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

Elective courses BDR-RI 2021/2022

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 two students. The lecturer of the course must not be the advisor or co-advisor of the student selecting the course. Each student can select only one individual research course.

Lecture type courses offered in 2021/2022:

Individual research type courses offered in 2021/2022:

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 problem. The course will stimulate competition between students to achieve the best possible learning results.

Modern Cryptography and Computer Security

(Selected Topics in Mathematical Methods in Computer Sciences 1, Polona Oblak)
Lecturer: Aleksandar Jurišić
Course code: 63828B
Course type: lectures, fall (first) semester

This course introduces modern cryptography, which provides maximum security while preserving the flexibility of digital media. It forms the foundation of Information Society (objectives: privacy, data integrity, digital authentication/signatures, digital cash, and other cryptographic protocols; it covers mathematics, computer science, electrical engineering, finances, policy, defense, *etc.*). We will discuss rigorous definitions of security in various situations. We will review basic mathematical problems used in cryptography and related complexity considerations. We will build more complex protocols from simpler primitives and prove that the more complex protocols meet their security objectives.

In this course, we will study selected topics from symmetric cryptography (classical ciphers and history of cryptography, stream ciphers, Shannon theory of information, block ciphers, cryptoanalysis and statistical methods, hash functions, authentication codes, birthday paradox attacks), asymmetric (public-key) cryptography (cryptosystems, digital signatures, key agreement protocols, identification schemes, other protocols), computer and information security (security of programs, security of databases, security of OS and real-time security management, security of network communications, privacy in CS, key managements, efficient and secure implementation of cryptosystems).

Advanced Topics in Ubiquitous Sensing and Learning

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

The course covers theoretical, system, and application aspects of the use of mobile, wearable, and the Internet of Things devices (here from 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 of these platforms, 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 component of the course 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.

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.

Optimization Methods for Large Networks

(Selected Topics in Mathematical Methods in Computer Sciences 1, Polona Oblak)
Lecturer: Aljaž Zalar
Course code: 63828C
Course type: lectures, spring (second) semester

In this course, we will study the theory behind polynomial optimization, which encompasses many NP-hard non-convex problems with applications in machine learning, statistics, control theory, quantum information theory, *etc*. We will focus on the so-called Lasserre type hierarchy schemes, which exploit certificates of positivity of polynomials and apply them to some concrete problems such as optimal power flow in a network, graph partitioning problems, global optimization of polynomials.

Heterogeneous Computing Platforms

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

In this course, we will deal with state-of-the-art platforms and technologies, which present an important 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. For a selected problem related to their doctoral thesis, they will have to recognize an interesting platform and then implement and evaluate their problem on it. In this course, we will study the speed-up of complex algorithms on modern hardware, how to combine CPU and custom FPGA circuits programmed in OpenCL, and how to analyze the effect of number representation to reduce computational cost and save energy.

Tensor Networks for Machine Learning

(Selected Topics in Artificial Intelligence 2, Zoran Bosnić)
Lecturer: Bojan Žunkovič
Course code: 63835E
Course type: individual research course, 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, and a vital ingredient of the state-of-the-art numerical techniques of many-body quantum mechanics. In the field of 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 automata and quadratic automata from the formal language's literature. The course will focus on recently developed tensor networks applications to 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 the techniques or applying the learned techniques to a slightly different problem. The proposed projects will be adapted to fit student interests, time, and expertise.

Selected Topics in Analysis of Sound Signals

(Selected Topics in Software Development 1, Matija Marot)
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*. Students will develop their 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 PhD students) around the world.

Computer Graphics and Visualization

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

Students will get to know 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 the application of appropriate algorithms and data structures for fast and real-time rendering, implementation of techniques on graphic processors, and remote rendering. The students will learn the benefits of modern graphics libraries (Vulkan, Web-GPU) 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 method 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 ad KAUST.

Low-Power Hardware Designs for Next-Generation Signal Processing and Machine Learning Applications

(Selected Topics in Computer Systems 1, Miha Mraz)
Lecturer: Patricio Bulić
Course code: 63830C
Course type: individual research course, spring (second) semester

The need to support various signal and media processing and recognition applications on energy-constrained mobile computing devices has steadily grown. There has been a growing interest in hardware neural networks in recent years, which express many benefits over conventional software models, mainly in applications where speed, cost, reliability, or energy efficiency are of great importance. Deep Neural Networks (DNN) are, presently, the most popular application models. Multilayered networks, characterized by many hidden layers and vast amounts of data to be trained, demand specialized, high-performance, low-power hardware architectures. DNN training and inference are computation-intensive processes: training requires high throughput, whereas inference needs a low latency.

In the last few years, FPGAs and GPUs vendors engaged in a race to offer the best hardware platform that runs computationally intensive algorithms quickly and efficiently. These algorithms' standard hardware implementations require many resources, power and time-consuming arithmetic operations (mainly multiplication). Hence the goal is to reduce the size and power consumption of internal arithmetic units. In particular, for large DNNs to run in real-time on resource-constrained systems, it is crucial to simplify/approximate tensor cores since they are usually responsible for the significant area, power and latency costs. One option to achieve this goal is to replace the complex exact arithmetic circuits with simpler, approximate ones. Approximate computing forms a design alternative that exploits the intrinsic error resilience of various applications and produces energy-efficient circuits with a small accuracy loss.

In the course, we will study the importance of low-power hardware designs, evaluate the accuracy of media processing algorithms and DNNs based on approximate computing, evaluate power reduction in approximate circuits and investigate training-time methodologies to compensate for the decrease of accuracy. During the course, the students will implement various circuits in FPGAs and evaluate them in terms of speed, area, and power consumption.