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

Elective courses BDR-RI 2023/2024

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 2023/2024:

Individual research type courses offered in 2023/2024:

Machine Learning and Artificial Intelligence

(Selected Topics in Artificial Intelligence 1, Zoran Bosnić)
Lecturer: 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.*

Incremental Learning from Data Streams

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

The course aims to teach students about state-of-the-art algorithms used to perform learning from data streams. The course will guide the students through the major open challenges in the field (supervised learning, data compression, concept drift detection, clustering from streams, specialized evaluation statistics). With this knowledge, the students will be able to apply their machine learning skills to a specialized and useful area connected to the abundance of data in our everyday lives (bank/weather/financial transactions, sensor readings, *etc.*). The course will be organized by mixing lectures with hands-on lab exercises in the Statistical package R. The students will apply the acquired knowledge to their problems. The course will stimulate competition between students to achieve the best possible learning results.

Restrictions/Prerequisites: Basic concepts of supervised and unsupervised learning in machine learning/data mining. Basic mathematical and statistical knowledge is also welcome for easier understanding of data sampling and descriptor techniques. It is mandatory that the students have programmed before in any procedural or object-oriented language as the practicals and tasks will be held using Statistical package R. *Having* R *installed at their own computers (notebooks) is desired.*

Tensor Networks for Machine Learning

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

Tensor networks are decompositions of multi-dimensional tensors with exponential reduction of parameters. They have been introduced to quantum mechanics approximately 20 ago. Since then, they have become one of the most important technical tools to understand quantum states' structure, especially in one dimension. They are a vital ingredient of the state-of-the-art numerical techniques of many-body quantum mechanics. In many-body quantum mechanics and quantum information, tensor networks are now a well-established and understood tool with well-known geometric properties and robust optimization algorithms. In the last seven years, they also appeared in mathematical literature (particularly matrix product states or tensor trains) in linear algebra with large matrices. Over the previous four years, they increasingly appear in machine learning literature, where they have been applied to various practical problems from parameter compression, classification to anomaly detection. Theoretically, tensor networks have been related to Born machines, hidden Markov models, and probabilistic and quadratic automata from the formal language's literature. The course will focus on recently developed tensor network applications for machine learning (mainly from an experimental/numerical perspective).

We will guide students in reproducing recent results involving tensor network decompositions in machine learning and then trying to go beyond by improving or applying the learned techniques to a slightly different problem. The proposed projects will be adapted to fit student interests, time, and expertise.

Restrictions/Prerequisites: Basic linear algebra, basic machine learning.

Advanced Topics in Ubiquitous Sensing and Learning

(Selected Topics in Computer Systems 2, Miha Mraz)
Lecturer: Veljko Pejović
Course code: 63831
Course type: lectures, spring (second) semester

The course covers theoretical, system, and application aspects of mobile, wearable, and Internet of Things devices (here referred to as "ubiquitous") for sensing and learning about the environment. The course starts with the overview of ubiquitous sensing platforms, covering topics such as the constraints and applications and the functioning of these platforms, thus touching upon the sampling theory (including the recent advances in sub-Nyquist sampling) and the "sampling - feature extraction - machine learning" pipeline. The course then thoroughly examines recent innovations that finally brought deep learning (DL) to a range of ubicomp devices before continuing with an in-depth investigation of applications of DL on this platform, *e.g.*, for human activity recognition, healthcare, authentication and security, and wireless inference. Deep learning is also the focus of the collective intelligence brought by distributed IoT deployments. The course covers distributed DL training via federated and split learning and solutions for local-cloud learning distribution. A key course component is a practical project that students will independently work on. The project harnesses modern tools for mobile sensing (*e.g.*, Android) and on-device deep learning (*e.g.*, TensorFlow Lite) and requires students to develop a full-fledged mobile deep learning application. Student participation is facilitated further by mandatory research paper presentations delivered by each student in the class. Finally, the course's relevance for preparing students to address global challenges is exacerbated through a dedicated lecture on ubicomp technologies for tackling COVID-19 pandemics and a guest lecture on mobile solutions for developing regions.

Restrictions/Prerequisites: Fundamentals of machine learning, Basics of deep learning, Object-oriented programming. Each student must have a laptop capable of running Android Studio. Each student should have an Android smartphone running Android API 26 or above for project demo purposes (two such phones are available in our lab). Additional equipment in the form of Android development boards (e.g. HiKey) or external sensors may be needed for certain projects. This is likely available in our lab or at FRIzider.

Selected Topics in Software Engineering

(Selected Topics in Software Development 2, Matija Marolt)
Lecturer: Vlado Stankovski and Petar Kochovski
Course code: 63833
Course type: individual research course, spring (second) semester

In the duration of the course, selected chapters will be presented and analyzed in the domains of the Internet of Things, Artificial Intelligence, Cloud-to-Things computing, semantic technologies, Digital Twins, and blockchain technologies with the goal of developing novel human-centric smart and trustworthy services and applications. The course will cover the following domains:
1. main concepts, vision, and goals of cloud, fog, and edge computing with respect to security, trust, and service reliability, 2. cognitive and context-aware computing and solving bias in artificial intelligence, 3. decentralized democratic governance mechanisms, consensus-based proofs, verification, and certification that contribute towards a Next Generation Trusted Internet supporting humanity in all aspects of life, 4. advanced concepts related to scalability, interoperability, energy efficiency, privacy, and security for making dependable building blocks of the Next Generation Internet, 5. advanced concepts of ontologies, resource models, reputation, and tokenization mechanisms for decentralised information and knowledge management, 6.

advanced concepts of Digital Twins.

Restrictions/Prerequisites: /

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.

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.

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.

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 the basic computer vision and deep learning approaches. Maximum number of students: 3.