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

Elective courses BDR-RI 2025/2026

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 2025/2026: 2
Machine Learning and Artificial Intelligence 2
ChatGPT for Researchers 2
Advanced Topics in Edge Sensing and Learning 3
Selected Topics in Software Engineering 3
Incremental Learning from Data Streams 4
Individual research type courses offered in 2025/26 4
Advanced Topics in AI for Medicine (IR) 4
Evaluation in information systems research (IR) 5
Qualitative Research Methods in Computer and Information Science (IR) 5
Intelligent Embedded Systems (IR) 6
Selected Topics in Analysis of Sound Signals (IR) 6
Selected Topics from Computer Graphics and Visualization (IR) 7

Lecture type courses offered in 2025/2026:

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.

Advanced Topics in Edge 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 pertaining to the use of mobile, wearable, and the Internet of Things devices (herefrom referred to as "edge devices") for sensing and learning about the environment. The course starts with the overview of edge 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 - compressive sensing. The course then focuses on deep learning (DL) on edge devices, more specifically, on advanced topics, such as dynamic model compression techniques and running contemporary large language models (LLMs) on mobile devices. We then discuss distributed machine learning training on edge devices through split and federated learning. The course then presents in-depth investigation of applications of DL on edge devices, e.g. for healthcare, authentication and security. Finally, we present a critical analysis of sustainability in edge computing. 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 either a full-fledged edge deep learning application or generate new insights from pre-collected sensor data. Student participation is facilitated further by mandatory research paper presentations that will be delivered by each student in the class.

Keywords: Mobile computing, Edge computing, Mobile deep learning, Federated learning, Split learning, mHealth, Mobile security, Computing and sustainability, Smartphones, Wearables, Internet of Things (IoT), Android

Selected Topics in Software Engineering

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

In the duration of this course, selected chapters will be presented and analyzed across the domains of the Internet of Things, Artificial Intelligence, Cloud-to-Things computing, semantic technologies, Digital Twins, and blockchain technologies. By incorporating human-centric approaches to software development, probabilistic mechanisms to tackle bias in decision-making, decentralized identities, and secure, trustworthy data management, the course aims to foster novel, smart, and reliable services and applications. The course will cover the following areas:

1. Foundational 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 bias mitigation in artificial intelligence. probabilistic decision-making techniques for addressing bias in algorithms, human-in-the-loop methodologies for transparent and fair AI.

3. Decentralized democratic governance mechanisms, consensus-based proofs, verification and certification for the Next Generation Trusted Internet supporting humanity in all aspects of life.
4. Advanced concepts related to scalability, interoperability, energy efficiency, privacy and security. Designing the building blocks of the Next Generation Internet and dependable frameworks and infrastructure for global connectivity.
5. Advanced concepts of ontologies, resource models, reputation and tokenization mechanisms for decentralised information and knowledge management.
6. Advanced concepts of Digital Twins, Secure, trustworthy handling of real-time data streams, ensuring integrity and reliability across distributed systems.

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

Since the data streams are characterized as open-ended and often infinite sequences of data examples, of which data distribution is likely to change through time, this indicates the requirement for alternative learning algorithms to typical (batch) learning algorithms that learn from stationary datasets. The goal of the proposed course is to teach the students about the state-of-the-art algorithms that are 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 that is 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 that the students will do in the Statistical package R. The lab exercises will include applying the acquired knowledge on their own problem and stimulating a competition between students to achieve the best possible learning results.

Individual research type courses offered in 2025/26

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.

Evaluation in information systems research (IR)

(Selected Topics in Informatics 1, Denis Trček)
Lecturer: Damjan Vavpotič
Course code: 63826A
Course type: individual research course, fall (first) semester

The course covers key models and approaches in the field of evaluating the success of information systems (IS), user adoption of IS, information technologies (IT) and related processes and user satisfaction with IS/IT features. The student will use the selected evaluation approach in connection with his/her research field and apply it in practice in the form of a case study.

Qualitative Research Methods in Computer and Information Science (IR)

(Selected Topics in Informatics 2, Denis Trček)
Lecturer: Damjan Fujs
Course code: 63827
Course type: individual research course, spring (second) semester

The field of computer and information science has traditionally emphasized quantitative research methods. This course introduces a novel perspective by exposing doctoral students to alternative research approaches, specifically the application of qualitative methods. By integrating these methods, students can gain deeper and complementary insights to enrich their predominantly quantitative research. This course introduces doctoral students to key concepts, approaches, and techniques in qualitative research, with a focus on their application within computer and information science. Students will gain an understanding of the research underpinnings of qualitative inquiry, explore common data collection and analysis methods, and critically assess their use in empirical research. The course encourages reflection on ethical considerations, the researcher's role, and the value of qualitative approaches in addressing complex questions in socio-technical environments. Through independent and guided work, students will be supported in developing research designs aligned with qualitative traditions, applicable to their own doctoral projects or problem domains. The idea is for students to use a qualitative method as a supplement to their primary research method, for example, using a focus group, interview, case study, observation, quasi-experiment, experiment, Delphi method, action research, grounded theory, etc.

Intelligent Embedded Systems (IR)

(Selected Topics in Computer Systems 2, Miha Mraz)
Lecturer: Patricio Bulić, Luka Čehovin Zajc
Course code: 63831A
Course type: individual research course, spring (second) semester

This course focuses on the design, implementation, and deployment of AI/ML driven edge computing systems in real-time, resource-constrained environments. Students will explore modern artificial intelligence (AI) and machine learning (ML) techniques, implement algorithms on dual-core STM32H7 microcontrollers, and gain hands-on experience deploying PyTorch-trained models (exported to ONNX) using STM32Cube.AI. The course integrates concepts from AI/ML, embedded systems, and energy-efficient programming to build end-to-end edge applications.
While AI/ML models are commonly deployed in high-level programming environments such as Python or on single-board computers running Linux and Python interpreter, real-time embedded systems often require a different approach. This course emphasizes the deployment of AI/ML models on resourceconstrained microcontrollers using low-level programming in C/C++. Such an approach ensures realtime performance, deterministic execution, and energy efficiency—critical factors for applications in robotics, automotive systems, and industrial automation. Students will learn how to optimize models for embedded execution, efficiently manage memory and compute resources, and implement inference pipelines without relying on an operating system. By writing software in C/C++, students will ensure real-time performance, deterministic execution, and energy efficiency.
The STM32 Model Zoo will be a key resource in our course, enabling students to explore AI applications on embedded devices and integrate machine learning models into real-world chip design projects. It will provide hands-on experience with deploying and customizing machine learning models on low-power, resource-constrained microcontrollers, an essential skill in modern embedded system development. The STM32 Model Zoo is a collection of pre-trained machine learning models optimized for deployment on STM32 microcontrollers. Developed by STMicroelectronics, it provides a valuable resource for integrating edge AI capabilities into embedded systems. This tool offers models tailored for various applications, such as image classification, object detection, and audio recognition, and is particularly suitable for use with STM32 microcontrollers, including those with the Neural-ART Accelerator (NPU).
We plan to utilize the STM32N6 and STM32H7 series in the course to demonstrate the practical application of AI models and machine learning on embedded microcontrollers, providing students with hands-on experience in deploying efficient AI solutions on low-power devices. The STM32N6 series is a family of microcontrollers from STMicroelectronics, designed specifically for AI and machine learning applications at the edge. These microcontrollers are equipped with the Neural-ART Accelerator, a hardware-based accelerator optimized for running machine learning algorithms efficiently. The STM32H7 dual-core series is also an excellent, low-cost platform for deploying AI and machine learning algorithms in embedded systems. Its unique dual-core architecture, combining a high-performance Arm Cortex-M7 core and a lower-power Arm Cortex-M4 core, enables efficient handling of both complex AI tasks and resource-constrained operations. This flexibility makes the STM32H7 series ideal for edge AI applications, where real-time decision-making and efficient resource management are crucial.

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: /