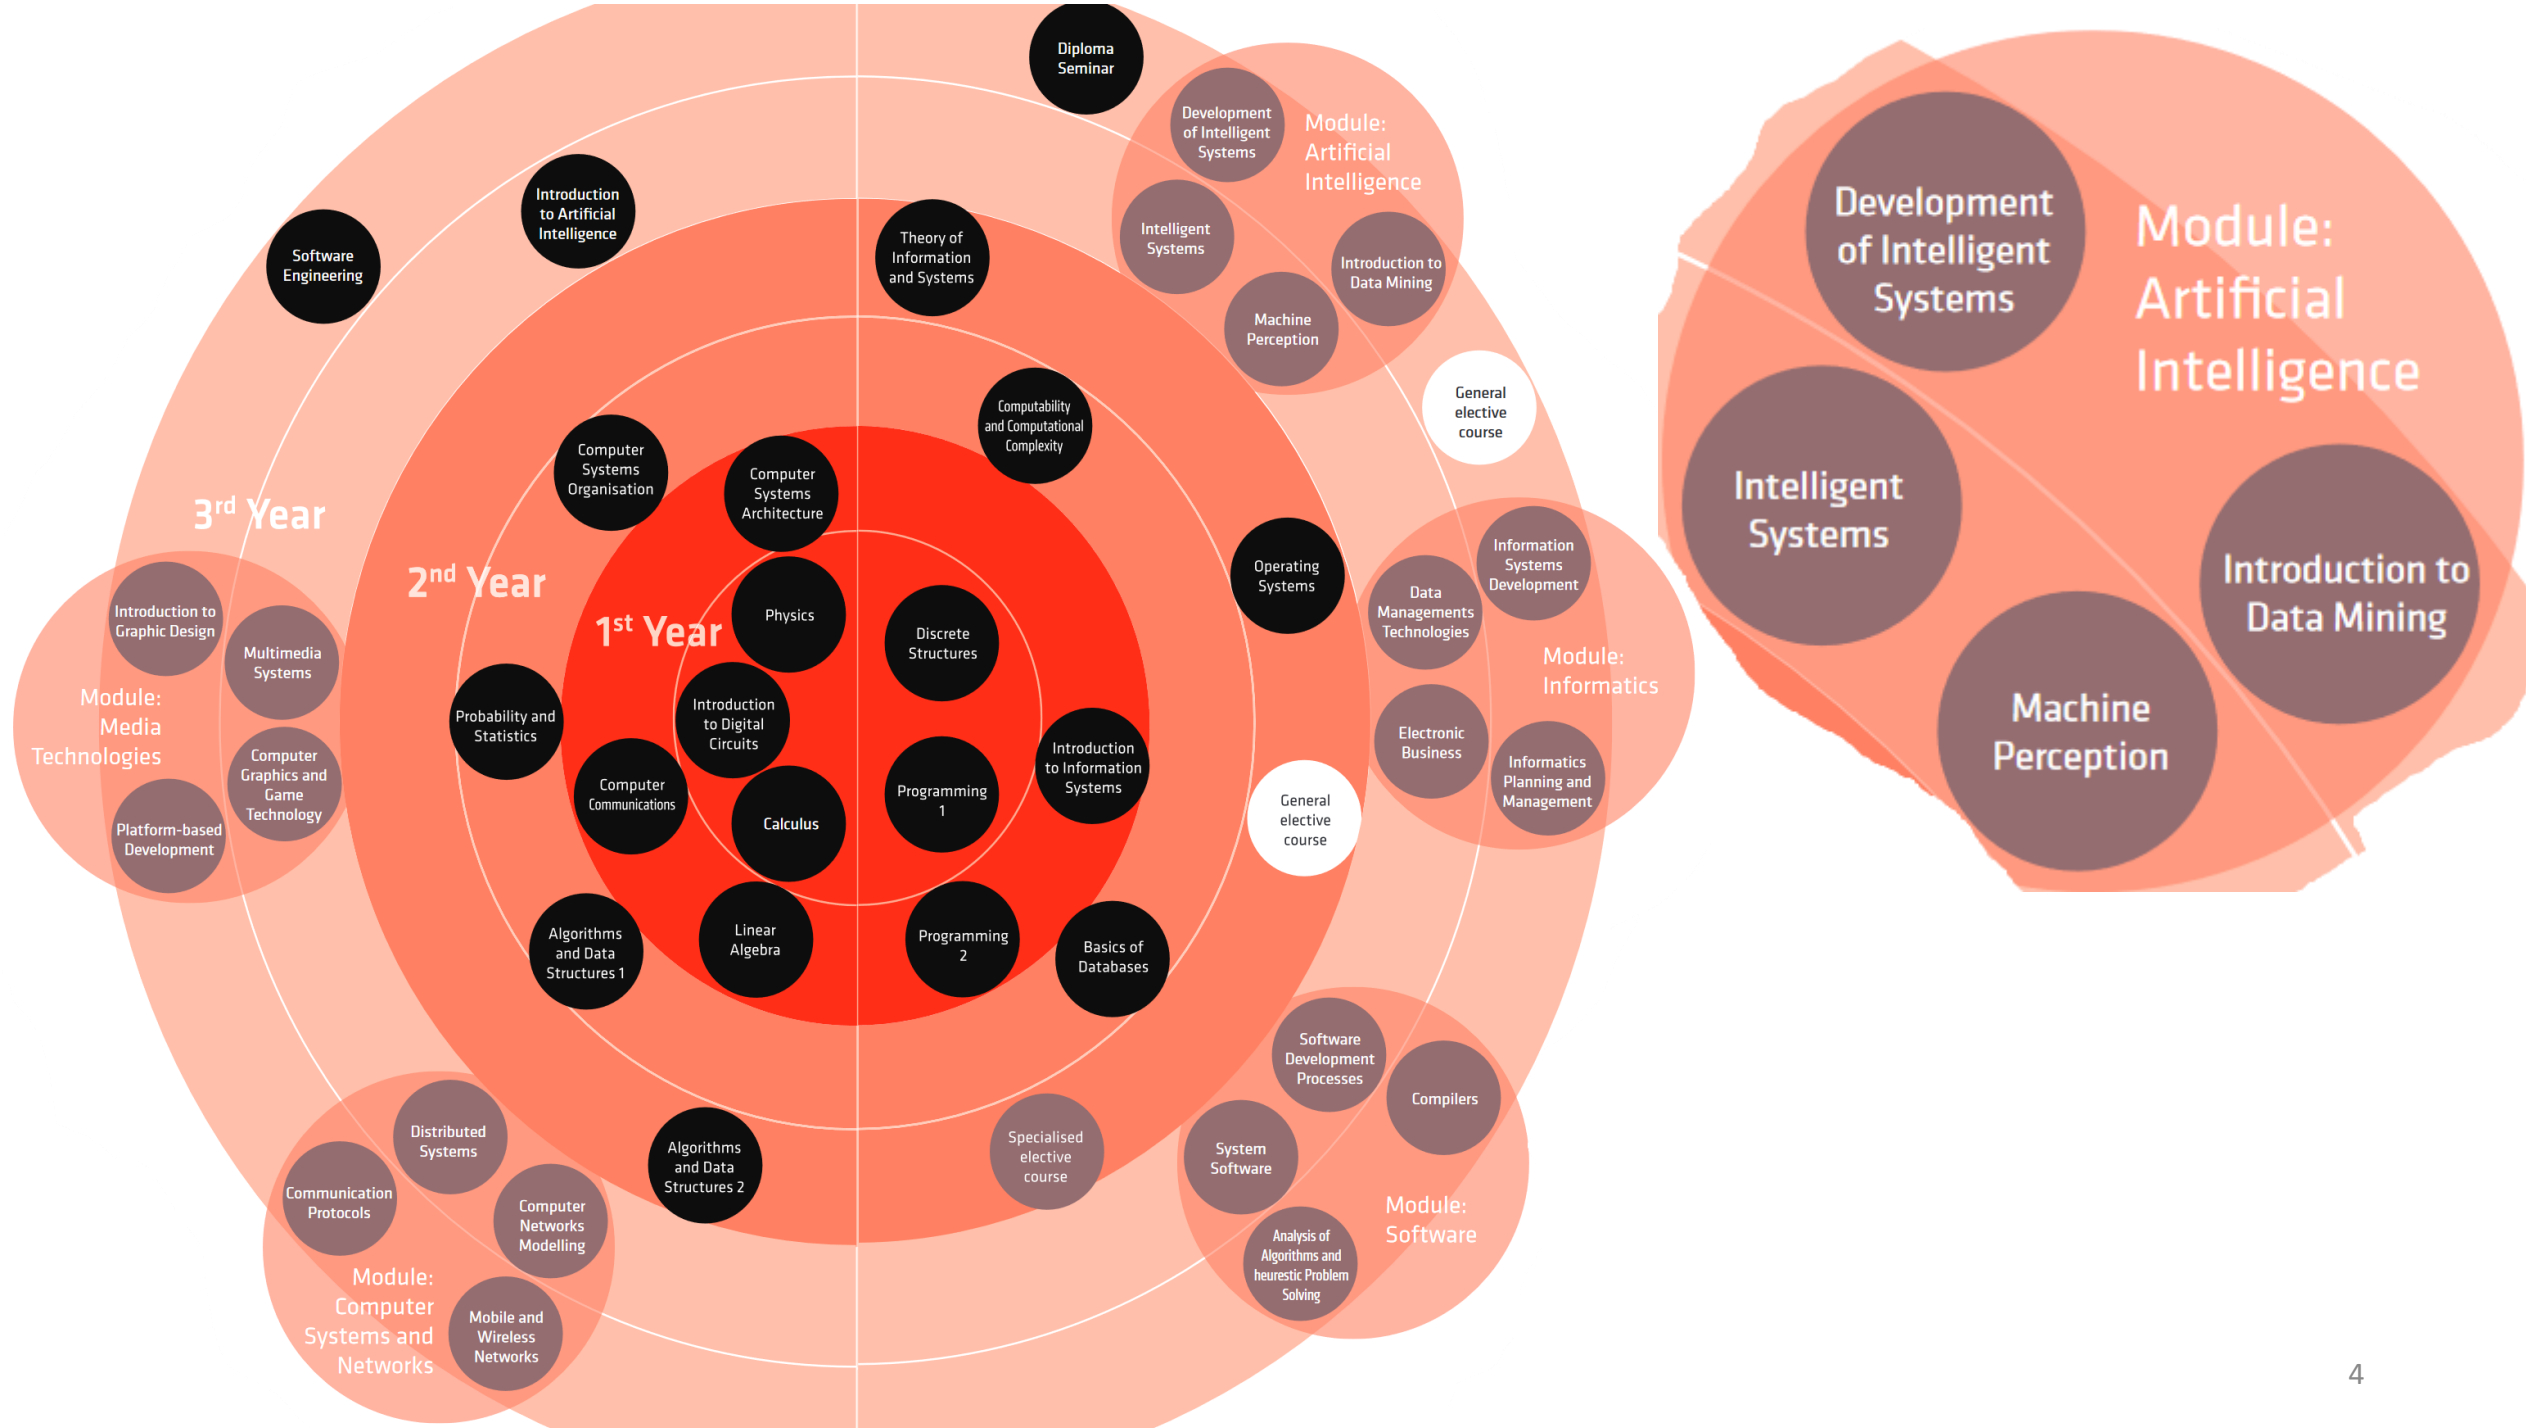


Assistants

- Dr Tadej Škvorc
- Timotej Knez, PhD student
- Boshko Koloski, PhD student



- tutorials, assignments, work in Python
- please, prepare questions!



Syllabus

- nature inspired computing (genetic algorithms, genetic programming)
- basics of machine learning,
- bias, variance, generalization error, and overfitting
- representation learning and feature selection
- neural networks
- natural language processing
- ensemble methods
- kernel methods
- automated machine learning
- transformers for tabular data and time series
- model inference and explanation
- reinforcement learning

Objectives

1. students shall become acquainted with

- nature inspired computing
- machine learning
 - predictive modeling approaches
 - model selection and evaluation techniques
 - model comprehensibility and explanation
 - practical application of predictive modeling in python
- natural language processing
- reinforcement learning

2. practical use of theoretical knowledge on (almost) real-world problems ; for a given prediction problem students shall be able to

- transform it to a form suitable for predictive modeling
- select and train an appropriate predictive model
- evaluate the model and present the results in a comprehensible form and language

3. awareness of domain expertise and ethical issues in data science

4. increase the (mental) problem-solving toolbox with

- predictive modeling techniques
- evolutionary optimization approaches
- large language models
- reinforcement learning
- experiment design, result understanding, visualization, and explanation approaches

Concrete course goals

Be able to explain

- difference between different types of machine learning models
- properties of models: bias, variance, generalization, hypothesis language
- properties of the following models: kNN, decision rules, bagging, boosting, random forests, stacking, SVM, neural networks
- properties and purpose of evaluation approaches and metrics: cross-validation, bootstrapping, ROC curves, sensitivity, specificity etc.
- When and how to apply AutoML techniques
- inference methods for predictive methods and explanation of predictions
- when and why to apply reinforcement learning
- how to prepare and process text
- when and how and to optimize a problem using evolutionary algorithms

Build and evaluate models in Python

- visualize datasets and created models
- prepare data into a suitable form for modeling algorithms
- apply classification and regression models to solve a prediction task with a given data set
- build natural language classifier
- estimate error of models using statistically valid approaches
- select models and tune their parameters using cross-validation and bootstrapping
- apply AutoML techniques
- visualize models and explain their predictions
- given a new dataset, select an appropriate modeling technique and evaluate the created model

Syllabus explained

Nature inspired computing

- genetic algorithms
- genetic programming
- neuro-evolution

Introduction to statistical predictive modelling

- Learning as modelling: data, evidence, background knowledge, predictive models, hypotheses, learning as optimization, learning as search, criteria of success, inductive learning, generalization.
- Classification and regression: supervised and unsupervised learning, learning discrete and numeric functions, learning relations, learning associations.
- Simple classification models: nearest neighbor, decision rules

Model selection

- Bias and variance: error decomposition, trade-off, estimating bias and variance
- Generalization performance: training and testing set error, cross-validation, evaluation set, bootstrapping.
- Performance measures: confusion matrix, sensitivity and specificity, ROC curves, AUC, cost-based classification.
- Parameter tuning: regularization, search
- AutoML

Kernel methods

- SVM for classification and regression: kernels, support vectors, hyperplanes.
- SVM for more than two classes: one vs. one, one vs. all.

Ensemble methods

- Model averaging, why ensembles work.
- Tree based ensembles: bagging, boosting, random forests.
- MARS and AODE ensembles.
- Stacking, mixture of experts.

Neural networks

- perceptron,
- backpropagation,
- setting structure of networks
- deep neural networks
- transformer architecture
- autoencoders
- GANs
- neural embeddings and representation learning

Explaining prediction models

- Model comprehensibility, visualization and knowledge discovery.
- General methodology for explaining predictive models.
- Model level and instance level explanations, methods SHAP, LIME, EXPLAIN, and IME.

Learning with special settings

- imbalanced data,
- multi-task learning,
- multi-label learning,
- Etc.

Natural language processing

- text preprocesing
- text representation
- text similarity
- text classification
- sentiment analysis
- generative models

Reinforcement learning

- basics
- Markov decision problem
- Q learning
- Deep RL

Course organisation

Obligations

- 5 quizzes
- Two projects, 50 points
- Written exam, 50 points

Grading

Obligation	% of total	subject to
Five quizzes	0%	$\geq 50\%$ altogether
Projects	50%	$\geq 50\%$ each
Written exam	50%	$\geq 50\%$

Learning materials

- [Learning materials in Moodle](#)
- Slides
- Quizzes
- Links and other materials


Readings (all freely available)

- James, G., Witten, D., Hastie, T., Tibshirani, R. and Taylor, J., 2023. [An Introduction to Statistical Learning: With Applications in Python](#). New York: Springer. (also exists for R)
- Chollet, F. and Watson, M., 2025. [Deep Learning with Python](#), 3rd edition. Manning.

Further readings:

- Jurafsky, Daniel and James, Martin (2025): [Speech and Language Processing](#), 3rd edition in progress
- Richard S. Sutton and Andrew G. Barto: [Reinforcement Learning, An Introduction](#), 2nd edition, MIT Press, 2018
- Kevin P. Murphy: [Probabilistic Machine Learning: An Introduction](#). MIT Press, 2022
- Kevin P. Murphy: [Probabilistic Machine Learning: Advanced Topics](#). MIT Press, 2023
- Friedman, J., Hastie, T., & Tibshirani, R., 2009). [The elements of statistical learning](#), 2nd edition. Springer, Berlin
- scientific papers
- many excellent machine learning and data science courses on Coursera, edX etc.


Retention of learning



Retention of Learning

Data Science is a part of Intelligent Systems

- good job perspective
- many jobs in this area regularly occupy list of the most promising jobs
- Thomas H. Davenport, D.J. Patil: Data Scientist: The Sexiest Job of the 21st Century. *Harvard Business Review*, October 2012



MODERN DATA SCIENTIST

Data Scientist, the sexiest job of the 21st century, requires a mixture of multidisciplinary skills ranging from an intersection of mathematics, statistics, computer science, communication and business. Finding a data scientist is hard. Finding people who understand who a data scientist is, is equally hard. So here is a little cheat sheet on who the modern data scientist really is.

MATH & STATISTICS

- ☆ Machine learning
- ☆ Statistical modeling
- ☆ Experiment design
- ☆ Bayesian inference
- ☆ Supervised learning: decision trees, random forests, logistic regression
- ☆ Unsupervised learning: clustering, dimensionality reduction
- ☆ Optimization: gradient descent and variants

PROGRAMMING & DATABASE

- ☆ Computer science fundamentals
- ☆ Scripting language e.g. Python
- ☆ Statistical computing packages, e.g., R
- ☆ Databases: SQL and NoSQL
- ☆ Relational algebra
- ☆ Parallel databases and parallel query processing
- ☆ MapReduce concepts
- ☆ Hadoop and Hive/Pig
- ☆ Custom reducers
- ☆ Experience with xaaS like AWS

DOMAIN KNOWLEDGE & SOFT SKILLS

- ☆ Passionate about the business
- ☆ Curious about data
- ☆ Influence without authority
- ☆ Hacker mindset
- ☆ Problem solver
- ☆ Strategic, proactive, creative, innovative and collaborative


COMMUNICATION & VISUALIZATION

- ☆ Able to engage with senior management
- ☆ Story telling skills
- ☆ Translate data-driven insights into decisions and actions
- ☆ Visual art design
- ☆ R packages like ggplot or lattice
- ☆ Knowledge of any of visualization tools e.g. Flare, D3.js, Tableau



DATA SCIENTIST

What my **CUSTOMERS** think I do



What my **MUM** thinks I do



What my **FRIENDS** think I do



I work with
Brad
Pitt!!!

What my **HUSBAND** thinks I do



What **I** think I do



What I **ACTUALLY** do



Intelligent systems and media



Will robots destroy us?

Will they take our jobs?

Will we still need a driving licence?

Will we still need doctors?

How will humanoid robots evolve?

What about cyborgs?

What is artificial general intelligence?

What is technological singularity?

New prophets of technological singularity

Elon Musk says humans must become cyborgs to stay relevant. Is he right?

Sophisticated artificial intelligence will make 'house cats' of humans, claims the entrepreneur, but his grand vision for mind-controlled tech may be a long way off



Some scientific opinions

- Rodney Brooks: The Seven Deadly Sins of Predicting the Future of AI. <https://rodneybrooks.com/the-seven-deadly-sins-of-predicting-the-future-of-ai/> also in MIT Technology Review
- Marko Robnik-Šikonja: Is artificial intelligence a (job) killer?. The Conversation, Jul. 2017 <https://theconversation.com/is-artificial-intelligence-a-job-killer-80473>
- ...



Short history of optimism

- starting in 1950s,
1956 Dartmouth conference
- great expectations, enormous underestimation
of problem difficulty
- AI winter (2 x)




1958, H. A. Simon and Allen Newell: "... within ten years a digital computer will discover and prove an important new mathematical theorem."

1965, H. A. Simon: "... machines will be capable, within twenty years, of doing any work a man can do."


1967, Marvin Minsky: "Within a generation ... the problem of creating 'artificial intelligence' will substantially be solved."

1970, Marvin Minsky: "In from three to eight years we will have a machine with the general intelligence of an average human being."

Hype Cycle for Artificial intelligence, 2025



Hype Cycle for Artificial Intelligence, 2024



Hype Cycle for Emerging Technologies, 2024

