# Take a unique specialisation in Quantum Computing at Maastricht University

**Quantum computing is a much-anticipated new technology that is still in its infancy. “As soon as quantum hardware becomes widely available, quantum computing has the potential to revolutionise research and technology”, says Georgios Stamoulis, assistant professor at the Department of Advanced Computing Sciences at Maastricht University. He will lecture about topics of Quantum Computing.**

**The specialisation in Quantum Computing…**
• Prepares you to use this future technology by knowing its limitations and possibilities when applied to Data Science and Artificial Intelligence
• Teaches you which computational problems can be tackled in the framework of quantum mechanics
• Focuses on the underlying theory, algorithms and possible applications, not on quantum hardware
• Is offered as part of master’s programmes in Data Science for Decision Making and Artificial Intelligence at Maastricht University only

> As soon as quantum hardware becomes widely available, quantum computing has the potential to revolutionise research and technology

Georgios Stamoulis - assistant professor at the Department of Advanced Computing Sciences at Maastricht University.

# Curriculum

As of the academic year 2023-2024, Maastricht University offers a specialisation in Quantum Computing. The specialisation is accessible for students of the master’s programmes in Artificial Intelligence and Data Science for Decision Making. During the last semester of the first year, you can opt to follow an introductory course. In year 2, you specialise further by taking more courses, and top it off with a research project on quantum computing.

If you complete the required courses and the corresponding research project, the specialisation in Quantum Computing will be listed on your master’s diploma. It’s also possible to follow some of the specialisation’s courses as electives.

- Year 1 - Period 5: Introduction to Quantum Computing for AI & Data Science
- Year 2 - Period 1: Quantum Algorithms
- Year 2 - Period 2: Quantum Artificial Intelligence | Quantum Information & Security
- Year 2 - Periods 1,2,3: Group Research Project on Quantum Computing

## Be prepared

As soon as the hardware becomes available, he says. So, why should we already study quantum computing now? “Well, we better be prepared”, Stamoulis answers. Quantum Delta NL, the Dutch quantum ecosystem, therefore started working with Dutch universities to begin educating specialists in quantum computing. The demand for specialists in this emerging field is already high.

"Companies and research groups ask for it", says Matúš Mihalák, programme director of the master's studies that will embed the new specialisation programme. "They want someone in their team who can advise them on the possibilities of quantum computing. Just think of financial institutions that want to know about quantum-proof encryption or scientists who research quantum computing."

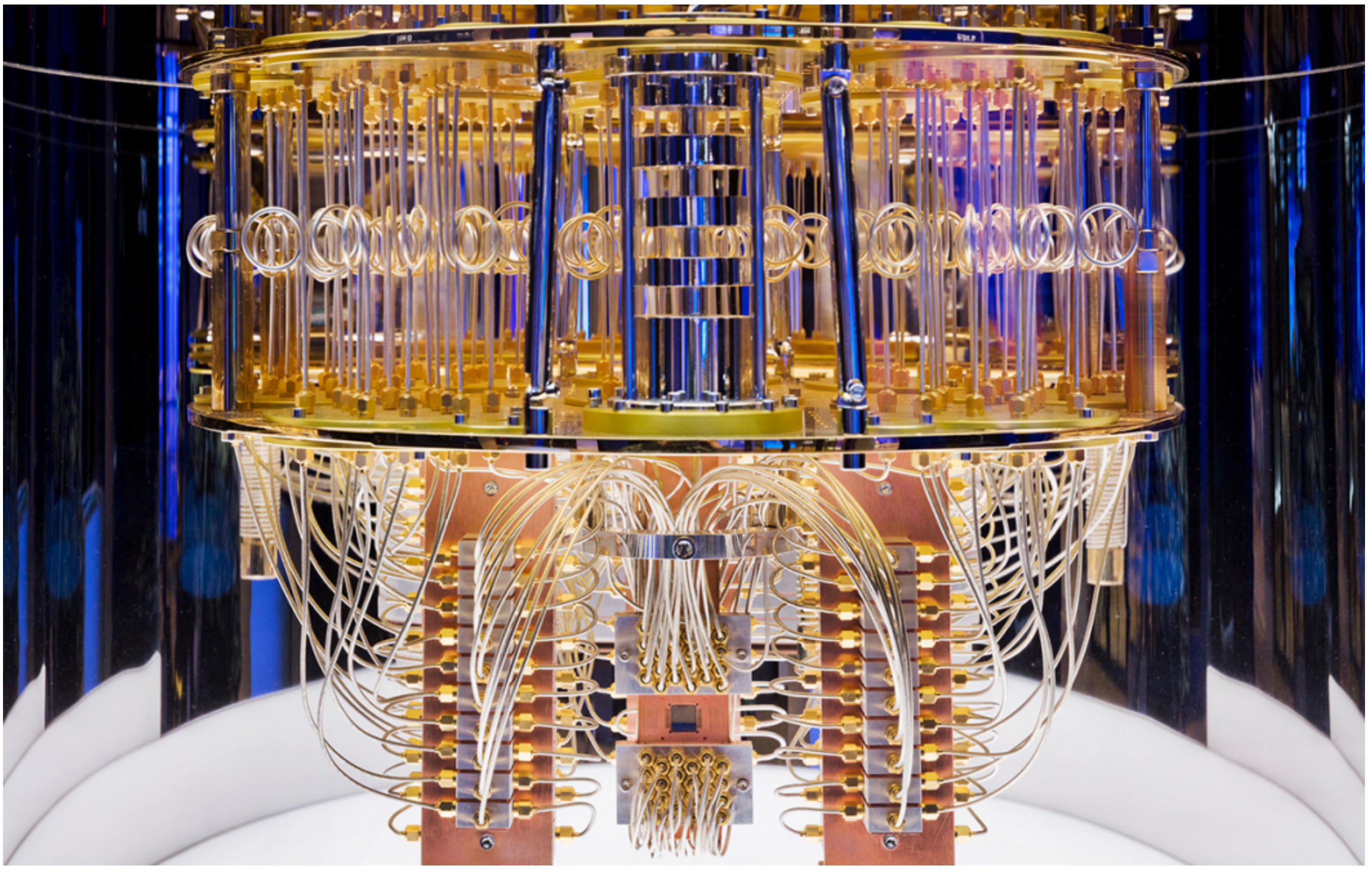

## Linear algebra

The specialisation in Quantum Computing will introduce you to realistic as well as theoretical possibilities of quantum computing. The technology relies on the behaviour of the smallest particles that make up our world, as defined by the theory of quantum mechanics. Quantum mechanics is a linear theory described by linear algebra. "With quantum computing you can theoretically work on problems that fit into the mathematical structure of quantum mechanics", states Stamoulis. "It will be our challenge to find linear problems, or parts of non-linear problems, like curing cancer or solving the climate problem, that can be in reach of quantum computing technology."

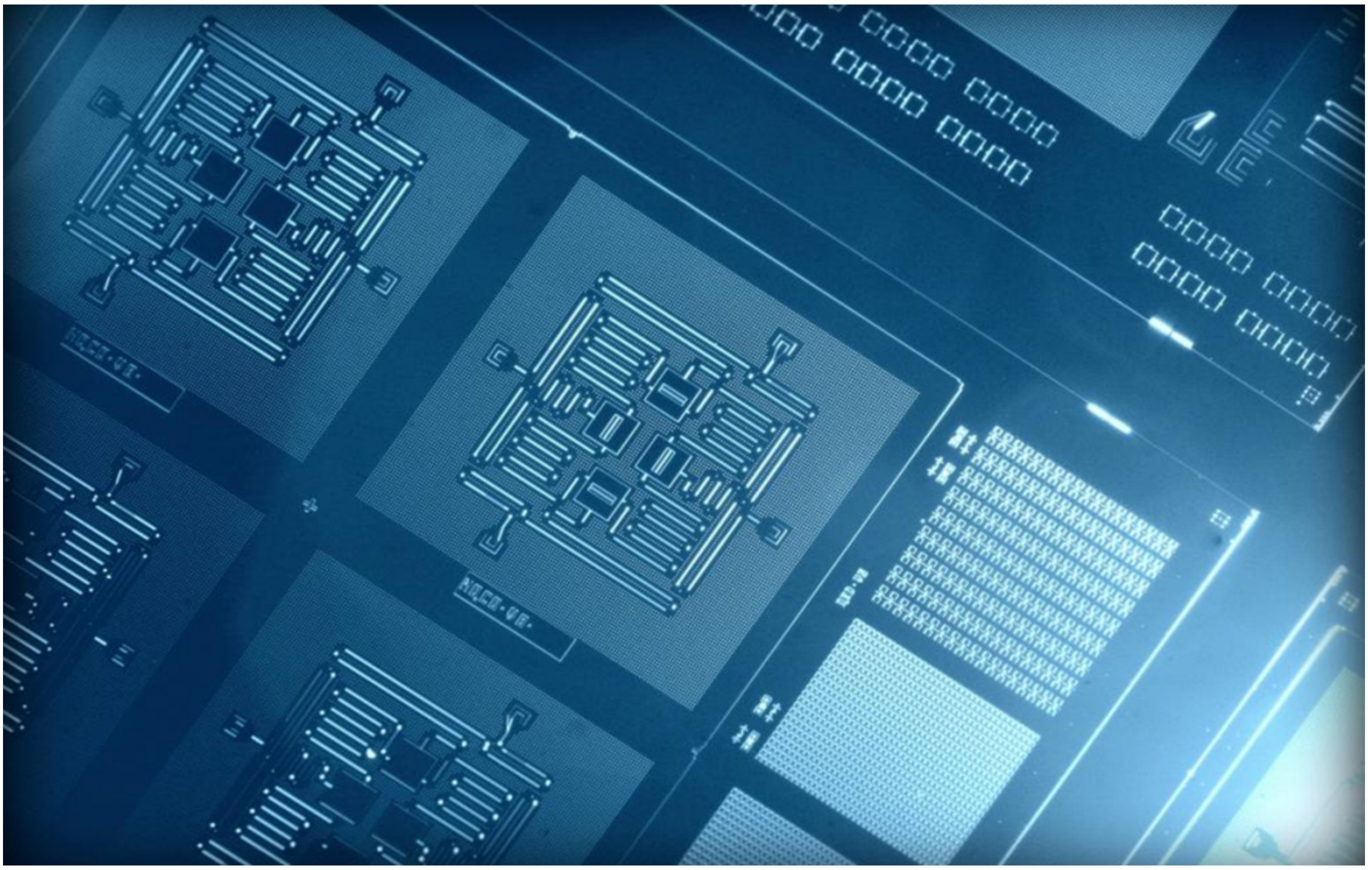

## Examples

Most examples of potentially solvable problems are in the field of quantum mechanics itself, like in quantum chemistry. Another, notorious, case is cryptography using prime factorisation (describing a number as the product of two prime numbers). Quantum computers, when available, can heavily exploit prime factorisation, whereas classical computers have a very hard time finding the primes of large numbers efficiently. Don’t worry: the money in your bank account is safe for now. The current best experimental quantum computer can factor 21 into its primes 7 and 3 (using Shor's algorithm).

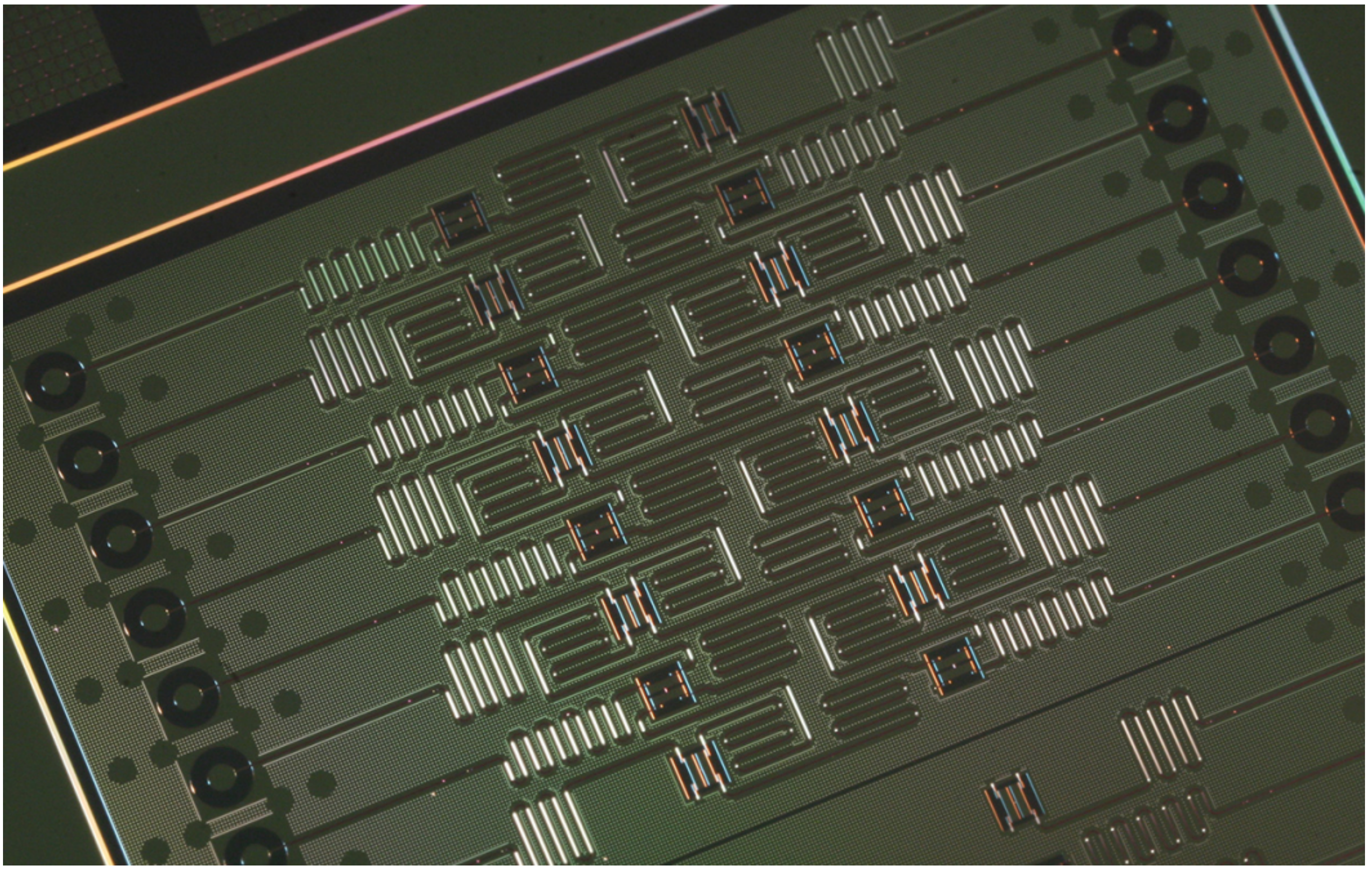

## Courses

The specialisation in Quantum Computing is accessible to master’s students of the Artificial Intelligence and Data Science for Decision Making programmes. It offers four courses. You will familiarise yourself with designing quantum algorithms and see how you might use them in quantum artificial intelligence, machine learning, cryptology and security. You will do a project on quantum computing as well. It’s also possible to follow some of the specialisation’s courses as electives.

*All pictures by IBM Research and are licensed under CC BY-ND 2.0.*

First year courses

# DSDM Core Courses

# Data Mining

### 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

KEN4113
Period 1
4 Sep 2023
27 Oct 2023
[Print course description]
ECTS credits:
6.0
Instruction language:
English
Coordinator:

- [E.N. Smirnov]

Teaching methods:
Project-Centered Learning
Assessment methods:
Written exam, Assignment

## Model Identification and Data Fitting

### 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.

KEN4242

Data Science for Decision Making

Period 2
30 Oct 2023
22 Dec 2023
Print course description
ECTS credits:
6.0
Instruction language:
English
Coordinators:

- R.L.M. Peeters
- P.W.L. Dreesen

Teaching methods:
Project-Centered Learning
Assessment methods:
Written exam

# Research Project Data Science for Decision Making 1

## 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 period, the goal is to start modelling, prototyping and developing. In period 3, the implementation, model and/or experiments set out in the project plan has to be finished and reported on. At the end of period 1 and 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

Justin Zobel (2004), Writing for Computer Science, Springer, ISBN:1852338024

KEN4230
Semester 1
4 Sep 2023
2 Feb 2024
Print course description
ECTS credits:
6.0

Data Science for Decision Making

Instruction language:
English
Coordinators:

- Linda Rieswijk
- L. Rieswijk

Teaching methods:
Project-Centered Learning, Work in subgroups, Presentation(s), Skills
Assessment methods:
Assignment, Presentation and paper, Participation

# Computational Statistics

## 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.

KEN4258
Period 4
5 Feb 2024
5 Apr 2024
Print course description
ECTS credits:
6.0
Coordinator:

- C.J. Seiler

Teaching methods:
Project-Centered Learning
Assessment methods:
Written exam, Assignment

# Algorithms for Big Data

## 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, elementary discrete probability

## Recommended reading

None.

KEN4254
Period 5
8 Apr 2024
7 Jun 2024
Print course description
ECTS credits:
6.0
Coordinator:

- M. Mihalak

Teaching methods:
Project-Centered Learning
Assessment methods:
Written exam, Assessment

# Research Project Data Science for Decision Making 2

## 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 period, the goal is to start modelling, prototyping and developing. In period 3, the implementation, model and/or experiments set out in the project plan has to be finished and reported on. At the end of period 1 and 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 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

Justin Zobel (2004), Writing for Computer Science, Springer, ISBN:1852338024

KEN4231
Semester 2
5 Feb 2024
5 Jul 2024
[Print course description]
ECTS credits:
6.0
Instruction language:
English
Coordinators:

- [Linda Rieswijk]
- [L. Rieswijk]

Teaching methods:
Project-Centered Learning, Work in subgroups, Presentation(s), Skills
Assessment methods:
Assignment, Presentation and paper, Participation

# DSDM Electives

# Mathematical Optimization

### Full course description

Optimization (or “Optimisation”) 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 two different areas of optimization: nonlinear optimization and combinatorial optimization. Nonlinear optimization deals with the situation that there is a
continuum of available solutions. A best solution is then usually approximated with one of several available general-purpose algorithms, such as Brent's method for one-dimensional problems, Newton, 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. Combinatorial optimization deals with situations that a best solution from a finite number of available solutions must be chosen. A variety of techniques, such as linear programming, branch and cut, Lagrange relaxation dynamic programming and approximation algorithms are employed to tackle this type of problems. Throughout the course, we aim to provide a coherent framework for the subject, with a focus on consideration of optimality conditions (notably the Karush-Kuhn-Tucker conditions), Lagrange multipliers and duality, relaxation and approximate problems, and on convergence rates and computational complexity.
The methods will be illustrated by in-class computer demonstrations, exercises illustrating the main concepts and algorithms, and modelling and computational work on case studies of practical interest, such as optimal control and network flow.

### Prerequisites

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

### Recommended reading

1. Nonlinear Programming, Theory and Algorithms, by Bazaraa, Sherali, and Shetty (Wiley). 2. Combinatorial Optimization, Algorithm and Complexity, by Papadimitriou and Steiglitz (Dover Publications).

KEN4211
Period 1
4 Sep 2023
27 Oct 2023
Print course description
ECTS credits:
6.0
Instruction language:
English
Coordinator:

- P.J. Collins

Teaching methods:
Project-Centered Learning
Assessment methods:
Written exam

## Stochastic Decision Making

## Full course description

Any realistic model of a real-world phenomenon must take into account the possibility of randomness. That is, more often than not, the quantities we are interested in will not be predictable in advance but, rather, will exhibit an inherent variation that should be taken into account by the model. Mathematically, this is usually accomplished by allowing the model to be probabilistic in nature. In this course, the following topics will be discussed:

(1) Basic concepts of probability theory: Probabilities, conditional probabilities, random variables, probability distribution functions, density functions, expectations and variances.

(2) Finding probabilities, expectations and variances of random variables in complex probabilistic experiments.

(3) Discrete and continuous time Markov chains and related stochastic processes like random walks, branching processes, Poisson processes, birth and death processes, queueing theory.

(4) Markov decision problems.

(5) Multi-armed bandit problems, bandit algorithms, contextual bandits, cumulative regret, and simple regret

## Prerequisites

Probability & Statistics.

## Recommended reading

Probability: A Lively Introduction by Henk Tijms; Reinforcement Learning by Richard S. Sutton and Andrew G. Barto (2nd ed.) (chapter 2); Bandit Algorithms by Tor Lattimore and Csaba Szepesvári

KEN4221
Period 1
4 Sep 2023
27 Oct 2023
Print course description
ECTS credits:
6.0
Instruction language:
English
Coordinator:

- G.M. Schoenmakers

Teaching methods:
Project-Centered Learning
Assessment methods:
Written exam

# Signal and Image Processing

## 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.

KEN4222
Period 1
4 Sep 2023
27 Oct 2023
Print course description
ECTS credits:
6.0
Instruction language:
English
Coordinators:

- J.M.H. Karel
- P. Bonizzi

Teaching methods:
Project-Centered Learning
Assessment methods:
Written exam, Computer test

# Advanced Concepts in Machine Learning

## 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.

KEN4154
Period 2
30 Oct 2023
22 Dec 2023
Print course description
ECTS credits:
6.0
Coordinator:

- E. Hortal Quesada

Teaching methods:
Project-Centered Learning
Assessment methods:
Written exam, Assignment

# Deep Learning for Image & Video Processing

## Full course description

Applications of image and video processing will be presented, and connections to basic algorithms will be demonstrated. We will examine some of the most popular and widespread applications, namely security, surveillance, medical, traffic monitoring, astronomy, farming, culture. The methods used in these applications will be analysed in class and common characteristics between them will be explained. Students will be able to suggest further applications of interest to them and bring relevant literature to the class.

### Prerequisites

**Desired prior knowledge:** Image and Video Processing, Calculus, Linear Algebra, Machine Learning.

### Recommended reading

Rafael C. Gonzalez and Richard E. Woods, Digital Image Processing (3rd Edition), Prentice Hall.
A. Bovik (Ed.), The Essential Guide to Video Processing. Academic Press, 2009.

KEN4244
Period 2
30 Oct 2023
22 Dec 2023
Print course description
ECTS credits:
6.0
Coordinator:

- A. Briassouli

Teaching methods:
Project-Centered Learning
Assessment methods:
Assignment, Presentation, Participation

# Advanced Natural Language Processing

### Full course description

For decades, teaching a computer to deal with natural language processing (NLP) was a long-time dream of humankind. Task such as machine translation, summarization, question-answering, speech recognition or chatting remained a challenge for computer program. Around 2020, 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 and GPT) in depth, which models works best for which tasks, their capacities, limitations and how to optimize these.

Although that we have algorithms that can deal with Natural Language Processing in ways that can no longer be distinguished from humans, we still have some 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. We need to develop models that are less computationally (and thus energy) intensive. (iii) Now that these algorithms operate at human-level quality, several ethical problems arise related to computer generated fake-news, fake profiles, bias, and other abuse. But there are also ethical, legal, regulatory and privacy challenges. In this courses, these important topics will also be discussed.

This course is closely related with the course Information Retrieval and Text-Mining (IRTM). In this course the focus is more on advanced methods and architectures to deal with complex natural language tasks. The IRTM course focusses more on building search engines and text-analytics, but also uses a number of the architectures which are discussed in more depth in this course. The overlap between the two courses is kept to a minimum. There is no need to follow the courses in a specific order.

### Prerequisites

None.

### Recommended reading

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

KEN4259
Period 2
30 Oct 2023
22 Dec 2023
[Print course description]
ECTS credits:
6.0
Coordinator:

- [J. Scholtes]

Teaching methods:
Project-Centered Learning
Assessment methods:
Written exam, Assignment

# Data Fusion

### 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.

KEN4223
Period 4
5 Feb 2024
5 Apr 2024
[Print course description]
ECTS credits:
6.0
Coordinators:

- [A.M. Wilbik]
- [M.C. ten Thij]

Teaching methods:
Project-Centered Learning
Assessment methods:
Written exam, Assignment

# Explainable AI

## 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

## Recommended reading

- Molnar, Christoph. Interpretable Machine Learning. Lulu. com, 2020.
- 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.

KEN4246
Period 4
5 Feb 2024
5 Apr 2024
Print course description
ECTS credits:
6.0
Coordinators:

- N. Tintarev

Data Science for Decision Making

- T.D. Rienstra

Teaching methods:
Project-Centered Learning
Assessment methods:
Assignment, Written exam

# Dynamic Game Theory

## 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.

KEN4251
Period 4
5 Feb 2024
5 Apr 2024
Print course description
ECTS credits:
6.0
Instruction language:
English
Coordinators:

- F. Thuijsman
- M. Salvioli
- M. Salvioli

Teaching methods:
Project-Centered Learning
Assessment methods:
Written exam

# Planning and Scheduling

## 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.

KEN4253
Period 4
5 Feb 2024
5 Apr 2024
Print course description
ECTS credits:
6.0
Instruction language:
English
Coordinators:

- S. Kelk
- S.M. Kelk

Teaching methods:
Project-Centered Learning
Assessment methods:
Written exam, Assessment

# Building and Mining Knowledge Graphs

## 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

KEN4256
Period 4
5 Feb 2024
5 Apr 2024
Print course description
ECTS credits:
6.0
Coordinators:

- M.J. Dumontier
- C.A.W. Brewster

Teaching methods:
Project-Centered Learning
Assessment methods:
Assignment, Assessment

# Information Retrieval and Text Mining

## Full course description

Using today’s search engines allows us to find the needle in the haystack much easier than before. But how do you find out what the needle looks like and where the haystack is? That is exactly the problem we will discuss in this course.  An important difference with standard information retrieval (search) techniques is that they require a user to know what he or she is looking for, while text mining attempts to discover information that is not known beforehand. This is very relevant, for example, in criminal investigations, legal discovery, (business) intelligence, sentiment- & emotion mining or clinical research. Text mining refers generally to the process of extracting interesting and non-trivial information and knowledge from unstructured text. Text mining encompasses several computer science disciplines with a strong orientation towards artificial intelligence in general, including but not limited to information retrieval (building a search engine), statistical pattern recognition, natural language processing, information extraction and different methods of machine learning (including deep learning), clustering and ultimately integrating it all using advanced data visualization and chatbots to make the search experience easier and better.

In this course we will also discuss ethical aspect of using Artificial Intelligence for the above tasks, including the need for eXplainable AI (XAI), training deep-learning large language-models more energy efficient, and several ethical problems that may arise related to bias, legal, regulatory and privacy challenges.

This course is closely related with the course Advanced Natural Language Processing (ANLP). In the ANLP course, the focus is more on advanced methods and architectures to deal with complex natural language tasks such as machine translation, and Q&A systems. IRTM focusses more on building search engines and using text-analytics to improve the search experience. In the IRTM course, we will use a number of the architectures that are discussed in more detail in ANLP. The overlap between the two courses is kept to a minimum. There is no need to follow the courses in a specific order.

## Prerequisites

None.

## 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.

KEN4153
Period 5
8 Apr 2024
7 Jun 2024
Print course description
ECTS credits:
6.0
Instruction language:
English
Coordinator:

- J.C. Scholtes

Teaching methods:
Project-Centered Learning
Assessment methods:

Data Science for Decision Making

Written exam, Assessment
Dept. of Advanced Computing Sciences

# Introduction to Quantum Computing for AI and Data Science

## 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 planned elective courses Quantum Algorithms, Quantum AI, and Quantum Information and Security, **which will be offered in Semester 1 of the upcoming academic year 2024-2025.** These four courses, together with a dedicated research project on quantum computing forms the specialization in Quantum Computing for AI and Data Science.

## Prerequisites

None.

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

!! This course is a prerequisite for the planned elective courses Quantum Algorithms, Quantum AI, and Quantum Information and Security, **which will be offered in Semester 1 of the upcoming academic year 2024-2025.** 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

To be announced.

KEN4155
Period 5
8 Apr 2024
7 Jun 2024
Print course description
ECTS credits:
6.0
Coordinators:

- G. Stamoulis
- M. Dibenedetto
- D. Dibenedetto

Data Science for Decision Making

Teaching methods:
Project-Centered Learning
Assessment methods:
Written exam

# Symbolic Computation and Control

## Full course description

This course consists of two interrelated parts. In the first part, we focus on basic techniques for the digital control of linear dynamical systems using feedback. We start by addressing system stability and we discuss the technique of pole placement by state feedback to solve the regulation problem. Then we introduce state observers to solve the regulation problem by output feedback. Next, we extend our scope to tracking problems. This involves the design of additional dynamics to characterize the relevant class of reference signals, which are then integrated with the earlier set-up for output feedback. Finally, we discuss the classical topic of optimal control, which can be employed to avoid using prototype systems for pole placement, and which allows the user to design a feedback law by trading off the cost involved in generating large inputs against the achieved tracking accuracy. In the second part, we address computational issues, related to the field of systems and control. Classically, computers have been designed primarily to perform approximate numerical arithmetic. Modern software packages for mathematical computation, such as Maple and Mathematica, allow one to perform exact and symbolic computation too. We shall explore this new area. It is demonstrated how speed, efficiency and memory usage considerations often lead to surprising and fundamentally different algorithmic solutions in a symbolic or exact context. Applications and examples involve stability of linear systems, model approximation, and linear matrix equations with free parameters. Practical classes serve to demonstrate the techniques and to make the student familiar with exact and symbolic computation.

## Prerequisites

**Desired Prior Knowledge:** Linear Algebra, Calculus, Mathematical Modelling.

## Recommended reading

Richard J. Vaccaro, Digital Control - A State-Space Approach, McGraw-Hill International Editions, 1995. ISBN 0-07-066781-0.

KEN4252
Period 5
8 Apr 2024
7 Jun 2024
Print course description
ECTS credits:
6.0
Instruction language:
English
Coordinator:

Data Science for Decision Making

- R.L.M. Peeters

Teaching methods:
Project-Centered Learning
Assessment methods:
Computer test, Take home exam

# Computer Vision

## 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? How do search engines index and retrieve billions of images? 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, 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 analysis, deep learning based object classification, 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.

KEN4255
Period 5
8 Apr 2024
7 Jun 2024
Print course description
ECTS credits:
6.0
Coordinator:

- M.C. Popa

Teaching methods:
Lecture(s), Project-Centered Learning

Data Science for Decision Making

Assessment methods:
Written exam, Assignment
Second year courses

# DSDM Electives

# Foundations of Agents

### Full course description

Agents are autonomous computer programs, robots, humans, etc. Agents operate in some environment, which they can observe, and in which they can realize objectives through the execution of actions. Examples of environment in which agents can operate, are computer game environments, the internet, and also the physical world is case of robots and humans.

In this course we address the problem of how an agent can act optimally in order to realize its objectives. We will answer this question by investigating how we can formally specify the agent's environment, the agent's objectives, the observations the agent can make and the actions it can execute. We use the formal model to investigate how the agent can determine an (optimal) behaviour realizing its objectives.

The following formal models will be investigated:

- Markov Decision Processes,
- Partially Observable Markov Decision Process,
- logic-based models such as Epistemic Logic, Doxastic Logic, Dynamic Logic, and BDI logics, and
- Game Theory.

Some examples of methods for determining the agents optimal behaviour addressed in the course are: Value and Policy Iteration, Q-Learning, Planning, etc.

### Prerequisites

Desired prior knowledge: A basic course in logic and in probability theory.

### Recommended reading

None

KEN4115
Period 1
4 Sep 2023
27 Oct 2023
[Print course description]
ECTS credits:
6.0
Instruction language:

Data Science for Decision Making

English
Coordinator:

- N. Roos

Teaching methods:
Project-Centered Learning
Assessment methods:
Written exam, Assignment

# Intelligent Search & Games

## 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:
(1) Basic search techniques. Alpha-beta; A*.
(2) Advanced search techniques. IDA*; B*, transposition tables; retrograde analysis and endgame databases; proof-number search and variants; multi-player search methods; Expectimax and *-minimax variants.
(3) Heuristics. World representations; killer moves; history heuristic, PVS; windowing techniques; null-moves; forward-pruning techniques; selective search, GOAP.
(4) Monte Carlo methods. Monte Carlo tree search (MCTS) techniques, enhancements and applications; AlphaGo and AlphaZero approaches.
(5) (5) Game design. Evolutionary game design; game quality metrics; self-play evaluation; procedural content generation (PCG); puzzle design.

## Prerequisites

None.

**Desired Prior Knowledge**: Data Structures & Algorithms.

## Recommended reading

- Millington, I. and Funge, J. (2009). Artificial Intelligence for Games, 2nd Edition Morgan Kaufmann Publishers, ISBN: 978-0123747310
- Russell, S.J. and Norvig, P. (2010). Artificial Intelligence: A Modern Approach, 3rd edition. Pearson Education, New Jersey. ISBN 0-13-207148-7.
- Yannakakis, G.N. and Togelius, J. (2018) Artificial Intelligence and Games, Springer, Berlin. ISBN 978-3-319-63519-4 (eBook) 978-3-319-63518-7 (hardcover)

KEN4123
Period 1
4 Sep 2023
27 Oct 2023
Print course description
ECTS credits:
6.0

Data Science for Decision Making

Instruction language:
English
Coordinators:

- M.H.M. Winands
- C.B. Browne

Teaching methods:
Project-Centered Learning
Assessment methods:
Written exam, Assessment

# Mathematical Optimization

## Full course description

Optimization (or “Optimisation”) 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 two different areas of optimization: nonlinear optimization and combinatorial optimization. Nonlinear optimization deals with the situation that there is a
continuum of available solutions. A best solution is then usually approximated with one of several available general-purpose algorithms, such as Brent’s method for one-dimensional problems, Newton, 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. Combinatorial optimization deals with situations that a best solution from a finite number of available solutions must be chosen. A variety of techniques, such as linear programming, branch and cut, Lagrange relaxation dynamic programming and approximation algorithms are employed to tackle this type of problems. Throughout the course, we aim to provide a coherent framework for the subject, with a focus on consideration of optimality conditions (notably the Karush-Kuhn-Tucker conditions), Lagrange multipliers and duality, relaxation and approximate problems, and on convergence rates and computational complexity.
The methods will be illustrated by in-class computer demonstrations, exercises illustrating the main concepts and algorithms, and modelling and computational work on case studies of practical interest, such as optimal control and network flow.

## Prerequisites

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

## Recommended reading

1. Nonlinear Programming, Theory and Algorithms, by Bazaraa, Sherali, and Shetty (Wiley). 2. Combinatorial Optimization, Algorithm and Complexity, by Papadimitriou and Steiglitz (Dover Publications).

KEN4211
Period 1

Data Science for Decision Making

4 Sep 2023
27 Oct 2023
Print course description
ECTS credits:
6.0
Instruction language:
English
Coordinator:

- P.J. Collins

Teaching methods:
Project-Centered Learning
Assessment methods:
Written exam

# Stochastic Decision Making

## Full course description

Any realistic model of a real-world phenomenon must take into account the possibility of randomness. That is, more often than not, the quantities we are interested in will not be predictable in advance but, rather, will exhibit an inherent variation that should be taken into account by the model. Mathematically, this is usually accomplished by allowing the model to be probabilistic in nature. In this course, the following topics will be discussed:

(1) Basic concepts of probability theory: Probabilities, conditional probabilities, random variables, probability distribution functions, density functions, expectations and variances.

(2) Finding probabilities, expectations and variances of random variables in complex probabilistic experiments.

(3) Discrete and continuous time Markov chains and related stochastic processes like random walks, branching processes, Poisson processes, birth and death processes, queueing theory.

(4) Markov decision problems.

(5) Multi-armed bandit problems, bandit algorithms, contextual bandits, cumulative regret, and simple regret

## Prerequisites

Probability & Statistics.

## Recommended reading

Probability: A Lively Introduction by Henk Tijms; Reinforcement Learning by Richard S. Sutton and Andrew G. Barto (2nd ed.) (chapter 2); Bandit Algorithms by Tor Lattimore and Csaba Szepesvári

KEN4221

Data Science for Decision Making

Period 1
4 Sep 2023
27 Oct 2023
Print course description
ECTS credits:
6.0
Instruction language:
English
Coordinator:

- G.M. Schoenmakers

Teaching methods:
Project-Centered Learning
Assessment methods:
Written exam

# Signal and Image Processing

## 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.

KEN4222
Period 1
4 Sep 2023
27 Oct 2023

Data Science for Decision Making

[Print course description](#)
ECTS credits:
6.0
Instruction language:
English
Coordinators:

- [J.M.H. Karel](#)
- [P. Bonizzi](#)

Teaching methods:
Project-Centered Learning
Assessment methods:
Written exam, Computer test

# Advanced Concepts in Machine Learning

## 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.

KEN4154
Period 2
30 Oct 2023
22 Dec 2023
[Print course description](#)
ECTS credits:
6.0
Coordinator:

- [E. Hortal Quesada](#)

Data Science for Decision Making

Teaching methods:
Project-Centered Learning
Assessment methods:
Written exam, Assignment

# Deep Learning for Image & Video Processing

## Full course description

Applications of image and video processing will be presented, and connections to basic algorithms will be demonstrated. We will examine some of the most popular and widespread applications, namely security, surveillance, medical, traffic monitoring, astronomy, farming, culture. The methods used in these applications will be analysed in class and common characteristics between them will be explained. Students will be able to suggest further applications of interest to them and bring relevant literature to the class.

## Prerequisites

**Desired prior knowledge:** Image and Video Processing, Calculus, Linear Algebra, Machine Learning.

## Recommended reading

Rafael C. Gonzalez and Richard E. Woods, Digital Image Processing (3rd Edition), Prentice Hall.
A. Bovik (Ed.), The Essential Guide to Video Processing. Academic Press, 2009.

KEN4244
Period 2
30 Oct 2023
22 Dec 2023
Print course description
ECTS credits:
6.0
Coordinator:

- A. Briassouli

Teaching methods:
Project-Centered Learning
Assessment methods:
Assignment, Presentation, Participation

# Advanced Natural Language Processing

## Full course description

For decades, teaching a computer to deal with natural language processing (NLP) was a long-time dream of humankind. Task such as machine translation, summarization, question-answering, speech recognition or chatting remained a challenge for computer program. Around 2020, 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 and GPT) in depth, which models works best for which tasks, their capacities, limitations and how to optimize these.

Although that we have algorithms that can deal with Natural Language Processing in ways that can no longer be distinguished from humans, we still have some 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. We need to develop models that are less computationally (and thus energy) intensive. (iii) Now that these algorithms operate at human-level quality, several ethical problems arise related to computer generated fake-news, fake profiles, bias, and other abuse. But there are also ethical, legal, regulatory and privacy challenges. In this courses, these important topics will also be discussed.

This course is closely related with the course Information Retrieval and Text-Mining (IRTM). In this course the focus is more on advanced methods and architectures to deal with complex natural language tasks. The IRTM course focusses more on building search engines and text-analytics, but also uses a number of the architectures which are discussed in more depth in this course. The overlap between the two courses is kept to a minimum. There is no need to follow the courses in a specific order.

## Prerequisites

None.

## Recommended reading

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

KEN4259
Period 2
30 Oct 2023
22 Dec 2023
Print course description
ECTS credits:
6.0
Coordinator:

- J. Scholtes

Teaching methods:
Project-Centered Learning
Assessment methods:
Written exam, Assignment

# Data Fusion

## 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.

KEN4223
Period 4
5 Feb 2024
5 Apr 2024
Print course description
ECTS credits:
6.0

Data Science for Decision Making

Coordinators:

- A.M. Wilbik
- M.C. ten Thij

Teaching methods:
Project-Centered Learning
Assessment methods:
Written exam, Assignment

# Explainable AI

## 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

## Recommended reading

- Molnar, Christoph. Interpretable Machine Learning. Lulu. com, 2020.

- 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.

KEN4246
Period 4
5 Feb 2024
5 Apr 2024
Print course description
ECTS credits:
6.0
Coordinators:

- N. Tintarev
- T.D. Rienstra

Teaching methods:
Project-Centered Learning
Assessment methods:
Assignment, Written exam

# Dynamic Game Theory

## 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.

KEN4251
Period 4
5 Feb 2024
5 Apr 2024
Print course description
ECTS credits:
6.0
Instruction language:
English
Coordinators:

Data Science for Decision Making

- [F. Thuijsman]
- M. Salvioli
- [M. Salvioli]

Teaching methods:
Project-Centered Learning
Assessment methods:
Written exam

# Planning and Scheduling

## 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.

KEN4253
Period 4
5 Feb 2024
5 Apr 2024
[Print course description]
ECTS credits:
6.0
Instruction language:
English
Coordinators:

Data Science for Decision Making

- S. Kelk
- S.M. Kelk

Teaching methods:
Project-Centered Learning
Assessment methods:
Written exam, Assessment

# Building and Mining Knowledge Graphs

## 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

KEN4256
Period 4
5 Feb 2024
5 Apr 2024
Print course description
ECTS credits:
6.0
Coordinators:

- M.J. Dumontier
- C.A.W. Brewster

Teaching methods:
Project-Centered Learning
Assessment methods:
Assignment, Assessment

# Agents and Multi-Agent Systems

## 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.

## 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.

KEN4111
Period 4
5 Feb 2024
5 Apr 2024
Print course description
ECTS credits:
6.0
Instruction language:
English
Coordinators:

- G.B. Weiss
- N. Roos

Teaching methods:
Project-Centered Learning
Assessment methods:
Written exam, Assignment

# Autonomous Robotic Systems

## 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

KEN4114
Period 5
8 Apr 2024
7 Jun 2024
Print course description
ECTS credits:
6.0
Instruction language:
English
Coordinator:

- R. Möckel

Teaching methods:
Project-Centered Learning
Assessment methods:
Written exam, Assignment

# Information Retrieval and Text Mining

## Full course description

Using today’s search engines allows us to find the needle in the haystack much easier than before. But how do you find out what the needle looks like and where the haystack is? That is exactly the problem we will discuss in this course.  An important difference with standard information retrieval (search) techniques is that they require a user to know what he or she is looking for, while text mining attempts to discover information that is not known beforehand. This is very relevant, for example, in criminal investigations, legal discovery, (business) intelligence, sentiment- & emotion mining or clinical research. Text mining refers generally to the process of extracting interesting and non-trivial information and knowledge from unstructured text. Text mining encompasses several computer science disciplines with a strong orientation towards artificial intelligence in general, including but not limited to information retrieval (building a search engine), statistical pattern recognition, natural language processing, information extraction and different methods of machine learning (including deep learning), clustering and ultimately integrating it all using advanced data visualization and chatbots to make the search experience easier and better.

In this course we will also discuss ethical aspect of using Artificial Intelligence for the above tasks, including the need for eXplainable AI (XAI), training deep-learning large language-models more energy efficient, and several ethical problems that may arise related to bias, legal, regulatory and privacy challenges.

This course is closely related with the course Advanced Natural Language Processing (ANLP). In the ANLP course, the focus is more on advanced methods and architectures to deal with complex natural language tasks such as machine translation, and Q&A systems. IRTM focusses more on building search engines and using text-analytics to improve the search experience. In the IRTM course, we will use a number of the architectures that are discussed in more detail in ANLP. The overlap between the two courses is kept to a minimum. There is no need to follow the courses in a specific order.

## Prerequisites

None.

## 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](http://informationretrieval.org).

KEN4153
Period 5
8 Apr 2024
7 Jun 2024
Print course description
ECTS credits:
6.0
Instruction language:
English
Coordinator:

Data Science for Decision Making

- [J.C. Scholtes]

Teaching methods:
Project-Centered Learning
Assessment methods:
Written exam, Assessment
Dept. of Advanced Computing Sciences

# Reinforcement Learning

## 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

KEN4157
Period 5
8 Apr 2024
7 Jun 2024
[Print course description]
ECTS credits:
6.0
Coordinator:

- [K. Driessens]

Data Science for Decision Making

Teaching methods:
Project-Centered Learning
Assessment methods:
Assignment, Take home exam

# Symbolic Computation and Control

## Full course description

This course consists of two interrelated parts. In the first part, we focus on basic techniques for the digital control of linear dynamical systems using feedback. We start by addressing system stability and we discuss the technique of pole placement by state feedback to solve the regulation problem. Then we introduce state observers to solve the regulation problem by output feedback. Next, we extend our scope to tracking problems. This involves the design of additional dynamics to characterize the relevant class of reference signals, which are then integrated with the earlier set-up for output feedback. Finally, we discuss the classical topic of optimal control, which can be employed to avoid using prototype systems for pole placement, and which allows the user to design a feedback law by trading off the cost involved in generating large inputs against the achieved tracking accuracy. In the second part, we address computational issues, related to the field of systems and control. Classically, computers have been designed primarily to perform approximate numerical arithmetic. Modern software packages for mathematical computation, such as Maple and Mathematica, allow one to perform exact and symbolic computation too. We shall explore this new area. It is demonstrated how speed, efficiency and memory usage considerations often lead to surprising and fundamentally different algorithmic solutions in a symbolic or exact context. Applications and examples involve stability of linear systems, model approximation, and linear matrix equations with free parameters. Practical classes serve to demonstrate the techniques and to make the student familiar with exact and symbolic computation.

## Prerequisites

**Desired Prior Knowledge:** Linear Algebra, Calculus, Mathematical Modelling.

## Recommended reading

Richard J. Vaccaro, Digital Control - A State-Space Approach, McGraw-Hill International Editions, 1995. ISBN 0-07-066781-0.

KEN4252
Period 5
8 Apr 2024
7 Jun 2024
[Print course description]
ECTS credits:
6.0
Instruction language:
English
Coordinator:

Data Science for Decision Making

- R.L.M. Peeters

Teaching methods:
Project-Centered Learning
Assessment methods:
Computer test, Take home exam

# Computer Vision

## 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? How do search engines index and retrieve billions of images? 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, 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 analysis, deep learning based object classification, 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.

KEN4255
Period 5
8 Apr 2024
7 Jun 2024
Print course description
ECTS credits:
6.0
Coordinator:

- M.C. Popa

Teaching methods:
Lecture(s), Project-Centered Learning

Data Science for Decision Making

Assessment methods:
Written exam, Assignment

# Study Abroad

## 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.

For current first year master's students, the next possibility to study a semester abroad is in the fall semester of the subsequent academic year or during the spring semester (for February intake students). This section will explain your possibilities and the necessary procedures to undertake. If you might be interested in studying a semester abroad, please read this complete section and all attached documents very carefully!

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

## Prerequisites

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

KEN3600
Semester 1
4 Sep 2023
2 Feb 2024
Semester 2
5 Feb 2024
5 Jul 2024
Print course description
ECTS credits:
30.0
Coordinators:

Data Science for Decision Making

- Luc Giezenaar
- J.M.H. Karel
- M. Musegaas

Assessment methods:
Written exam, Attendance, Assignment

# Master Internship

## Full course description

In the master elective semester, 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 send 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 Claire van Doorn.

## Prerequisites

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.

KEN4176
Year
1 Sep 2023
31 Aug 2024
Print course description
ECTS credits:
30.0
Instruction language:
English
Coordinators:

- C. van Doorn

Data Science for Decision Making

- K. Driessens

Assessment methods:
Final paper

# DSDM Thesis

# Master's Thesis DSDM

### Full course description

The Master Data Science for Decision Making 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.

KEN4260
Year
1 Sep 2023
31 Aug 2024
Print course description
ECTS credits:
30.0
Instruction language:
English
Coordinator:

- M. Mihalak

Teaching methods:
Paper(s)
Assessment methods:
Presentation and paper