University of Ljubljana, Faculty of Computer and Information Science
Doctoral study programme Computer and Information Science

Elective courses BDR-RI 2024/2025

All courses are 5 ECTS. There are two types of courses available.

The **lecture** type courses are delivered as regular lectures and follow the format 15-20-15 (lectures-seminar-tutorial hours).

The **individual research** type courses introduce advanced technological breakthroughs or practical solutions in computer and information science. Students work under the lecturer's supervision on a seminar topic related to the student's doctoral research topic. Each course can be selected by at most six students. The course lecturer can be the advisor or co-advisor of the student selecting the course. Each student can take at most three individual research courses.

Lecture type courses offered in 2024/2025:

Machine Learning and Artificial Intelligence 2
ChatGPT for Researchers 2
Effective Theory of Deep Learning........ 3
Physics-Informed Machine Learning........ 3
Advanced Algorithms for Search and Planning........ 4
Large Language Models: Machine Learning View........ 4
Predictive Analytics for Structured Data........ 5
INFOSEC of Socio-Technical Systems 5

Individual research type courses offered in 2024/25:

Advanced Topics in AI for Medicine (IR) 6
Recent Advances in Combinatorial Solvers (IR) 6
Selected Topics in Analysis of Sound Signals (IR) 7
Selected Topics from Computer Graphics and Visualization (IR) 7
Deep Learning for Computer Vision (IR)........ 8
Heterogeneous Computing Platforms (IR) 8
In-Depth Computer Vision Research (IR)........ 9
Advanced Image-based Biometrics (IR)........ 10
Machine Learning for Remote Sensing (IR) 10

Machine Learning and Artificial Intelligence

(Selected Topics in Artificial Intelligence 1, Zoran Bosnić)
Lecturers: Blaž Zupan and Janez Demšar
Course code: 63834E
Course type: lectures, fall (first) semester

This course is an introduction to data science for non-computer scientists. The course covers topics from data preparation, clustering, regression and classification, model evaluation, and embedding of unstructured data.

Restrictions/Prerequisites: The course will be promoted by the University's Doctoral School, and we expect enrollment from students of engineering, natural sciences, and humanities. ***The course is not intended for computer science students or students whose curricula already included courses on machine learning or data science.*** *No prior knowledge on the topics is assumed. This course will not use computer programming and no prior knowledge on statistics or data science is required.*

ChatGPT for Researchers

(Selected Topics in Artificial Intelligence 2, Zoran Bosnić)
Lecturer: Blaž Zupan
Course code: 63835F
Course type: lectures, fall (first) semester

This course is an introduction to ChatGPT and similar large language models. It will cover an introduction with intuitive explanation of what are large language models. We will continue with use cases of ChatGPT's web-based interface, focusing on how it can assist researchers in various tasks, including providing instant access to a vast range of information, facilitating brainstorming, generating ideas, and summarizing complex concepts. It can also assist in reviewing and editing research documents, proposing research questions, and helping researchers understand complex methodologies and techniques in various disciplines. We will discuss the deficiencies of the technology, including the provision of inaccurate or outdated information and lack of understanding or context awareness, reflecting limitations in its training data and the absence of real-world experience or subjective perception. The course will showcase the use of ChatGPT's application programming interface (API) and its advanced uses, including AutoGPT.

Restrictions/Prerequisites: The course will be promoted by the University's Doctoral School, and we expect enrollment from students of engineering, natural sciences, and humanities. ***The course is not intended for computer science students or students whose curricula already included courses on machine learning or data science.*** *No prior knowledge on the topics of large language models or computer programming is assumed. This is an introductory course intended for general audience. Students from humanities, social sciences, natural sciences and engineering are welcome.*

Effective Theory of Deep Learning

(Selected Topics in Artificial Intelligence 2, Zoran Bosnić)
Lecturer: Bojan Žunkovič
Course code: 63835I
Course type: lectures, fall (first) semester

Deep learning architectures have many hyperparameters. Considering only the paradigmatic multi-layer perceptron (MLP) architecture, we can change the number of layers, their width, activation functions, and initial parameter distributions. In the last two decades, these hyperparameters have been experimentally optimized towards their optimal values. In this course, we will build a rigorous effective theory that quantitatively describes various choices of hyperparameters, namely the height-to-width ratio, initialization distribution, and activation functions. Our theory will use Gaussian perturbation theory to describe the correlations of activations through various layers of the MLP network. Our rigorous results will provide a firm theoretical ground for many well-established deep-learning practices. Considered techniques can be extended to various other architectures, e.g., transformers. Importantly, they can also be used to derive the scaling laws for large models, which is increasingly important due to the huge cost of their training.

Restrictions/Prerequisites: Linear algebra, probability theory, multidimensional Gaussian integrals, perturbation theory, Taylor expansion.

Physics-Informed Machine Learning

(Selected Topics in Artificial Intelligence 2, Zoran Bosnić)
Lecturer: Bojan Žunkovič
Course code: 63835J
Course type: lectures, fall (first) semester

We will discuss several modern physics-informed approaches to data-driven dynamical systems and architecture design. The key principles that will guide our exploration will be parsimony and symmetry. Based on those principles, we will discover correct dynamical descriptions of complicated systems only from (sparse and noisy) data and modern deep-learning architectures only from several basic ingredients: symmetry, stability and scale separation. Presented methods are an integral part of modern data-driven physics discovery and state-of-the-art deep-learning system design. We will provide Jupyter notebooks (or Colab) with a demonstration code on simple examples.

Restrictions/Prerequisites: Basics of deep neural networks, group and representational theory, linear algebra.

Advanced Algorithms for Search and Planning

(Selected Topics in Artificial Intelligence 1, Zoran Bosnić)
Lecturer: Ivan Bratko
Course code: 63834A
Course type: lectures, fall (first) semester

The course covers advanced algorithms for the following problem-solving paradigms: state-space search, symbolic search, planning based on means-ends analysis (mostly approaches to partial order planning) and planning under uncertainty. The course assumes introductory knowledge of search and planning algorithms in AI (uninformed search, bestfirst search, and planning principles based on preconditions and effects of actions). The course includes advanced versions and extensions of search and planning algorithms, their theoretical properties, implementations and practical projects in combinatorial optimisation and robot programming. The course includes home assignments in experimentation with existing or own implementations of heuristic search on practical problems (with own ideas for heuristics), and a project in robot planning (or a planning problem of a student's choice, possibly related to the topics of their doctoral research).

Restrictions/Prerequisites: Basics of search and planning techniques.

Large Language Models: Machine Learning View

(Selected Topics in Artificial Intelligence 1, Zoran Bosnić)
Lecturer: Marko Robnik Šikonja
Course code: 63834D
Course type: lectures, fall (first) semester

The course presents large language models (LLMs) from the point of view of machine learning, covering the main ideas used in their creation, adaptation, and application. In this context, it presents advanced machine learning topics used in representation learning, including topics from ensemble learning, explainable artificial intelligence, transfer learning, and graph learning. The contents commonly apply in natural language processing, natural language understanding, knowledge graphs, multi-relational learning, digital humanities, bioinformatics, etc. The specific technical contents of the course cover neural network architectures for large language models, relevant large language model adaptions and extensions, retrieval augmented generation, and multimodal models. The course requires students to apply large language models, typically to language and graph processing tasks, preferably in the context of their research work.

Restrictions/Prerequisites: Prior knowledge of machine learning, statistics and programming in Python.

Predictive Analytics for Structured Data

(Selected Topics in Artificial Intelligence 2, Zoran Bosnić)
Lecturer: Sašo Džeroski
Course code: 63835B
Course type: lectures, spring (second) semester
The course will introduce different tasks of structured output prediction and describe various approaches for solving such tasks. The students will get to know some state-of-the-art tools for solving such tasks and examples of their use in practice. Within the course, the students will learn to apply predictive analytics methods for structured data in the context of their research. In this course, we will study the different tasks of structured output prediction, such as multi-target classification/regression and (hierarchical) multi-label classification, predictive clustering methods (tree and rule-based) for structured output prediction, ontologies for data mining and their use for describing structured output prediction, ensemble methods for structured output prediction (tree and rule ensembles), applications of structured output prediction to different practical problems from areas such as environmental and life sciences, and image annotation and retrieval. In the practical hands-on work, students will be guided through a series of methods for predicting structured outputs. They will analyze relevant data sets (from ecology and systems biology) that represent different tasks of predicting structured outputs, *e.g.*, multi-target regression, multi-label classification, hierarchical multi-label classification. In the last part of the course, each student will apply and test methods for predicting structured outputs on a selected relevant doctoral research problem.

Restrictions/Prerequisites: Basic knowledge of machine learning, data mining, statistics and programming skills.

INFOSEC of Socio-Technical Systems

(Selected Topics in Computer and Information Science, Tomaž Curk)
Lecturer: David Modic
Course code: 63818A
Course type: lectures, spring (second) semester

In this course, we will briefly discuss the foundations of informational security (INFOSEC), define socio-technical systems, discover ethics and laws as they pertain to cybersec kill-chain, and penetration testing activities. Our focus will be on the so-called soft INFOSEC skills. Some of the topics that will be covered include: Open Source Intelligence Gathering; Social engineering; phishing; intelligence gathering toolkits; and the psychology of effective manipulation, and security in general. An operational understanding of the Internet and familiarity with computers is a bonus. There will be a certain amount of practical work required and active participation in the course will be mandatory. The course is open to PhD students from any infosec related field (upon confirmation). Previous years have thought us that there is also strong interest from masters and undergraduate students, who are welcome to audit the course upon confirmation from the module convenor.

Restrictions/Prerequisites: The course is intended for doctoral students in computer science or related fields.

Advanced Topics in AI for Medicine (IR)

(Selected Topics in Artificial Intelligence 1, Zoran Bosnić)
Lecturer: Zoran Bosnić
Course code: 63834F
Course type: individual research course, fall (first) semester

This course will offer an exploration of the integration of artificial intelligence (AI) and machine learning (ML) methodologies in healthcare and medicine. Through individual research work, students will have an opportunity to address topics, examining the use of AI in medical diagnosis, treatment, and patient management. Specifically, students will investigate the application of deep learning techniques in analyzing medical imaging data, the development of predictive models for personalized treatment strategies, and the utilization of natural language processing algorithms for extracting valuable insights from clinical text repositories. The examples of such AI applications include existing applications in cardiology, focusing on the development of predictive models for early detection of cardiac diseases and risk stratification of patients; and applications in the intersection of AI and neurodegenerative diseases, particularly Parkinson's and Alzheimer's, investigating the potential of AI-driven approaches in early detection and disease progression monitoring through analysis of biomarkers, neuroimaging data, and patient records. The goal of the course would be to enable PhD students develop their research skills in an actual and particular research areas, while also meaningfully contributing to the ongoing advancement of this interdisciplinary field.

Restrictions/Prerequisites: A good foundation in machine learning and programming. Beneficial is a basic understanding of concepts and healthcare systems. Maximum number of students: 6.

Recent Advances in Combinatorial Solvers (IR)

(Selected Topics in Architectures and Algorithms 2, Borut Robič)
Lecturer: Uroš Čibej
Course code: 63825
Course type: individual research course, fall (first) semester

The last decade has witnessed a small revolution in the field of solvers for hard combinatorial problems. The guiding problem has been SAT, which has pushed to boundaries of solvable problems for a few orders of magnitudes and modern SAT solvers can achieve truly remarkable performance. However, solvers for other problems have seen big breakthroughs, e.g., solvers for the clique problem, subgraph isomorphism problem, vertex cover, and graph isomorphism problem. In this course, the students will investigate and empirically evaluate these advances on real-life problems.

Restrictions/Prerequisites: Basic knowledge of graph theory, combinatorial algorithms, computational complexity.

Selected Topics in Analysis of Sound Signals (IR)

(Selected Topics in Software Development 1, Matija Marolt)
Lecturer: Matija Marolt
Course code: 63832
Course type: individual research course, fall (first) semester

Students will have the opportunity to explore the use of different methods for pattern recognition and machining learning (for example, deep neural networks) to solve the problems that we encounter when analyzing sound signals, such as identification of events in sound recordings, classification of sound recordings, transcription of music, detection of samples in music, etc. In the course, students will develop their own algorithm for solving a problem and send it to one of the evaluation campaigns (e.g., Mirex or DCASE), where its performance can be compared with approaches developed by other researchers (mostly doctoral students) worldwide.

Restrictions/Prerequisites: /

Selected Topics from Computer Graphics and Visualization (IR)

(Selected Topics in Software Development 1, Matija Marolt)
Lecturer: Ciril Bohak
Course code: 63832B
Course type: individual research course, fall (first) semester

Students will learn the current methods and technologies in the field of three-dimensional computer graphics. Emphasis will be given to rendering different types of data: volumetric data, point clouds, mesh geometry, and logically defined geometry in the fields of medicine, biology, geodesy, and high-energy physics. Because the rendered data can be very large, emphasis will also be given to applying appropriate algorithms and data structures for fast and real-time rendering, implementation of techniques on graphic processors, and remote rendering. The students will get to know the benefits of modern graphics libraries (Vulkan, WebGPU) for addressing these challenges. In addition to the techniques, the students will also get acquainted with the different ways of visualizing such data, how to utilize various deep learning tools on the data for visualization preparation or visual parameter estimation, and how to select suitable visualization methods for an individual domain. Students will have an opportunity to collaborate and interact with other students and staff from one of the world's best Visual computing groups in the world at KAUST.

Restrictions/Prerequisites: /

Deep Learning for Computer Vision (IR)

(Selected Topics in Artificial Intelligence 2, Zoran Bosnić)
Lecturer: Danijel Skočaj
Course code: 63835
Course type: lectures, spring (second) semester

The research field of computer vision addresses the problems related to acquiring, processing, analyzing, and understanding visual information such as images, videos, and 3D point clouds. One of the core problems in computer vision is visual learning and recognition, i.e., learning the representations (of objects, faces, rooms, actions, etc.) that are later used to classify unknown instances that appear in new images. This problem has been tackled since the beginning of the computer vision. However, no previously proposed method has increased the performance beyond the current state of the art like deep learning approaches in recent years. Convolutional neural networks and related deep learning approaches have proven to be a very efficient way of finding the representations and building a classifier in a unified framework that yields excellent results in various computer vision tasks. The main goal of this course is to introduce students to the field of deep learning, with a special emphasis on its application in computer vision. The students will be acquainted with the main principles of computer vision and machine learning, relating them to neural network methods and showing them how to train and use neural networks, emphasizing Convolutional Neural Networks. It will be shown how these approaches can be used for object classification, localization, and detection, as well as for other tasks in computer vision and beyond.

Restrictions/Prerequisites: Solid knowledge of computer vision and machine learning, programming skills.

Heterogeneous Computing Platforms (IR)

(Selected Topics in Computer Systems 1, Miha Mraz)
Lecturer: Uroš Lotrič
Course code: 63830A
Course type: individual research course, spring (second) semester

In this course, we will deal with state-of-the-art platforms and technologies that present a prevailing direction in ensuring enough computing performance for increasing computational requirements. Students will work with different types of hardware accelerators like GPU, FPGA, multicore CPU, and their combinations, preferably integrated into the high-performance computing system. For a selected application related to their doctoral thesis, they will have to recognize a suitable platform and then implement and evaluate the chosen application on it. In this course, we will study the speed-up of complex algorithms on modern hardware architectures, how to combine CPU and custom FPGA circuits to excel computation, and how to analyze the effect of number representation (fixed point, half-precision) to reduce computational cost and save energy.

Restrictions/Prerequisites: Background in computer science, electronics, mathematics, or physics; programming skills, favorable in concurrent programming.

In-Depth Computer Vision Research (IR)

(Selected Topics in Artificial Intelligence 2, Zoran Bosnić)
Lecturer: Matej Kristan
Course code: 63835G
Course type: lectures, spring (second) semester

The course will focus on a selected topic in computer vision that connects to the candidate's doctoral thesis. The main aim of the course is to expand the research with intensive training on how to tease out the most relevant related works in computer vision, analytically or experimentally discover their drawbacks, make original contributions, and validate them.

Computer vision involves the development of algorithms that can abstract complex unstructured data, such as images and videos, in a broad sense. This encompasses tasks like semantic segmentation, visual object tracking, object detection, extraction of 3D information, visual data generation, and their application to various downstream tasks like image/video manipulation and mobile robotics. The student will begin by selecting a suitable topic within their broader tentative doctoral research area, with the guidance of their supervisor. A specific goal, achievable within a single semester (approximately 150 hours of work), will be defined. The student will then establish a research plan and provide progress reports through bi-weekly meetings.

The initial reports will involve breaking down and critically analyzing existing works closely related to the chosen topic from a methodological perspective. A series of test-and-hypothesis-generation sessions will be conducted, challenging the student to identify a localized big-picture idea that can lead to a scientific contribution in their selected area. This intensive training aims to equip the student with a robust methodology for approaching scientific discovery within computer vision and encourage them to generate their contributions. The ultimate objective is a high-caliber scientific contribution meeting publication quality standards necessary at a prominent computer vision venue.

Moreover, the student will develop the ability to critically review relevant papers on arXiv as part of their daily routines (expected to spend an hour a day), ensuring they stay updated with daily scientific advancements. A secondary outcome of the course is to foster an understanding of the characteristics that distinguish top publications. The student will gain the knowledge and skills to reproduce the style and rigor required for high-quality scientific reporting in computer vision.

Restrictions/Prerequisites: The course is primarily aimed at doctoral candidates who have started their doctoral training under the mentorship of the course lecturer. ***The students' doctoral topic must be from the core computer vision topic. The appropriateness of the topic will be judged by the lecturer, and the candidate is invited to consult the lecturer before applying for the course.***

Advanced Image-based Biometrics (IR)

(Selected Topics in Artificial Intelligence 2, Zoran Bosnić)
Lecturer: Peter Peer
Course code: 63835H
Course type: lectures, spring (second) semester

With the advances of computer vision with deep learning biometrics gained a lot of extra traction. This is evident from the papers published in the top-notch conferences like CVPR, ICCV, ECCV, FG, IJCB, journals like PAMI, TIFS, TIP, IVC, KBS, and also in a number of competitions organised each year. The research conducted in the field of biometrics improves security and privacy issues in society - two of very relevant questions that even the general public is very interested in.

Based on the student doctoral topic, we will link it to selected biometrical modality to address open research questions, for instance, in: segmentation and recognition models, bias assessment, generation of synthetic datasets, privacy-preserving biometrics, deidentification, input quality assessment and enhancement, deepfake detection.

Restrictions/Prerequisites: Familiarity with basic computer vision and deep learning approaches. Maximum number of students: 3.

Machine Learning for Remote Sensing (IR)

(Selected Topics in Artificial Intelligence 2, Zoran Bosnić)
Lecturer: Luka Čehovin Zajc
Course code: 63835K
Course type: lectures, spring (second) semester

The course will combine two very complementary topics. On the one hand, we have machine learning, a methodological framework for building prediction models based on statistics of data. On the other hand, we have remote sensing, a prime example of which is Earth observation using a growing number of satellites. These satellites produce enormous quantities of sensory data that can only be analyzed using computers. Through individual research projects, the course will explore the opportunities and challenges of using machine learning methods on remote sensing data.

The course is primarily intended for students who already work with machine learning (e.g., computer vision, language) in their primary research but would like to extend the applicability of their methodological skills and for students (from other faculties) who are primarily remote sensing practitioners who would like to improve their machine learning skills and deepen their understanding of ML methodology and best practices.

Some example research topics that could be explored: Temporal data for event detection or vegetation classification. Counting objects of interest in large areas. Self-supervised learning of meaningful features on multi-spectral data. Handling noisy training data by incorporating context. Large-language models and their relation to remote sensing data. Change detection as detection of anomalies. Satellite-based visual localization of drones.

Restrictions/Prerequisites: Basic machine learning and deep learning.