DSDM Year 1 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
1 Sep 2021
22 Oct 2021
[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
Period 2
25 Oct 2021
17 Dec 2021
Print course description
ECTS credits:
6.0
Instruction language:
English
Coordinators:

- R.L.M. Peeters
- P. Bonizzi

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

Research Project DSDM 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
1 Sep 2021
28 Jan 2022
Print course description
ECTS credits:
6.0
Instruction language:
English
Coordinators:

- G.M. Schoenmakers
- K. Driessens

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

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 4
1 Feb 2022
1 Apr 2022
Print course description
ECTS credits:
6.0
Coordinator:

- M. Mihalak

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

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

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

Recommended reading

None.

KEN4253
Period 5
4 Apr 2022
3 Jun 2022
Print course description
ECTS credits:
6.0
Instruction language:
English
Coordinator:

- M. Mihalak

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

Research Project DSDM 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
1 Feb 2022
1 Jul 2022
Print course description
ECTS credits:
6.0
Instruction language:
English
Coordinator:

- K. Driessens

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

DSDM Year 1 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
1 Sep 2021
22 Oct 2021
Print course description
ECTS credits:
6.0
Instruction language:
English
Coordinator:

- P.J. Collins

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
1 Sep 2021
22 Oct 2021
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

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.

Prerequisites

Probability & Statistics.

Recommended reading

Probability: A Lively Introduction by Henk Tijms.

KEN4221
Period 1
1 Sep 2021
22 Oct 2021
Print course description
ECTS credits:
6.0
Instruction language:
English
Coordinator:

- G.M. Schoenmakers

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

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
25 Oct 2021
17 Dec 2021

Data Science for Decision Making

Print course description
ECTS credits:
6.0
Coordinators:

- K. Driessens
- D. Dibenedetto

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

Applications of Image and 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
25 Oct 2021
17 Dec 2021
Print course description
ECTS credits:
6.0
Coordinator:

- A. Briassouli

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

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
1 Feb 2022
1 Apr 2022
Print course description
ECTS credits:
6.0
Coordinators:

- M.J. Dumontier
- R. Celebi
- K. Moodley

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

Computational Statistics

Full course description

Complex and high dimensional data are abundant in academia and industry. At the same time, computers are cheap and powerful. These developments enable us to fit increasingly complex statistical models using computer intensive methods. We will model and analyze both independent and dependent data from real world problems. The course is hands-on; we will use Stan (platform for statistical modelling and high-performance statistical computation) and R (statistical programming language). Key topics: Statistical modelling, uncertainty quantification, Markov chain Monte Carlo, bootstrap resampling, permutation tests, and causal inference.

Prerequisites

Desired prior knowledge: Probability and Statistics (Code: KEN2130)

Recommended reading

Selected chapters:

- Efron and Tibshirani (1993), An Introduction to the Bootstrap
- Hoff (2009), A First Course in Bayesian Statistical Methods
- Grolemund and Wickham (2017), R for Data Science
- Hernán and Robins (2019, forthcoming), Causal Inference

KEN4258
Period 4
1 Feb 2022
1 Apr 2022
Print course description
ECTS credits:
6.0
Coordinator:

- C.J. Seiler

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

Desired prior knowledge: statistics and basic machine learning

Recommended reading

None.

KEN4223
Period 4
1 Feb 2022
1 Apr 2022
Print course description
ECTS credits:
6.0
Coordinator:

- A.M. Wilbik

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

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
1 Feb 2022
1 Apr 2022
[Print course description]
ECTS credits:
6.0
Instruction language:
English
Coordinators:

- F. Thuijsman
- K. Stankova

Teaching methods:
Project-Centered Learning
Assessment methods:
Written 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, object classification and face recognition, camera models, epipolar geometry and 3D reconstruction from 2D views.

Prerequisites

Desired prior knowledge: Basic knowledge of Matlab, 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, 2002 / Digital Image Processing using MATLAB. 2° Edition Rafael C. Gonzalez, Richard E. Woods, Steven L. Eddins. Gatesmark Publishing

KEN4255
Period 5
4 Apr 2022
3 Jun 2022
[Print course description]
ECTS credits:
6.0
Coordinators:

- [S. Asteriadis]
- [M.C. Popa]

Teaching methods:
Lecture(s), Project-Centered Learning
Assessment methods:
Written exam, Assignment

Deep Learning

Full course description

Conventional machine learning techniques were limited in processing data in their raw forms and many domain experts were required in transforming raw data into meaningful features or representations. Deep Learning techniques have revolutionized many application domains ranging from auditory to vision signal processing. In this course, we will study various concepts in deep architectures using both artificial neural networks as well as kernel-based models. Several deep learning models such as convolutional neural networks, auto-encoders, generative adversarial networks and their variants among other state-of-the-art models will be covered in depth. We will further study different types of deep architectures used for domain adaptation problems where one is encountered with heterogeneous datasets as well as multi-modal datasets. The regularization and optimization methods used in deep learning framework will be discussed. Introduction to open-source deep learning platforms will be given. This course will be equipped with a practical component, and students are expected to write their own deep learning code and test its performance on various problems. In addition they are strongly encouraged to participate in mini-projects (in a group or individual) targeting a conference paper.

Prerequisites

Advanced Concepts of Machine Learning.

- Advanced Concepts in Machine Learning

Recommended reading

- Ian Goodfellow and Yoshua Bengio and Aaron Courville (2016). Deep Learning, MIT Press.
- Research Papers Published in high ranked journals and conferences.

KEN4257
Period 5
4 Apr 2022
3 Jun 2022
Print course description
ECTS credits:
6.0
Coordinator:

- S. Mehrkanoon

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

Information Retrieval and Text Mining

Full course description

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, clustering and ultimately data visualization. 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 in a pattern that is not known beforehand. This is very relevant, for example, in criminal investigations, legal discovery, (business) intelligence, sentiment- & emotion mining or clinical research.

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 and Feldman, R., and Sanger, J. (2006). The Text Mining Handbook: Advanced Approaches in Analyzing Unstructured Data. Cambridge University Press.

KEN4153
Period 5
4 Apr 2022
3 Jun 2022
Print course description
ECTS credits:
6.0
Instruction language:
English
Coordinator:

- J.C. Scholtes

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

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
4 Apr 2022
3 Jun 2022
Print course description
ECTS credits:
6.0
Instruction language:
English
Coordinator:

- R.L.M. Peeters

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

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

Data Mining or Advanced Concepts in Machine Learning.

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 5
4 Apr 2022
3 Jun 2022
Print course description
ECTS credits:
6.0
Coordinators:

- N. Tintarev
- T.D. Rienstra

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

DSDM Year 2 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

A basic course in logic and in probability theory.

Recommended reading

None

KEN4115
Period 1
1 Sep 2021
22 Oct 2021
Print course description
ECTS credits:
6.0
Instruction language:
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

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
1 Sep 2021
22 Oct 2021
Print course description
ECTS credits:
6.0
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
1 Sep 2021
22 Oct 2021
Print course description
ECTS credits:
6.0
Instruction language:
English
Coordinator:

- P.J. Collins

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
1 Sep 2021
22 Oct 2021
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

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:

Data Science for Decision Making

(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.

Prerequisites

Probability & Statistics.

Recommended reading

Probability: A Lively Introduction by Henk Tijms.

KEN4221
Period 1
1 Sep 2021
22 Oct 2021
Print course description
ECTS credits:
6.0
Instruction language:
English
Coordinator:

- G.M. Schoenmakers

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

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
25 Oct 2021
17 Dec 2021
Print course description
ECTS credits:
6.0
Coordinators:

- K. Driessens
- D. Dibenedetto

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

Applications of Image and 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

Data Science for Decision Making

Period 2
25 Oct 2021
17 Dec 2021
Print course description
ECTS credits:
6.0
Coordinator:

- A. Briassouli

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

Multi-Agent Systems

Full course description

Multi-agent systems are systems composed of multiple interacting intelligent agents, where an agent is 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 field of multi-agent systems has is origin in the late 1970s and today is an established and vibrant topic in computer science. Multi-agent systems are an enabling technology for applications that rely on distributed and parallel processing of data, information and knowledge in complex – networked, open and large-scale – computing environments. With advancing technological progress in inter-connectivity and interoperability of computers and software such applications are becoming standard in a variety of domains such as e-commerce, logistics, supply chain management, telecommunication, health care, and manufacturing. The course covers important conceptual, theoretical and practical foundations of multi-agent systems. Examples of topics treated in the course are agent-agent communication, automated negotiation and argumentation in cooperative and competitive settings, multi-agent learning and planning, automated decision making based on mechanisms such as voting and auctioning, and development and engineering of agent-based systems. In the practical part of the course students have the opportunity to apply the learnt multi-agent concepts, algorithms and methods.

Prerequisites

Desired Prior Knowledge: Introduction to Computer Science 1 and 2.

Recommended reading

• Gerhard Weiss (Ed.) (2013, 2nd edition): Multi-agent Systems. MIT Press.
• Mike Wooldridge (2009, 2nd edition): An Introduction to Multi Agent Systems, Michael Wooldridge, John Wiley & Sons Ltd.
• Y oav Shoham and Kevin Leyton-Brown (2009): Multi-agent Systems Algorithmic, Game-Theoretic, and Logical Foundations", Cambridge University Press.

KEN4111
Period 2

Data Science for Decision Making

25 Oct 2021
17 Dec 2021
[Print course description]
ECTS credits:
6.0
Instruction language:
English
Coordinator:

- [G.B. Weiss]

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

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 4
1 Feb 2022
1 Apr 2022
[Print course description]
ECTS credits:
6.0
Instruction language:
English

Data Science for Decision Making

Coordinator:

- R. Möckel

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

Desired prior knowledge: statistics and basic machine learning

Recommended reading

None.

KEN4223

Data Science for Decision Making

Period 4
1 Feb 2022
1 Apr 2022
[Print course description]
ECTS credits:
6.0
Coordinator:

- [A.M. Wilbik]

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

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
1 Feb 2022
1 Apr 2022
[Print course description]
ECTS credits:
6.0
Coordinators:

- [M.J. Dumontier]

Data Science for Decision Making

- R. Celebi
- K. Moodley

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

Computational Statistics

Full course description

Complex and high dimensional data are abundant in academia and industry. At the same time, computers are cheap and powerful. These developments enable us to fit increasingly complex statistical models using computer intensive methods. We will model and analyze both independent and dependent data from real world problems. The course is hands-on; we will use Stan (platform for statistical modelling and high-performance statistical computation) and R (statistical programming language). Key topics: Statistical modelling, uncertainty quantification, Markov chain Monte Carlo, bootstrap resampling, permutation tests, and causal inference.

Prerequisites

Desired prior knowledge: Probability and Statistics (Code: KEN2130)

Recommended reading

Selected chapters:

- Efron and Tibshirani (1993), An Introduction to the Bootstrap
- Hoff (2009), A First Course in Bayesian Statistical Methods
- Grolemund and Wickham (2017), R for Data Science
- Hernán and Robins (2019, forthcoming), Causal Inference

KEN4258
Period 4
1 Feb 2022
1 Apr 2022
Print course description
ECTS credits:
6.0
Coordinator:

- C.J. Seiler

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

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
1 Feb 2022
1 Apr 2022
Print course description
ECTS credits:
6.0
Instruction language:
English
Coordinators:

- F. Thuijsman
- K. Stankova

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

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

Data Mining or Advanced Concepts in Machine Learning.

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 5
4 Apr 2022
3 Jun 2022
Print course description
ECTS credits:
6.0
Coordinators:

- N. Tintarev
- T.D. Rienstra

Teaching methods:
Project-Centered Learning
Assessment methods:
Assignment, Written 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, object classification and face recognition, camera models, epipolar geometry and 3D reconstruction from 2D views.

Prerequisites

Desired prior knowledge: Basic knowledge of Matlab, 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, 2002 / Digital Image Processing using MATLAB. 2° Edition Rafael C. Gonzalez, Richard E. Woods, Steven L. Eddins. Gatesmark Publishing

KEN4255
Period 5
4 Apr 2022
3 Jun 2022
[Print course description]
ECTS credits:
6.0
Coordinators:

- [S. Asteriadis]
- [M.C. Popa]

Teaching methods:
Lecture(s), Project-Centered Learning
Assessment methods:
Written exam, Assignment

Deep Learning

Full course description

Conventional machine learning techniques were limited in processing data in their raw forms and many domain experts were required in transforming raw data into meaningful features or representations. Deep Learning techniques have revolutionized many application domains ranging from auditory to vision signal processing. In this course, we will study various concepts in deep architectures using both artificial neural networks as well as kernel-based models. Several deep learning models such as convolutional neural networks, auto-encoders, generative adversarial networks and their variants among other state-of-the-art models will be covered in depth. We will further study different types of deep architectures used for domain adaptation problems where one is encountered with heterogeneous datasets as well as multi-modal datasets. The regularization and optimization methods used in deep learning framework will be discussed. Introduction to open-source deep learning platforms will be given. This course will be equipped with a practical component, and students are expected to write their own deep learning code and test its performance on various problems. In addition they are strongly encouraged to participate in mini-projects (in a group or individual) targeting a conference paper.

Prerequisites

Advanced Concepts of Machine Learning.

- Advanced Concepts in Machine Learning

Recommended reading

- Ian Goodfellow and Yoshua Bengio and Aaron Courville (2016). Deep Learning, MIT Press.
- Research Papers Published in high ranked journals and conferences.

KEN4257
Period 5
4 Apr 2022
3 Jun 2022
Print course description
ECTS credits:
6.0
Coordinator:

- S. Mehrkanoon

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

Information Retrieval and Text Mining

Full course description

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, clustering and ultimately data visualization. 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 in a pattern that is not known beforehand. This is very relevant, for example, in criminal investigations, legal discovery, (business) intelligence, sentiment- & emotion mining or clinical research.

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 and Feldman, R., and Sanger, J. (2006). The Text Mining Handbook: Advanced Approaches in Analyzing Unstructured Data. Cambridge University Press.

KEN4153
Period 5
4 Apr 2022
3 Jun 2022
Print course description
ECTS credits:
6.0
Instruction language:
English
Coordinator:

- J.C. Scholtes

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

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
4 Apr 2022
3 Jun 2022
Print course description
ECTS credits:
6.0
Instruction language:
English
Coordinator:

- R.L.M. Peeters

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

Study Abroad

Full course description

Students can apply to study abroad for a semester, at another University with whom Maastricht University has an Agreement of Exchange.
Nomination is decided on by the Board of Examiners based on study progress as mentioned in

Data Science for Decision Making

Article 5.3.1 of the Student Handbook and motivation of the student.

This study abroad will take place in Semester 1 of year 2 and has a study load of 30 ECTS.
The selected course programme has to be approved by the Board of Examiners.

For more information you can contact our Study Adviser Wendy Brandt.

Prerequisites

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

KEN3600
Semester 1
1 Sep 2021
28 Jan 2022
Semester 2
1 Feb 2022
1 Jul 2022
Print course description
ECTS credits:
30.0
Coordinators:

- W. Brandt
- J.M.H. Karel

Assessment methods:
Written exam, Attendance, Assignment

DSDM Year 2

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 60 ECTS (among

Data Science for Decision Making

which are 40 credits of the first year).

Recommended reading

None.

KEN4260
Year
1 Sep 2021
31 Aug 2022
[Print course description]
ECTS credits:
30.0
Instruction language:
English
Coordinator:

- [M. Mihalak]

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

Master Internship

Full course description

DKE can offer internship vacancies or the student may find a relevant internship and DKE supervisor him/herself.
For each internship, an internship proposal must be send to the Board of Examiners for approval. .

The Board of Examiners appoints an examiner for the internship.

The student may undertake an internship supervised by DKE at most once during the programme.

The Board of Examiners may formulate guidelines for internships. The guidelines will be included in the Rules and Regulations.

For more information you contact our Study Advisor Wendy Brandt.

Prerequisites

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

KEN4176
Year

Data Science for Decision Making

1 Sep 2021
31 Aug 2022
[Print course description]
ECTS credits:
30.0
Instruction language:
English
Coordinators:

- K. Driessens
- C. van Doorn

Assessment methods:
Final paper

Electives

KEN4170
Semester 1
1 Sep 2021
28 Jan 2022
[Print course description]
ECTS credits:
6.0
Coordinator:

- J.M.H. Karel