

Research Project Artificial Intelligence 1

Dept. of Advanced Computing Sciences

KEN4130

Semester 1:

1 Sep 2025

23 Jan 2026

Credits:

6.0

Coordinator:

L. Rieswijk

Teaching methods:

Skills, Work in subgroups, Presentation(s), Project-Centered Learning

Assessment methods:

Participation, Assignment, Presentation and paper

Full course description

The research project takes place during the three periods of the semester.

Project topics are presented at the start of the semester and assigned to students based on their preferences and availability.

The emphasis in the first phase is on initial study of the context set out for the project and the development of a project plan.

In the second phase, the goal is to start modelling, prototyping and developing.

In phase 3, the implementation, model and/or experiments set out in the project plan has to be finished and reported on.

In week 4 of period 2 a progress presentation takes place. The project results in a project presentation, a project report and possibly a public website and/or product.

The Research Project 1 will start in period 1.1 and 1.2 with weekly meetings. The credits for the project will become available at the end of period 1.3.

Prerequisites

None.

Recommended reading

Recommended readings are project dependent

Intelligent Search & Games

Dept. of Advanced Computing Sciences

KEN4123

Period 1:

1 Sep 2025

24 Oct 2025

Credits:

6.0

Coordinator:

D.J.N.J. Soemers M.H.M. Winands

Teaching methods:

Lecture(s)

Assessment methods:

Written exam, Assignment

Full course description

In this course the students learn how to apply advanced techniques in the framework of game-playing programs. Depending on the nature of the game these techniques can be of a more or less algorithmic nature. The following subjects will be discussed:

- a. Basic search techniques. Alpha-beta; A*.
- b. (2) Advanced search techniques. IDA*; transposition tables; retrograde analysis; proof-number search variants; multi-player search methods; Expectimax and *-minimax variants.
- c. (3) Heuristics. Move ordering, windowing techniques (PVS); forward-pruning techniques; selective search.
- d. (4) Monte Carlo methods. Monte Carlo Tree Search (MCTS) techniques, enhancements, and applications; AlphaGo and AlphaZero approaches.
- e. (5) Video game AI techniques: World representations, GOAP, hierarchical task networks, behavior trees.

Prerequisites

None.

Desired Prior Knowledge : Data Structures & Algorithms.

Recommended reading

- Millington, I. (2019). Artificial Intelligence for Games, 3rd Edition, CRC Press, ISBN: 978-1138483972
- Russell, S.J. and Norvig, P. (2020). Artificial Intelligence: A Modern Approach, 4th edition. Pearson. ISBN 0-13-461099-7.
- Yannakakis, G.N. and Togelius, J. (2025) Artificial Intelligence and Games, 2nd edition, Springer, Berlin. ISBN-13. 978-3031833465

Advanced Concepts in Machine Learning

Dept. of Advanced Computing Sciences

KEN4154

Period 2:

27 Oct 2025

12 Dec 2025

Credits:

6.0

Coordinator:

E. Hortal Quesada

Teaching methods:

Project-Centered Learning

Assessment methods:

Written exam, Assignment

Full course description

This course will introduce a number of advanced concepts in the field of machine learning such as Support Vector Machines, Gaussian Processes, Deep Neural Networks, etc. All of these are approached from the view that the right data representation is imperative for machine learning solutions. Additionally, different knowledge representation formats used in machine learning are introduced. This course counts on the fact that basics of machine learning were introduced in other courses so that it can focus on more recent developments and state of the art in machine learning research. Labs and assignments will give the students the opportunity to implement or work with these techniques and will require them to read and understand published scientific papers from recent Machine Learning conferences.

Prerequisites

Desired Prior Knowledge: Machine Learning

Recommended reading

Pattern Recognition and Machine Learning - C.M. Bishop; Bayesian Reasoning and Machine Learning - D. Barber; Gaussian Processes for Machine Learning - C.E. Rasmussen & C. Williams; The Elements of Statistical Learning - T. Hastie et al.

Research Project Artificial Intelligence 2

Dept. of Advanced Computing Sciences

KEN4131

Semester 2:

26 Jan 2026

19 Jun 2026

Credits:

6.0

Coordinator:

L. Rieswijk

Teaching methods:

Skills, Work in subgroups, Presentation(s), Project-Centered Learning

Assessment methods:

Participation, Assignment, Presentation and paper

Full course description

The research project takes place during the three periods of the semester. Project topics are presented at the start of the semester and assigned to students based on their preferences and availability.

The emphasis in the first phase is on initial study of the context set out for the project and the development of a project plan. In the second phase, the goal is to start modelling, prototyping and developing. In phase 3, the implementation, model and/or experiments set out in the project plan has to be finished and reported on. In week 4 of period 5 a progress presentation takes place. The project results in a project presentation, a project report and possibly a public website and/or product.

The Research Project 2 will start in period 1.4 and 1.5 with weekly meetings. The credits for the project will become available at the end of period 1.6.

Prerequisites

None.

Recommended reading

Recommended readings are project dependent

Agents and Multi-Agent Systems

Dept. of Advanced Computing Sciences

KEN4111

Period 4:

26 Jan 2026

27 Mar 2026

Credits:

6.0

Coordinator:

G.B. Weiss

Teaching methods:

Project-Centered Learning

Assessment methods:

Written exam, Assignment

Full course description

The notion of an (intelligent) agent is fundamental to the field of artificial intelligence. Thereby an agent is viewed as a computational entity such as a software program or a robot that is situated in some environment and that to some extent is able to act autonomously in order to achieve its design objectives. The course covers important conceptual, theoretical and practical foundations of single-agent systems (where the focus is on agent-environment interaction) and multi-agent systems (where the focus is on agent-agent interaction). Both types of agent-based systems have found their way to real-world applications in a variety of domains such as e-commerce, logistics, supply chain management, telecommunication, health care, and manufacturing. Examples of topics treated in the course are agent architectures, computational autonomy, game-theoretic principles of agent-based systems, coordination mechanisms (including auctions and voting), and automated negotiation and argumentation. Other topics such as ethical or legal aspects raised by computational agency may also be covered. In the exercises and in the practical part of the course students have the opportunity to apply the covered concepts and methods.

Prerequisites

Desired Prior Knowledge: Basic knowledge and skills in programming.

Late enrollment for this course is not permitted, as participation in mandatory assignments and group work is required from the start.

Recommended reading

- Stuart Russell and Peter Norvig (2010). Artificial Intelligence. A Modern Approach. 3rd edition. Prentice Hall.
- Gerhard Weiss (Ed.) (2013, 2nd edition): Multi-agent Systems. MIT Press.
- Mike Wooldridge (2009, 2nd edition): An Introduction to Multi Agent Systems, John Wiley & Sons Ltd.
- Yoav Shoham and Kevin Leyton-Brown (2009): Multi-agent Systems. Algorithmic, Game-Theoretic, and Logical Foundations, Cambridge University Press.

Autonomous Robotic Systems

Dept. of Advanced Computing Sciences

KEN4114

Period 5:

30 Mar 2026

22 May 2026

Credits:

6.0

Coordinator:

R. Möckel

Teaching methods:

Project-Centered Learning

Assessment methods:

Written exam, Assignment

Full course description

Operating autonomously in unknown and dynamically changing environments is a core challenge that all robotic systems must solve to work successfully in industrial, public, and private areas. Currently popular systems that must demonstrate such capabilities include self-driving cars, autonomously operating drones, and personal robotic assistants. In this course, students obtain deep knowledge in creating autonomous robotic systems that can operate in and manipulate unknown and dynamically changing environments by autonomously planning, analysing, mapping, and modelling of such environments. Students learn to approach these challenging tasks through three main techniques: swarm intelligence, model-based probabilistic frameworks, and (mostly) model-free techniques from artificial evolution and machine learning.

Prerequisites

None.

Desired Prior Knowledge: Discrete Mathematics, Linear Algebra, Probabilities and Statistics, Data Structures and Algorithms, Machine Learning, Search Techniques.

Recommended reading

- Floreano and Nolfi (2000), Evolutionary Robotics, The MIT press. ISBN-13: 978-0262640565.

- Dario Floreano und Claudio Mattiussi (2008), Bio-Inspired Artificial Intelligence: Theories, Methods, and Technologies, ISBN-13: 978-0262062718

Signal and Image Processing

Dept. of Advanced Computing Sciences

KEN4222

Period 1:

1 Sep 2025

24 Oct 2025

Credits:

6.0

Coordinator:

J.M.H. KarelP. Bonizzi

Teaching methods:

Lecture(s), Computer Labs

Assessment methods:

Written exam, Computer test

Full course description

This course offers the student a hands-on introduction into the area of digital signal and image processing. We start with the fundamental concepts and mathematical foundation. This includes a brief review of Fourier analysis, z-transforms and digital filters. Classical filtering from a linear systems perspective is discussed. Next wavelet transforms and principal component analysis are introduced. Wavelets are used to deal with morphological structures in signals. Principal component analysis is used to extract information from high-dimensional datasets. We then discuss Hilbert-Huang Transform to perform detailed time-frequency analysis of signals. Attention is given to a variety of objectives, such as detection, noise removal, compression, prediction, reconstruction and feature extraction. We discuss a few cases from biomedical engineering, for instance involving ECG and EEG signals. The techniques are explained for both 1D and 2D (images) signal processing. The subject matter is clarified through exercises and examples involving various applications. In the practical classes, students will apply the techniques discussed in the lectures using the software package Matlab.

Prerequisites

Desired Prior Knowledge: Linear algebra, Calculus, basic knowledge of Matlab. Some familiarity with linear systems theory and transforms (such as Fourier and Laplace) is helpful.

Recommended reading

Principal Component Analysis, Ian T. Jolliffe, Springer, ISBN13: 978-0387954424.

Optimization

Dept. of Advanced Computing Sciences

KEN4211

Period 1:

1 Sep 2025

24 Oct 2025

Credits:

6.0

Coordinator:

P.J. Collins

Teaching methods:

Project-Centered Learning

Assessment methods:

Written exam

Full course description

Optimization is the subject of finding the best or optimal solution to a problem from a set of potential or feasible solutions. Optimization problems are fundamental in all forms of decision-making, since one wishes to make the best decision in any context, and in the analysis of data, where one wishes to find the best model describing experimental data.

This course treats one of the main branches of optimization theory, namely nonlinear programming, which covers a wide range of real life problems including the big data analysis for data analysis and machine learning. It builds on knowledge of linear programming using the simplex algorithm. We will study many general-purpose algorithms, including Brent's method for one-dimensional problems, (quasi-)Newton and conjugate gradient methods for unconstrained problems, and Lagrangian methods, including active-set methods, sequential quadratic programming and interior-point methods for general constrained problems. We will also look at algorithms for global optimization, including branch-and-bound and heuristic approaches such as simulated annealing. Finally, we will cover algorithms for big data, including stochastic gradient descent, and advanced batch optimization methods.

Throughout the course, we aim to provide a coherent framework for the subject, with a focus on optimality conditions (notably the Karush-Kuhn-Tucker conditions), convexity, Lagrange multipliers and duality, and convergence rates. The methods will be illustrated by in-class computer

demonstrations, exercises illustrating the main concepts and algorithms, modelling and computational work on case studies of practical interest in data science and machine learning, including training neural networks.

Prerequisites

Desired Prior Knowledge: Simplex algorithm. Calculus, Linear Algebra.

Recommended reading

J. Nocedal and S.J. Wright, “Numerical Optimization”, Springer, 2006; ISBN: 978-0-387-30303-1.

Léon Bottou, Frank E. Curtis, and Jorge Nocedal. “Optimization methods for large-scale machine learning”, *SIAM Review*, 60(2):223–311, 2018.

Data Mining

Dept. of Advanced Computing Sciences

KEN4113

Period 1:

1 Sep 2025

24 Oct 2025

Credits:

6.0

Coordinator:

E.N. Smirnov

Teaching methods:

Project-Centered Learning

Assessment methods:

Written exam, Assignment

Full course description

Data mining is a major frontier field of computer science. It allows extracting useful and interesting patterns and knowledge from large data repositories such as databases and the Web. Data mining integrates techniques from the fields of databases, machine learning, statistics, and artificial intelligence. This course will present the state-of-the-art techniques of data mining. The lectures and labs will emphasize the practical use of the presented techniques and the problems of developing real data-mining applications. A step-by-step introduction to data-mining environments will enable the students to achieve specific skills, autonomy, and hands-on experience. A number of real data sets will be analysed and discussed.

Prerequisites

None.

Recommended reading

Pang-Ning, T., Steinbach, M., Karpatne, A., and Kumar, V. (2018). Introduction to Data Mining, 2nd Edition, Pearson, ISBN-10: 0133128903, ISBN-13: 978-0133128901

Control and Intelligent Systems

Dept. of Advanced Computing Sciences

KEN4252

Period 1:

1 Sep 2025

24 Oct 2025

Credits:

6.0

Coordinator:

B. SakçakP.W.L. Dreesen

Teaching methods:

Project-Centered Learning

Assessment methods:

Written exam, Computer test, Take home exam

Full course description

This course explores control theory from classical principles to modern state-of-the-art methods, with applications to robotics, autonomous vehicles, and intelligent systems. Starting from foundational concepts, we explore the control systems that underlie modern engineering and technology. We will start by reviewing relevant models of dynamical systems and analysis and synthesis of feedback control systems. Then, we move the focus to modern state-space control methods, including state feedback, state observers, output tracking, and optimal control. In the last part of the course, we will discuss advanced concepts such as control of uncertain systems, nonlinear control, and data-driven control, as well as the links between control theory and reinforcement learning, and robotics.

Emphasizing digital control and data-driven approaches, the course connects 'Control and Intelligent Systems' to the broader fields of engineering and science, with specific attention to data science and robotics. Hands-on computer sessions and case studies on applications reinforce key concepts. At the end of the course, the students are expected to build a strong theoretical foundation. Furthermore, through selected case studies, the students will gain hands-on experience to design, analyze, and deploy control solutions for complex, real-world problems.

Prerequisites

None

Recommended reading

None

Algorithmic Game Theory

Dept. of Advanced Computing Sciences

KEN4251

Period 2:

27 Oct 2025

12 Dec 2025

Credits:

6.0

Coordinator:

M. Salvioli

Teaching methods:

Project-Centered Learning

Assessment methods:

Written exam

Full course description

The course will focus on non-cooperative games and on dynamic games in the following order: matrix and bimatrix games, repeated games, Stackelberg games, differential games, specific models of stochastic games, evolutionary games. These are games in which the players are acting as strategic decision makers, who cannot make binding agreements to achieve their goals. Instead, threats may be applied to establish stable outcomes. Besides, relations with population dynamics and with “learning” will be examined. Several examples will be taken from biological settings.

Prerequisites

Desired Prior Knowledge: Students are expected to be familiar with basic concepts from linear algebra, calculus, Markov chains and differential equations.

Recommended reading

None.

Advanced Natural Language Processing

Dept. of Advanced Computing Sciences

KEN4259

Period 2:

27 Oct 2025

12 Dec 2025

Credits:

6.0

Coordinator:

A.S. Härmä J.C. Scholtes

Teaching methods:

Project-Centered Learning

Assessment methods:

Written exam, Assignment

Full course description

For decades, teaching a computer to deal with natural language processing (NLP) was a long-time dream of humankind. Tasks such as machine translation, summarization, question-answering, speech recognition or chatting remained a challenge for computer programs. Around 2018, major improvements were made. Starting with machine translation and ultimately in late 2022 with ChatGPT. Why were these large-language models suddenly so good? How did we get here? What can we do with these new algorithms to improve them even more?

This course will provide the skills and knowledge to understand and develop state-of-the-art (SOTA) solutions for these natural language processing (NLP) tasks. After a short introduction to traditional generative grammars and statistical approaches to NLP, the course will focus on deep learning techniques. We will discuss Transformers, variations on their architecture (including BERT, GPT and multi-modal) in depth, which models work best for which linguistic tasks, their capacities, limitations and how to address these.

Although computer systems can now use Natural Language in ways that can no longer be distinguished from humans, there are still major problems to address: (i) we do not fully understand what these algorithms know and what they do not know. So, there is a strong need for explainable AI (XAI) in NLP. (ii) Training the deep-learning large language-models costs too much energy. In the lectures we will discuss models and methods that are less computationally (and thus energy)

intensive. (iii) New methods aim to go from having a simple conversation to reasoning and solving complex problems using agentic techniques and reinforcement learning. In this course, we will also discuss these topics.

This course is closely related with the course Information Retrieval and Text-Mining (IRTM). The ANLP course is more focussed on the inner workings of the methods and architectures to deal with complex natural language tasks. The IRTM course focusses more on building search engines, using text-analytics for deeper understanding and the requirements needed for responsible conversational search. The overlap between the two courses is kept to a minimum. There is no need to follow the courses in a specific order.

Prerequisites

Proficient coding skills in Python are required to participate in the tutorials.

Google CoLAB pro subscription for the tutorials.

Recommended reading

Papers published in top international conferences and journals in machine learning field.

Network Science

Dept. of Advanced Computing Sciences

KEN4275

Period 2:

27 Oct 2025

12 Dec 2025

Credits:

6.0

Coordinator:

A.I. Iamnitchi

Teaching methods:

Assignment(s), Work in subgroups, Project-Centered Learning

Assessment methods:

Assignment

Full course description

Many aspects of everyday life and science can be represented as networks: social networks represent relationship (links) between people (nodes); brain activity can be represented via synapses (links) between neurons (nodes); the street map is formed of roads (links) that connect at intersections (nodes); authors of scientific papers connect to each others in a citation network, with directed links from the paper cited to the paper citing it; communication networks connect routers via physical or logical links; etc. Network analysis plays a significant role in the “big data” analytics because of size, data velocity, or computational complexity.

This course focuses on the study of network structures and dynamic processes on networks using real data from various disciplines, including socio-technological platforms, biology, social science, and economics. Topics cover the analysis and modeling of complex networks, network dynamics, community detection, network resilience and contagion, as well as processing of network structures for machine learning tasks.

Prerequisites

Desired prior knowledge: Introductory knowledge of programming for data analysis, particularly in Python; algorithms; algorithmic complexity.

Prerequisites: Introductory courses in algorithms, data structures, and data analytics.

Recommended reading

"Networks, Crowds, and Markets: Reasoning About a Highly Connected World" by David Easley and Jon Kleinberg.

Additional literature: Graph Theory and Complex Networks: An Introduction by Maarten van Steen

Planning and Scheduling

Dept. of Advanced Computing Sciences

KEN4253

Period 4:

26 Jan 2026

27 Mar 2026

Credits:

6.0

Coordinator:

S.M. Kelk

Teaching methods:

Project-Centered Learning

Assessment methods:

Written exam, Assessment

Full course description

In many real-world processes, particularly in industrial processes and logistics, decisions need to be taken about the time of the completion of (sub)tasks, and the decision about what production machines complete the tasks. There are often constraints on the order in which tasks, or 'jobs' can be performed, and there are usually capacity constraints of the machines. This leads to natural, industrially critical optimization problems. For example, a company might choose to buy many machines to process jobs, but then there is a risk that the machines will be underused, which is economically inefficient. On the other hand, too few machines, or an inappropriate ordering of tasks, may lead to machines spending a significant amount of time standing idle, waiting for the output of other machines, which are overcrowded with tasks. In this course, we look at various mathematical models and techniques for optimizing planning and scheduling problems, subject to different optimality criteria. We will discuss, among others, single-machine models, parallel-machine models, job-shop models, and algorithms for planning and scheduling (exact, approximate, heuristic) and we also touch upon the computational complexity (distinguishing between 'easy' and 'difficult' problems) of the underlying problems. Last but not least, we will also introduce integer linear programming as a uniform and generic tool to model and solve planning and scheduling problems.

Prerequisites

None.

Desired prior knowledge: Data Structures & Algorithms. Discrete Mathematics. Graph Theory

Recommended reading

None.

Building and Mining Knowledge Graphs

Dept. of Advanced Computing Sciences

KEN4256

Period 4:

26 Jan 2026

27 Mar 2026

Credits:

6.0

Coordinator:

Y. YangM.J. Dumontier

Teaching methods:

Project-Centered Learning

Assessment methods:

Assignment, Assessment

Full course description

Knowledge graphs are large-scale, machine-processable representations of entities, their attributes, and their relationships. Knowledge graphs enable both people and machines to explore, understand, and reuse information in a wide variety of applications such as answering questions, finding relevant content, understanding social structures, and making scientific discoveries. However, the sheer size and complexity of these graphs present a formidable challenge particularly when mining across different topic areas.

In this course, we will examine approaches to construct and use knowledge graphs across a diverse set of applications using cutting-edge technologies such as machine learning and deep learning, graph databases, ontologies and automated reasoning, and other relevant techniques in the area of data mining and knowledge representation.

Prerequisites

Desired Prior Knowledge: Introduction to Computer Science

Recommended reading

Aggarwal, C.C. and Wang, H. eds., (2010) Managing and mining graph data (Vol. 40). New York: Springer. ISBN 978-1-4419-6045-0

Explainable AI

Dept. of Advanced Computing Sciences

KEN4246

Period 4:

26 Jan 2026

27 Mar 2026

Credits:

6.0

Coordinator:

M.F.M. SondagN. Tintarev

Teaching methods:

Project-Centered Learning

Assessment methods:

Group Assignment, Written Exam (Literature Review)

Full course description

A key component of an artificially intelligent system is the ability to explain to a human agent the decisions, recommendations, predictions, or actions made by it and the process through which they are made. Such explainable artificial intelligence (XAI) can be required in a wide range of applications. For example, a regulator of waterways may use a decision support system to decide which boats to check for legal infringements, a concerned citizen might use a system to find reliable information about a new disease, or an employer might use an artificial advice-giver to choose between potential candidates fairly.

For explanations from intelligent systems to be useful, they need to be able to justify the advice they give in a human-understandable way. This creates a necessity for techniques for automatic generation of satisfactory explanations that are intelligible for users interacting with the system. This interpretation goes beyond a literal explanation. Further, understanding is rarely an end-goal in itself. Pragmatically, it is more useful to operationalize the effectiveness of explanations in terms of a specific notion of usefulness or explanatory goals such as improved decision support or user trust. One aspect of intelligibility of an explainable system (often cited for domains such as health) is the ability for users to accurately identify, or correct, an error made by the system. In that case it may be preferable to generate explanations that induce appropriate levels of reliance (in contrast to over- or under-reliance), supporting the user in discarding advice when the system is incorrect, but also accepting correct advice.

The following subjects will be discussed:

- (1) Intrinsically interpretable models, e.g., decision trees, decision rules, linear regression.
- (2) Identification of violations of assumptions; such as distribution of features, feature interaction, non-linear relationships between features; and what to do about them.
- (3) Model agnostic explanations, e.g., LIME, scoped Rules (Anchors), SHAP (and Shapley values)
- (4) Ethics for explanations, e.g., fairness and bias in data, models, and outputs.
- (5) (Adaptive) User Interfaces for explainable AI
- (6) Evaluation of explanation understandability

Prerequisites

Desired Prior Knowledge: Data Analysis and Data Mining

Maximum number of 60 students can follow this course from MSc AI or MSc MSDM.

Mandatory for MSc RDS (no cap).

Late enrollment for this course is not permitted, as participation in mandatory assignments and group work is required from the start.

Recommended reading

- Molnar, Christoph. Interpretable Machine Learning (**3rd edition**) . Lulu. com, **2025**.
- Rothman, Denis. Hands-On Explainable AI (XAI) with Python: Interpret, visualize, explain, and integrate reliable AI for fair, secure, and trustworthy AI apps, Packt, 2020.

Information Retrieval and Text Mining

Dept. of Advanced Computing Sciences

KEN4153

Period 5:

30 Mar 2026

22 May 2026

Credits:

6.0

Coordinator:

J.C. Scholtes

Teaching methods:

Project-Centered Learning

Assessment methods:

Written exam, Assessment

Full course description

Information Retrieval and Text Mining aka Conversational Search Algorithms

This course is all about algorithms for Conversational Search

Traditional search (aka Information Retrieval) is about “Finding the needle in the haystack”. This course focusses on more complex problems: “How does the needle look like and where is the haystack?” using text-mining, topic modeling and data visualization techniques. Today, search is more about, “Having a guide handing you the needle and explaining you its significance and context” using conversational AI in combination with search engines and text-mining.

Building a full-text search engine may look trivial, but it is not! How do you search hundreds of billions of documents that can be located anywhere, with sub-second responds times? How do you find exactly what you are looking for without missing relevant information or having to plough through hundreds of irrelevant documents? How can you find if you do not know exactly what you are looking for? How can you find information which is deliberately hidden? How do you know that your search engine has given you the right information? Where does it come from? Is the answer factually correct?

In this course, we will teach you how to address these questions in three steps: (1) how is a search engine is constructed, optimized and used effectively, (2) How can techniques from the world of text-

mining, information extraction, text classification, clustering, topic modelling and data visualization lead to better search, and (3) What is the best way to integrate chatbots with search engines resulting in responsible conversational search. How to best guarantee factuality, measure and prevent hallucinations, provide provenance and explainability of the chatbots' recommendations. How can we integrate knowledge graphs and retrieval augmented generation (RAG) in the conversation? How can agentic architectures help us to overcome the limitations of LLMs?

Interestingly, traditional search-engine and text-mining techniques are essential components of such responsible conversational search solutions. In this course we will discuss all components and how they best work together.

Linear Algebra, Statistics, Deep Learning and Natural Language Processing play an important role in this course.

This course is complementary to the course Advanced Natural Language Processing (ANLP). Overlap is reduced to the necessary minimum. Both courses can be followed in any particular order. In the Information Retrieval and Text Mining course we focus more on creating an optimal modern search experience, in the Advanced Natural Language Processing course, we do a deep dive into the algorithms and models used for different language-related problems such as machine translation, abstracting, dialogs with chatbots, and LLM-based agents. Tutorials are shared between the two courses.

Prerequisites

Proficient coding skills in Python are required to participate in the tutorials.

Google CoLAB pro subscription for the tutorials.

Recommended reading

Introduction to Information Retrieval. Christopher D. Manning, Prabhakar Raghavan and Hinrich Schütze. Cambridge University Press, 2008. In bookstore and online: <http://informationretrieval.org> .

Computer Vision

Dept. of Advanced Computing Sciences

KEN4255

Period 5:

30 Mar 2026

22 May 2026

Credits:

6.0

Coordinator:

M.C. Popa

Teaching methods:

Project-Centered Learning

Assessment methods:

Written exam, Assignment

Full course description

Can we make machines look, understand and interpret the world around them? Can we make cars that can autonomously navigate in the world, robots that can recognize and grasp objects and, ultimately, recognize humans and communicate with them? This course will provide the knowledge and skills that are fundamental to core vision tasks of one of the fastest growing fields in academia and industry: visual computing. Topics include introduction to fundamental problems of computer vision, mathematical models and computational methodologies for their solution including both traditional and newer techniques such as Visual Language Models (VLMs), implementation of real-life applications and experimentation with various techniques in the field of scene analysis and understanding. In particular, after a recap of basic image analysis tools (enhancement, restoration, color spaces, edge detection), students will learn about feature detectors and trackers, fitting, image geometric transformation and mosaicing techniques, texture analysis and classification using unsupervised techniques, face and emotion analysis, deep learning based object classification using both CNNs and transformers, detection and tracking, camera models, epipolar geometry and 3D reconstruction from 2D views.

Prerequisites

None.

Desired prior knowledge: Basic knowledge of Python, linear algebra and machine learning. This course offers the basics on image processing although prior knowledge is also a plus.

Recommended reading

“Digital Image Processing”, Rafael C. Gonzalez & Richard E. Woods, Addison-Wesley, “Computer Vision: Models, Learning and Inference”, Simon J.D. Prince 2012.

Introduction to Quantum Computing for AI and Data Science

Dept. of Advanced Computing Sciences

KEN4155

Period 5:

30 Mar 2026

22 May 2026

Credits:

6.0

Coordinator:

D.O. MestelG. Stamoulis

Teaching methods:

Project-Centered Learning

Assessment methods:

Written exam

Full course description

In this course we lay down the foundations and basic concepts of quantum computing. We will use the mathematical formalism borrowed from quantum mechanics to describe quantum systems and their interactions. We introduce the concept of a quantum bit and discuss different physical realizations of it. We then introduce the basic building blocks of quantum computing: quantum measurements and quantum circuits, single and multi-qubit gates, the difference between correlated (entangled) and uncorrelated states and their representation, quantum communication, and basic quantum protocols and quantum algorithms. Finally, we discuss the different types of noise involved in real quantum computers (coherent and incoherent errors, state preparation, projection and measurement) and their effect on performance, and outline current efforts for mitigating the issues.

This course is a prerequisite for the elective courses Quantum Algorithms, Quantum AI, and Quantum Information and Security. These four courses, together with a dedicated research project quantum computing forms the specialization Quantum Computing.

Prerequisites

Desired prior knowledge: probability theory, linear algebra, design and analysis of algorithms.

This course is a prerequisite for the elective courses "Quantum Algorithms", "Quantum AI", "Quantum Information and Security" and the project "Research Project Quantum Computing".

These four courses, together with a dedicated research project on Quantum Computing forms the specialization in Quantum Computing for AI and Data Science.

Recommended reading

"Quantum Computation and Quantum Information" Michael Nielsen and Isaac Chuang, Cambridge University Press, 10th Anniversary edition.

Reinforcement Learning

Dept. of Advanced Computing Sciences

KEN4157

Period 5:

30 Mar 2026

22 May 2026

Credits:

6.0

Coordinator:

D.J.N.J. SoemersK. Driessens

Teaching methods:

Project-Centered Learning

Assessment methods:

Assignment, Take home exam

Full course description

Reinforcement learning is a type of machine learning problem in which the learner gets a (delayed) numerical feedback signal about its demonstrated performance. It is the toughest type of machine learning problem to solve, but also the one that best encompasses the idea of artificial intelligence as a whole. In this course we will define the components that make up a reinforcement learning problem and will see what the important concepts are when trying to solve such a problem, such as state and action values, policies and performance feedback. We will look at the different properties a reinforcement learning problem can have and what the consequences of these properties are with respect to solvability. We will discuss value based techniques as well as direct policy learning and learn how to implement these techniques. We will study the influence of generalisation on learning performance and see how supervised learning (and specifically deep learning) can be used to help reinforcement learning techniques tackle larger problems. We will also look at the evaluation of learned policies and the development of performance over time.

Prerequisites

No hard prerequisites but having some background in Machine Learning and/or Data Mining will be helpful.

Recommended reading

Lecture slides will be uploaded before each lecture. These slides are designed and intended as support during teaching, not as study material by themselves. They are supplied as a service, but additional note taking will be necessary to pass the class.

The book “Reinforcement Learning – An Introduction” by Sutton and Barto is freely available at: <https://www.andrew.cmu.edu/course/10-703/textbook/BartoSutton.pdf>

Research Project Quantum Computing

Dept. of Advanced Computing Sciences

KEN4330

Semester 1:

1 Sep 2025

23 Jan 2026

Credits:

6.0

Coordinator:

L. Rieswijk

Teaching methods:

Skills, Work in subgroups, Project-Centered Learning

Assessment methods:

Presentation and paper

Full course description

The research project takes place during the three periods of the semester. Project topics are presented at the start of the semester and assigned to students based on their preferences and availability.

The emphasis in the first phase is on initial study of the context set out for the project and the development of a project plan. In the second phase, the goal is to start modelling, prototyping and developing. In phase 3, the implementation, model and/or experiments set out in the project plan has to be finished and reported on. . In week 4 of period 2 a progress presentation takes place. The project results in a project presentation, a project report and possibly a public website and/or product.

The Research Project 1 will start in period 1.1 and 1.2 with weekly meetings. The credits for the project will become available at the end of period 1.3.

Prerequisites

Prerequisites: Introduction to Quantum Computing for AI and Data Science

Recommended reading

Recommended readings are project dependent

Study Abroad

Dept. of Advanced Computing Sciences

KEN3600

Semester 1:

1 Sep 2025

Semester 2:

26 Jan 2026

Credits:

30.0

Coordinator:

M. Musegaas

Assessment methods:

Written exam, Attendance, Assignment

23 Jan 2026

19 Jun 2026

Full course description

DACS offers its students the possibility to study a semester abroad at one of DACS partner universities. Third year bachelor's students and 2nd year master's students can get the opportunity to study a semester abroad, as part of their education programme in Maastricht. The credits received abroad will be transferred / part of your programme at DACS in Maastricht. Of course, this is only possible after approval of the Board of Examiners. There are several universities where DACS can send its students to.

If you still have questions afterwards you may contact our Exchange Coordinator Luc Giezenaar via: dacs-iro@maastrichtuniversity.nl

Prerequisites

You have to obtain at least 40 ECTS of year 1 courses.

Master Internship

Dept. of Advanced Computing Sciences

KEN4176

Year:

1 Sep 2025

19 Jun 2026

Credits:

30.0

Coordinator:

K. Driessens

Assessment methods:

Final paper

Full course description

In the master elective semester of year 2, students have the opportunity to do an internship at a company or research institution. The internship can be full-time (30 ECTS) or part-time (at least 10 ECTS). Please be aware that 1 ECTS is the equivalent of 28 working hours (which means a full-time internship involves 840 hours).

Internship vacancies can be found on Canvas and Intranet. In addition to the vacancies offered by DACS, students are also allowed to find something themselves. For each internship, an internship proposal must be sent to the Board of Examiners for approval. This proposal must be checked, approved and signed by a DACS supervisor and a company supervisor before submitted to the Board of Examiners. After receiving the official approval of the Board of Examiners, the student is allowed to start.

See <https://intranet.maastrichtuniversity.nl/en/dacs-students/my-studies/elective-semester/master-internship> for the complete procedure and relevant forms. The guidelines for internships are included in the Rules and Regulations.

More information or questions? Contact our internship coordinator Anouk Quaden.

Prerequisites

The Master Internship can only be done in the second year of the Master. You need to have obtained at least 40 ECTS of year 1 courses.

Only after you received the official approval by the Board of Examiners, you are allowed to start your internship (!). Never start before the approval; you are taking the risk that the internship might not be approved at all, and in that case, you will not receive study credits of the internship for the period prior to having obtained official approval.

Optimization

Dept. of Advanced Computing Sciences

KEN4211

Period 1:

1 Sep 2025

24 Oct 2025

Credits:

6.0

Coordinator:

P.J. Collins

Teaching methods:

Project-Centered Learning

Assessment methods:

Written exam

Full course description

Optimization is the subject of finding the best or optimal solution to a problem from a set of potential or feasible solutions. Optimization problems are fundamental in all forms of decision-making, since one wishes to make the best decision in any context, and in the analysis of data, where one wishes to find the best model describing experimental data.

This course treats one of the main branches of optimization theory, namely nonlinear programming, which covers a wide range of real life problems including the big data analysis for data analysis and machine learning. It builds on knowledge of linear programming using the simplex algorithm. We will study many general-purpose algorithms, including Brent's method for one-dimensional problems, (quasi-)Newton and conjugate gradient methods for unconstrained problems, and Lagrangian methods, including active-set methods, sequential quadratic programming and interior-point methods for general constrained problems. We will also look at algorithms for global optimization, including branch-and-bound and heuristic approaches such as simulated annealing. Finally, we will cover algorithms for big data, including stochastic gradient descent, and advanced batch optimization methods.

Throughout the course, we aim to provide a coherent framework for the subject, with a focus on optimality conditions (notably the Karush-Kuhn-Tucker conditions), convexity, Lagrange multipliers and duality, and convergence rates. The methods will be illustrated by in-class computer

demonstrations, exercises illustrating the main concepts and algorithms, modelling and computational work on case studies of practical interest in data science and machine learning, including training neural networks.

Prerequisites

Desired Prior Knowledge: Simplex algorithm. Calculus, Linear Algebra.

Recommended reading

J. Nocedal and S.J. Wright, “Numerical Optimization”, Springer, 2006; ISBN: 978-0-387-30303-1.

Léon Bottou, Frank E. Curtis, and Jorge Nocedal. “Optimization methods for large-scale machine learning”, *SIAM Review*, 60(2):223–311, 2018.

Signal and Image Processing

Dept. of Advanced Computing Sciences

KEN4222

Period 1:

1 Sep 2025

24 Oct 2025

Credits:

6.0

Coordinator:

J.M.H. KarelP. Bonizzi

Teaching methods:

Lecture(s), Computer Labs

Assessment methods:

Written exam, Computer test

Full course description

This course offers the student a hands-on introduction into the area of digital signal and image processing. We start with the fundamental concepts and mathematical foundation. This includes a brief review of Fourier analysis, z-transforms and digital filters. Classical filtering from a linear systems perspective is discussed. Next wavelet transforms and principal component analysis are introduced. Wavelets are used to deal with morphological structures in signals. Principal component analysis is used to extract information from high-dimensional datasets. We then discuss Hilbert-Huang Transform to perform detailed time-frequency analysis of signals. Attention is given to a variety of objectives, such as detection, noise removal, compression, prediction, reconstruction and feature extraction. We discuss a few cases from biomedical engineering, for instance involving ECG and EEG signals. The techniques are explained for both 1D and 2D (images) signal processing. The subject matter is clarified through exercises and examples involving various applications. In the practical classes, students will apply the techniques discussed in the lectures using the software package Matlab.

Prerequisites

Desired Prior Knowledge: Linear algebra, Calculus, basic knowledge of Matlab. Some familiarity with linear systems theory and transforms (such as Fourier and Laplace) is helpful.

Recommended reading

Principal Component Analysis, Ian T. Jolliffe, Springer, ISBN13: 978-0387954424.

Control and Intelligent Systems

Dept. of Advanced Computing Sciences

KEN4252

Period 1:

1 Sep 2025

24 Oct 2025

Credits:

6.0

Coordinator:

B. SakçakP.W.L. Dreesen

Teaching methods:

Project-Centered Learning

Assessment methods:

Written exam, Computer test, Take home exam

Full course description

This course explores control theory from classical principles to modern state-of-the-art methods, with applications to robotics, autonomous vehicles, and intelligent systems. Starting from foundational concepts, we explore the control systems that underlie modern engineering and technology. We will start by reviewing relevant models of dynamical systems and analysis and synthesis of feedback control systems. Then, we move the focus to modern state-space control methods, including state feedback, state observers, output tracking, and optimal control. In the last part of the course, we will discuss advanced concepts such as control of uncertain systems, nonlinear control, and data-driven control, as well as the links between control theory and reinforcement learning, and robotics.

Emphasizing digital control and data-driven approaches, the course connects 'Control and Intelligent Systems' to the broader fields of engineering and science, with specific attention to data science and robotics. Hands-on computer sessions and case studies on applications reinforce key concepts. At the end of the course, the students are expected to build a strong theoretical foundation. Furthermore, through selected case studies, the students will gain hands-on experience to design, analyze, and deploy control solutions for complex, real-world problems.

Prerequisites

None

Recommended reading

None

Data Mining

Dept. of Advanced Computing Sciences

KEN4113

Period 1:

1 Sep 2025

24 Oct 2025

Credits:

6.0

Coordinator:

E.N. Smirnov

Teaching methods:

Project-Centered Learning

Assessment methods:

Written exam, Assignment

Full course description

Data mining is a major frontier field of computer science. It allows extracting useful and interesting patterns and knowledge from large data repositories such as databases and the Web. Data mining integrates techniques from the fields of databases, machine learning, statistics, and artificial intelligence. This course will present the state-of-the-art techniques of data mining. The lectures and labs will emphasize the practical use of the presented techniques and the problems of developing real data-mining applications. A step-by-step introduction to data-mining environments will enable the students to achieve specific skills, autonomy, and hands-on experience. A number of real data sets will be analysed and discussed.

Prerequisites

None.

Recommended reading

Pang-Ning, T., Steinbach, M., Karpatne, A., and Kumar, V. (2018). Introduction to Data Mining, 2nd Edition, Pearson, ISBN-10: 0133128903, ISBN-13: 978-0133128901

Quantum Algorithms

Dept. of Advanced Computing Sciences

KEN4235

Period 1:

1 Sep 2025

24 Oct 2025

Credits:

6.0

Coordinator:

G. Stamoulis

Teaching methods:

Project-Centered Learning

Assessment methods:

Written exam, Presentation

Full course description

This course will provide a thorough examination of the most important Quantum Algorithms. We will see how the quantum mechanical formalism gives rise to a new algorithmic design paradigm with the potential of performing certain computational tasks faster than we could do using a classical computer. The course will start with some basic algorithms like Bernstein-Vazirani and Simon's algorithm, then we will move on to Quantum Fourier Transform and Phase Estimation. Then, a thorough discussion of Shor's celebrated algorithm for factoring will follow, together with a detailed coverage of Grover's unstructured search algorithm, its optimality, adaptations, and applications. Further, we will move on to the HHL algorithm for solving systems of linear equations, a crucial component of many quantum algorithms, including Machine Learning quantum algorithms. In the last part of the course, we will present algorithms for quantum simulation, discuss quantum walks, and basics of quantum complexity theory by introducing and discussing the BQP and QMA complexity classes.

Prerequisites

Desired prior knowledge: Fundamentals of Quantum Computation, Very Good command of Linear Algebra, Algorithms and Complexity

Prerequisites: Introduction to Quantum Computing for AI and Data Science

Recommended reading

Recommended literature: Quantum Computation and Quantum Information: 10th Anniversary Edition Anniversary Edition, Michael A. Nielsen, and Isaac L. Chuang

Additional literature: Papers, notes and other relevant material will be distributed in class.

Model Identification and Data Fitting

Dept. of Advanced Computing Sciences

KEN4242

Period 2:

27 Oct 2025

12 Dec 2025

Credits:

6.0

Coordinator:

R.L.M. PeetersP.W.L. Dreesen

Teaching methods:

Project-Centered Learning

Assessment methods:

Written exam

Full course description

This course is devoted to the various practical and theoretical aspects which involve the estimation (the identification) of a mathematical model within a given model class, starting from a record of observed measurement data (input-output data). First, we address distance measures, norms, and criterion functions. Then we discuss the prediction error identification of linear regression models, with special emphasis on the various interpretations of such models (deterministic, stochastic with Gaussian white noise and maximum likelihood estimation, stochastic in a Bayesian estimation context) and on numerical implementation aspects (recursion, numerical complexity, numerical conditioning and square root filtering). Next, we study identification within the important class of auto-regressive dynamical models, to which the Levinson algorithm applies. Other related topics receiving attention are identifiability, model reduction and model approximation. Some techniques for the estimation of linear dynamical i/o-systems are illustrated with the system identification toolbox in Matlab.

Prerequisites

Linear Algebra, Mathematical Modelling, Probability and Statistics.

Recommended reading

- L. Ljung, System Identification: Theory for the User (2nd ed.), Prentice-Hall, 1999.
- T. Soderstrom and P. Stoica, System Identification, Prentice-Hall, 1989.

Algorithmic Game Theory

Dept. of Advanced Computing Sciences

KEN4251

Period 2:

27 Oct 2025

12 Dec 2025

Credits:

6.0

Coordinator:

M. Salvioli

Teaching methods:

Project-Centered Learning

Assessment methods:

Written exam

Full course description

The course will focus on non-cooperative games and on dynamic games in the following order: matrix and bimatrix games, repeated games, Stackelberg games, differential games, specific models of stochastic games, evolutionary games. These are games in which the players are acting as strategic decision makers, who cannot make binding agreements to achieve their goals. Instead, threats may be applied to establish stable outcomes. Besides, relations with population dynamics and with “learning” will be examined. Several examples will be taken from biological settings.

Prerequisites

Desired Prior Knowledge: Students are expected to be familiar with basic concepts from linear algebra, calculus, Markov chains and differential equations.

Recommended reading

None.

Advanced Natural Language Processing

Dept. of Advanced Computing Sciences

KEN4259

Period 2:

27 Oct 2025

12 Dec 2025

Credits:

6.0

Coordinator:

A.S. Härmä J.C. Scholtes

Teaching methods:

Project-Centered Learning

Assessment methods:

Written exam, Assignment

Full course description

For decades, teaching a computer to deal with natural language processing (NLP) was a long-time dream of humankind. Tasks such as machine translation, summarization, question-answering, speech recognition or chatting remained a challenge for computer programs. Around 2018, major improvements were made. Starting with machine translation and ultimately in late 2022 with ChatGPT. Why were these large-language models suddenly so good? How did we get here? What can we do with these new algorithms to improve them even more?

This course will provide the skills and knowledge to understand and develop state-of-the-art (SOTA) solutions for these natural language processing (NLP) tasks. After a short introduction to traditional generative grammars and statistical approaches to NLP, the course will focus on deep learning techniques. We will discuss Transformers, variations on their architecture (including BERT, GPT and multi-modal) in depth, which models work best for which linguistic tasks, their capacities, limitations and how to address these.

Although computer systems can now use Natural Language in ways that can no longer be distinguished from humans, there are still major problems to address: (i) we do not fully understand what these algorithms know and what they do not know. So, there is a strong need for explainable AI (XAI) in NLP. (ii) Training the deep-learning large language-models costs too much energy. In the lectures we will discuss models and methods that are less computationally (and thus energy)

intensive. (iii) New methods aim to go from having a simple conversation to reasoning and solving complex problems using agentic techniques and reinforcement learning. In this course, we will also discuss these topics.

This course is closely related with the course Information Retrieval and Text-Mining (IRTM). The ANLP course is more focussed on the inner workings of the methods and architectures to deal with complex natural language tasks. The IRTM course focusses more on building search engines, using text-analytics for deeper understanding and the requirements needed for responsible conversational search. The overlap between the two courses is kept to a minimum. There is no need to follow the courses in a specific order.

Prerequisites

Proficient coding skills in Python are required to participate in the tutorials.

Google CoLAB pro subscription for the tutorials.

Recommended reading

Papers published in top international conferences and journals in machine learning field.

Quantum Information and Security

Dept. of Advanced Computing Sciences

KEN4237

Period 2:

27 Oct 2025

12 Dec 2025

Credits:

6.0

Coordinator:

D.O. Mestel

Teaching methods:

Project-Centered Learning

Assessment methods:

Written exam, Presentation

Full course description

In this course we will consider the power of quantum mechanics not in accomplishing computational or ‘algorithmic’ tasks, but instead for communication- and security-related tasks. The strange properties of the quantum world turn out to be remarkably useful for these. For example, we can exchange secret messages in a way that is unconditionally secure: secrecy is guaranteed by the physical laws of nature, rather than (as in ordinary cryptography) based on an assumption that a particular computational problem is too hard for the adversary.

We will begin by covering the theoretical techniques needed to study security-related protocols, where it is fundamental that some parties will not know what state a particular quantum system is in. After a thorough grounding in the ‘density matrix’ formalism which is used to represent this uncertainty, we will cover quantitative measures of this kind of uncertainty, for instance quantum versions of classical entropy. We will then look at a variety of protocols (e.g. quantum money, quantum key distribution,...), and how to define and prove the desired properties.

Prerequisites

Desired prior knowledge: Quantum states, operators and measurements

Prerequisites: Introduction to Quantum Computing for AI and Data Science

Recommended reading

Recommended literature: T. Vidick and S. Wehner, 'Introduction to quantum cryptography', Cambridge University Press 2024

Additional literature: M. Wilde, 'Quantum information theory', Cambridge University Press 2013

Network Science

Dept. of Advanced Computing Sciences

KEN4275

Period 2:

27 Oct 2025

12 Dec 2025

Credits:

6.0

Coordinator:

A.I. Iamnitchi

Teaching methods:

Assignment(s), Work in subgroups, Project-Centered Learning

Assessment methods:

Assignment

Full course description

Many aspects of everyday life and science can be represented as networks: social networks represent relationship (links) between people (nodes); brain activity can be represented via synapses (links) between neurons (nodes); the street map is formed of roads (links) that connect at intersections (nodes); authors of scientific papers connect to each others in a citation network, with directed links from the paper cited to the paper citing it; communication networks connect routers via physical or logical links; etc. Network analysis plays a significant role in the “big data” analytics because of size, data velocity, or computational complexity.

This course focuses on the study of network structures and dynamic processes on networks using real data from various disciplines, including socio-technological platforms, biology, social science, and economics. Topics cover the analysis and modeling of complex networks, network dynamics, community detection, network resilience and contagion, as well as processing of network structures for machine learning tasks.

Prerequisites

Desired prior knowledge: Introductory knowledge of programming for data analysis, particularly in Python; algorithms; algorithmic complexity.

Prerequisites: Introductory courses in algorithms, data structures, and data analytics.

Recommended reading

"Networks, Crowds, and Markets: Reasoning About a Highly Connected World" by David Easley and Jon Kleinberg.

Additional literature: Graph Theory and Complex Networks: An Introduction by Maarten van Steen

Data Visualization

Dept. of Advanced Computing Sciences

KEN4224

Period 2:

27 Oct 2025

12 Dec 2025

Credits:

6.0

Coordinator:

M.F.M. Sondag

Teaching methods:

Project-Centered Learning

Assessment methods:

Written exam, Assignment

Full course description

When data becomes sufficiently large and complex that going through it manually becomes difficult, we can make use a combination of automated analysis techniques with data visualization to better understand and analyze the underlying data, and make decisions based on these insights obtained from this analysis.

In this course, we provide a deeper understanding of data visualization, as well as how to effectively combine it with automated algorithmic approaches. Starting from the fundamentals of data visualizations, we cover how we can map the human visual perception into visual channels, which can be used a building blocks for new data visualizations. Through the lens of these channels, the course provides a comprehensive exploration of visualization techniques analyzing how to use these techniques effectively such that they do not mislead the viewer for different data-types such as geo-spatial, time-varying, and network data. We also cover how to integrate automated analysis techniques such as clustering and degree-of-interest functions through interactive visualizations, focusing on the connection between the algorithm and the visualization.

Prerequisites

Desired prior knowledge : Basic programming skills, Knowledge about data analytics.

Recommended reading

Tamara Munzner: "Visualization Analysis and Design", *A K Peters Visualization Series*, CRC Press,
ISBN: 9781466508910, 1st edition

Quantum AI

Dept. of Advanced Computing Sciences

KEN4236

Period 2:

27 Oct 2025

12 Dec 2025

Credits:

6.0

Coordinator:

D. Dibenedetto

Teaching methods:

Project-Centered Learning

Assessment methods:

Written exam, Assignment

Full course description

This course explores the groundbreaking intersection of quantum computing and artificial intelligence, focusing on how quantum technologies can potentially revolutionize AI paradigms. The curriculum delves into quantum algorithms tailored for AI tasks, addressing complex problems that are currently intractable for classical computers. Students will gain an understanding of how quantum principles can enhance machine learning algorithms, improve optimization tasks, and facilitate data processing capabilities. Through theoretical lessons and practical laboratory sessions, students will learn about quantum mechanics fundamentals applicable to AI, quantum circuit design, and quantum algorithm development. Special emphasis will be placed on hybrid models that integrate classical and quantum computing techniques to solve real-world problems. The course will provide a mix of both theoretical and technical insights, as well as practical implementation details by using the main quantum programming languages and quantum software available.

Prerequisites

Desired prior knowledge: Linear Algebra, Classical Machine Learning

Prerequisites: Introduction to Quantum Computing for AI and Data Science

Recommended reading

Recommended literature: "Machine Learning with Quantum Computers" by M. Schuld, F. Petruccione, Second Edition

Additional literature: Research articles and papers will be provided throughout the course.

Planning and Scheduling

Dept. of Advanced Computing Sciences

KEN4253

Period 4:

26 Jan 2026

27 Mar 2026

Credits:

6.0

Coordinator:

S.M. Kelk

Teaching methods:

Project-Centered Learning

Assessment methods:

Written exam, Assessment

Full course description

In many real-world processes, particularly in industrial processes and logistics, decisions need to be taken about the time of the completion of (sub)tasks, and the decision about what production machines complete the tasks. There are often constraints on the order in which tasks, or 'jobs' can be performed, and there are usually capacity constraints of the machines. This leads to natural, industrially critical optimization problems. For example, a company might choose to buy many machines to process jobs, but then there is a risk that the machines will be underused, which is economically inefficient. On the other hand, too few machines, or an inappropriate ordering of tasks, may lead to machines spending a significant amount of time standing idle, waiting for the output of other machines, which are overcrowded with tasks. In this course, we look at various mathematical models and techniques for optimizing planning and scheduling problems, subject to different optimality criteria. We will discuss, among others, single-machine models, parallel-machine models, job-shop models, and algorithms for planning and scheduling (exact, approximate, heuristic) and we also touch upon the computational complexity (distinguishing between 'easy' and 'difficult' problems) of the underlying problems. Last but not least, we will also introduce integer linear programming as a uniform and generic tool to model and solve planning and scheduling problems.

Prerequisites

None.

Desired prior knowledge: Data Structures & Algorithms. Discrete Mathematics. Graph Theory

Recommended reading

None.

Building and Mining Knowledge Graphs

Dept. of Advanced Computing Sciences

KEN4256

Period 4:

26 Jan 2026

27 Mar 2026

Credits:

6.0

Coordinator:

Y. YangM.J. Dumontier

Teaching methods:

Project-Centered Learning

Assessment methods:

Assignment, Assessment

Full course description

Knowledge graphs are large-scale, machine-processable representations of entities, their attributes, and their relationships. Knowledge graphs enable both people and machines to explore, understand, and reuse information in a wide variety of applications such as answering questions, finding relevant content, understanding social structures, and making scientific discoveries. However, the sheer size and complexity of these graphs present a formidable challenge particularly when mining across different topic areas.

In this course, we will examine approaches to construct and use knowledge graphs across a diverse set of applications using cutting-edge technologies such as machine learning and deep learning, graph databases, ontologies and automated reasoning, and other relevant techniques in the area of data mining and knowledge representation.

Prerequisites

Desired Prior Knowledge: Introduction to Computer Science

Recommended reading

Aggarwal, C.C. and Wang, H. eds., (2010) Managing and mining graph data (Vol. 40). New York: Springer. ISBN 978-1-4419-6045-0

Computational Statistics

Dept. of Advanced Computing Sciences

KEN4258

Period 4:

26 Jan 2026

27 Mar 2026

Credits:

6.0

Coordinator:

A. WodeyarC.J. Seiler

Teaching methods:

Project-Centered Learning

Assessment methods:

Written exam, Assignment

Full course description

In this course, we will review basic concepts in statistical inference (confidence intervals, parameter estimation, and hypothesis testing). We will then study computer-intensive methods that work without imposing unrealistic or unverifiable assumptions about the data generating mechanism (randomization tests, the bootstrap, and Markov chain Monte Carlo). This will provide us with the foundations to study modern inference problems in statistics and machine learning (false discovery rates, Benjamini-Hochberg procedure, and causal inference).

Prerequisites

None.

Desired prior knowledge: Probability and Statistics

Recommended reading

None.

Data Fusion

Dept. of Advanced Computing Sciences

KEN4223

Period 4:

26 Jan 2026

27 Mar 2026

Credits:

6.0

Coordinator:

A.M. Wilbik

Teaching methods:

Project-Centered Learning

Assessment methods:

Written exam, Assignment

Full course description

ICT development, e.g., remote sensing, IoT, lead to an enormous growth of available data for analysis. To integrate this heterogeneous or multimodal data, data fusion approaches are used. Data fusion can be understood as a framework for the joint analysis of data from multiple sources (modalities) that allows achieving information/knowledge not recoverable by the individual ones.

During this course, several approaches to data fusion will be discussed, such as:

1. Low level data fusion, where data fusion methods are directly applied to raw data sets for exploratory or predictive purposes. A main advantage is the possibility to interpret the results directly in terms of the original variables. An example of a low level data fusion is measuring the same signal or phenomena with different sensors, in order to discover the original one. Traditionally, PCA based methods are used for this type of data fusion.
2. Mid level data fusion, where data fusion operates on features extracted from each data set. The obtained features are then fused in a “new” data set, which is modeled to produce the desired outcome. A main advantage is that the variance can be removed in the features extraction step, and thus the final models may show better performance. An example of a mid level data fusion is extracting numerical features from an image, and building a decision model based on those features.
3. High level data fusion, also known as decision fusion, where decisions (models outcome) from processing of each data set are fused. It is used when the main objective is to improve the performance of the final model and reach an automatic decision. Several methods can be used for high-level DF, such as weighted decision methods, Bayesian inference, Dempstere Shafer’s

theory of evidence, and fuzzy set theory. There is a link between high-level data fusion and ensemble methods.

4. Federated learning. Federated learning enables multiple parties jointly train a machine learning model without exchanging the local data. In case of federated learning, we can talk about model fusion.

Prerequisites

None.

Desired prior knowledge: statistics and basic machine learning

Recommended reading

None.

Explainable AI

Dept. of Advanced Computing Sciences

KEN4246

Period 4:

26 Jan 2026

27 Mar 2026

Credits:

6.0

Coordinator:

M.F.M. SondagN. Tintarev

Teaching methods:

Project-Centered Learning

Assessment methods:

Group Assignment, Written Exam (Literature Review)

Full course description

A key component of an artificially intelligent system is the ability to explain to a human agent the decisions, recommendations, predictions, or actions made by it and the process through which they are made. Such explainable artificial intelligence (XAI) can be required in a wide range of applications. For example, a regulator of waterways may use a decision support system to decide which boats to check for legal infringements, a concerned citizen might use a system to find reliable information about a new disease, or an employer might use an artificial advice-giver to choose between potential candidates fairly.

For explanations from intelligent systems to be useful, they need to be able to justify the advice they give in a human-understandable way. This creates a necessity for techniques for automatic generation of satisfactory explanations that are intelligible for users interacting with the system. This interpretation goes beyond a literal explanation. Further, understanding is rarely an end-goal in itself. Pragmatically, it is more useful to operationalize the effectiveness of explanations in terms of a specific notion of usefulness or explanatory goals such as improved decision support or user trust. One aspect of intelligibility of an explainable system (often cited for domains such as health) is the ability for users to accurately identify, or correct, an error made by the system. In that case it may be preferable to generate explanations that induce appropriate levels of reliance (in contrast to over- or under-reliance), supporting the user in discarding advice when the system is incorrect, but also accepting correct advice.

The following subjects will be discussed:

- (1) Intrinsically interpretable models, e.g., decision trees, decision rules, linear regression.
- (2) Identification of violations of assumptions; such as distribution of features, feature interaction, non-linear relationships between features; and what to do about them.
- (3) Model agnostic explanations, e.g., LIME, scoped Rules (Anchors), SHAP (and Shapley values)
- (4) Ethics for explanations, e.g., fairness and bias in data, models, and outputs.
- (5) (Adaptive) User Interfaces for explainable AI
- (6) Evaluation of explanation understandability

Prerequisites

Desired Prior Knowledge: Data Analysis and Data Mining

Maximum number of 60 students can follow this course from MSc AI or MSc MSDM.

Mandatory for MSc RDS (no cap).

Late enrollment for this course is not permitted, as participation in mandatory assignments and group work is required from the start.

Recommended reading

- Molnar, Christoph. Interpretable Machine Learning (**3rd edition**) . Lulu. com, **2025**.
- Rothman, Denis. Hands-On Explainable AI (XAI) with Python: Interpret, visualize, explain, and integrate reliable AI for fair, secure, and trustworthy AI apps, Packt, 2020.

Data Privacy and Security

Dept. of Advanced Computing Sciences

KEN4225

Period 4:

26 Jan 2026

27 Mar 2026

Credits:

6.0

Coordinator:

T. Schnitzler

Teaching methods:

Project-Centered Learning

Assessment methods:

Participation, Assignment

Full course description

This course explores the intersection of data privacy, anonymity, and security in the modern digital landscape. Students will gain a deep understanding of how personal data is collected, shared, and sometimes leaks from online communication through platforms like social media, mobile apps, and web services.

The course addresses personal information exposure by means of voluntary and involuntary data disclosure, including real-world threats to invade privacy, such as traffic analysis to match user identities and activities on the Internet. Students will also be introduced to different levels of privacy-enhancing technologies including statistical disclosure control, anonymization, and differential privacy but also network- and application-level PETs.

Through working with case studies and a selection of state-of-the-art research addressing both offensive and defensive work in the context of privacy and anonymity, students will develop critical skills to understand, assess, and mitigate privacy risks in data-driven applications.

Prerequisites

Desired prior knowledge: Databases or Data Management, Probability and Statistics, Computer Networks

Recommended reading

Helen Nissenbaum: Privacy in Context. Stanford University Press, 2010. (Digital copies available through UM library)

Additional literature:

Research articles and papers will be provided throughout the course.

Algorithms for Big Data

Dept. of Advanced Computing Sciences

KEN4254

Period 5:

30 Mar 2026

22 May 2026

Credits:

6.0

Coordinator:

M. Mihalak

Teaching methods:

Project-Centered Learning

Assessment methods:

Written exam, Assignment

Full course description

The emergence of very large datasets poses new challenges for the algorithm designer. For example, the data may not fit into the main memory anymore, and caching from a hard-drive becomes a new bottleneck that needs to be addressed. Similarly, algorithms with larger than linear running time take simply too long on very large datasets. Moreover, simple sensory devices can observe large amount of data over time, but cannot store all the observed information due to insufficient storage, and an immediate decision of what to store and compute needs to be made. Classical algorithmic techniques do not address these challenges, and a new algorithmic toolkit needs to be developed. In this course, we will look at a number of algorithmic responses to these problems, such as: algorithms with (sub-)linear running times, algorithms where the data arrive as a stream, computational models where memory is organized hierarchically (with larger storage units, such as hard-drives, being slower to access than smaller, faster storage such as CPU cache memory). New programming paradigms and models such as MapReduce/Hadoop will be discussed. We will also look at a number of topics from classical algorithm design that have undiminished relevance in the era of big data such as approximation algorithms and multivariate algorithmic analysis.

Prerequisites

Desired prior knowledge: Discrete mathematics, algorithm design and analysis, probability theory

Recommended reading

None.

Information Retrieval and Text Mining

Dept. of Advanced Computing Sciences

KEN4153

Period 5:

30 Mar 2026

22 May 2026

Credits:

6.0

Coordinator:

J.C. Scholtes

Teaching methods:

Project-Centered Learning

Assessment methods:

Written exam, Assessment

Full course description

Information Retrieval and Text Mining aka Conversational Search Algorithms

This course is all about algorithms for Conversational Search

Traditional search (aka Information Retrieval) is about “Finding the needle in the haystack”. This course focusses on more complex problems: “How does the needle look like and where is the haystack?” using text-mining, topic modeling and data visualization techniques. Today, search is more about, “Having a guide handing you the needle and explaining you its significance and context” using conversational AI in combination with search engines and text-mining.

Building a full-text search engine may look trivial, but it is not! How do you search hundreds of billions of documents that can be located anywhere, with sub-second responds times? How do you find exactly what you are looking for without missing relevant information or having to plough through hundreds of irrelevant documents? How can you find if you do not know exactly what you are looking for? How can you find information which is deliberately hidden? How do you know that your search engine has given you the right information? Where does it come from? Is the answer factually correct?

In this course, we will teach you how to address these questions in three steps: (1) how is a search engine is constructed, optimized and used effectively, (2) How can techniques from the world of text-

mining, information extraction, text classification, clustering, topic modelling and data visualization lead to better search, and (3) What is the best way to integrate chatbots with search engines resulting in responsible conversational search. How to best guarantee factuality, measure and prevent hallucinations, provide provenance and explainability of the chatbots' recommendations. How can we integrate knowledge graphs and retrieval augmented generation (RAG) in the conversation? How can agentic architectures help us to overcome the limitations of LLMs?

Interestingly, traditional search-engine and text-mining techniques are essential components of such responsible conversational search solutions. In this course we will discuss all components and how they best work together.

Linear Algebra, Statistics, Deep Learning and Natural Language Processing play an important role in this course.

This course is complementary to the course Advanced Natural Language Processing (ANLP). Overlap is reduced to the necessary minimum. Both courses can be followed in any particular order. In the Information Retrieval and Text Mining course we focus more on creating an optimal modern search experience, in the Advanced Natural Language Processing course, we do a deep dive into the algorithms and models used for different language-related problems such as machine translation, abstracting, dialogs with chatbots, and LLM-based agents. Tutorials are shared between the two courses.

Prerequisites

Proficient coding skills in Python are required to participate in the tutorials.

Google CoLAB pro subscription for the tutorials.

Recommended reading

Introduction to Information Retrieval. Christopher D. Manning, Prabhakar Raghavan and Hinrich Schütze. Cambridge University Press, 2008. In bookstore and online: <http://informationretrieval.org> .

Computer Vision

Dept. of Advanced Computing Sciences

KEN4255

Period 5:

30 Mar 2026

22 May 2026

Credits:

6.0

Coordinator:

M.C. Popa

Teaching methods:

Project-Centered Learning

Assessment methods:

Written exam, Assignment

Full course description

Can we make machines look, understand and interpret the world around them? Can we make cars that can autonomously navigate in the world, robots that can recognize and grasp objects and, ultimately, recognize humans and communicate with them? This course will provide the knowledge and skills that are fundamental to core vision tasks of one of the fastest growing fields in academia and industry: visual computing. Topics include introduction to fundamental problems of computer vision, mathematical models and computational methodologies for their solution including both traditional and newer techniques such as Visual Language Models (VLMs), implementation of real-life applications and experimentation with various techniques in the field of scene analysis and understanding. In particular, after a recap of basic image analysis tools (enhancement, restoration, color spaces, edge detection), students will learn about feature detectors and trackers, fitting, image geometric transformation and mosaicing techniques, texture analysis and classification using unsupervised techniques, face and emotion analysis, deep learning based object classification using both CNNs and transformers, detection and tracking, camera models, epipolar geometry and 3D reconstruction from 2D views.

Prerequisites

None.

Desired prior knowledge: Basic knowledge of Python, linear algebra and machine learning. This course offers the basics on image processing although prior knowledge is also a plus.

Recommended reading

“Digital Image Processing”, Rafael C. Gonzalez & Richard E. Woods, Addison-Wesley, “Computer Vision: Models, Learning and Inference”, Simon J.D. Prince 2012.

Introduction to Quantum Computing for AI and Data Science

Dept. of Advanced Computing Sciences

KEN4155

Period 5:

30 Mar 2026

22 May 2026

Credits:

6.0

Coordinator:

D.O. MestelG. Stamoulis

Teaching methods:

Project-Centered Learning

Assessment methods:

Written exam

Full course description

In this course we lay down the foundations and basic concepts of quantum computing. We will use the mathematical formalism borrowed from quantum mechanics to describe quantum systems and their interactions. We introduce the concept of a quantum bit and discuss different physical realizations of it. We then introduce the basic building blocks of quantum computing: quantum measurements and quantum circuits, single and multi-qubit gates, the difference between correlated (entangled) and uncorrelated states and their representation, quantum communication, and basic quantum protocols and quantum algorithms. Finally, we discuss the different types of noise involved in real quantum computers (coherent and incoherent errors, state preparation, projection and measurement) and their effect on performance, and outline current efforts for mitigating the issues.

This course is a prerequisite for the elective courses Quantum Algorithms, Quantum AI, and Quantum Information and Security. These four courses, together with a dedicated research project quantum computing forms the specialization Quantum Computing.

Prerequisites

Desired prior knowledge: probability theory, linear algebra, design and analysis of algorithms.

This course is a prerequisite for the elective courses "Quantum Algorithms", "Quantum AI", "Quantum Information and Security" and the project "Research Project Quantum Computing".

These four courses, together with a dedicated research project on Quantum Computing forms the specialization in Quantum Computing for AI and Data Science.

Recommended reading

"Quantum Computation and Quantum Information" Michael Nielsen and Isaac Chuang, Cambridge University Press, 10th Anniversary edition.

Reinforcement Learning

Dept. of Advanced Computing Sciences

KEN4157

Period 5:

30 Mar 2026

22 May 2026

Credits:

6.0

Coordinator:

D.J.N.J. SoemersK. Driessens

Teaching methods:

Project-Centered Learning

Assessment methods:

Assignment, Take home exam

Full course description

Reinforcement learning is a type of machine learning problem in which the learner gets a (delayed) numerical feedback signal about its demonstrated performance. It is the toughest type of machine learning problem to solve, but also the one that best encompasses the idea of artificial intelligence as a whole. In this course we will define the components that make up a reinforcement learning problem and will see what the important concepts are when trying to solve such a problem, such as state and action values, policies and performance feedback. We will look at the different properties a reinforcement learning problem can have and what the consequences of these properties are with respect to solvability. We will discuss value based techniques as well as direct policy learning and learn how to implement these techniques. We will study the influence of generalisation on learning performance and see how supervised learning (and specifically deep learning) can be used to help reinforcement learning techniques tackle larger problems. We will also look at the evaluation of learned policies and the development of performance over time.

Prerequisites

No hard prerequisites but having some background in Machine Learning and/or Data Mining will be helpful.

Recommended reading

Lecture slides will be uploaded before each lecture. These slides are designed and intended as support during teaching, not as study material by themselves. They are supplied as a service, but additional note taking will be necessary to pass the class.

The book “Reinforcement Learning – An Introduction” by Sutton and Barto is freely available at: <https://www.andrew.cmu.edu/course/10-703/textbook/BartoSutton.pdf>

Master's Thesis AI

Dept. of Advanced Computing Sciences

KEN4160

Year:

1 Sep 2025

19 Jun 2026

Credits:

30.0

Coordinator:

M. Mihalak

Teaching methods:

Paper(s)

Assessment methods:

Presentation and paper

Full course description

The Master Artificial Intelligence will be completed by writing a master's thesis.

The thesis is produced individually and is the result of a master's research project that runs during the second semester of year 2 of the master's programme.

In the first phase, the emphasis is on self-study, subject determination, planning and some preliminary research. Then the actual research is started.

The final phase is used to finalize the master's thesis.

The master's project is completed by a presentation of the results.

The master's project will be supervised by one of the senior researchers.

Prerequisites

In order to start working on the thesis, a student needs to have obtained at least 70 ECTS (among which are 40 credits of the first year).

Recommended reading

None.

