

MAI YR1

Data Science & Knowledge Engineering

Algorithms for Big Data

Full course description

Algorithms for Big Data presents an algorithmic toolkit to efficiently deal with the challenges that the ever growing amount of data pose. For example, the data may not fit into 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 inputs. 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. Standard algorithmic solutions do not address these challenges, and new algorithmic techniques are needed. This course looks at a number of algorithmic responses to these problems, such as algorithms with linear or sublinear running time, algorithms where the data arrives in a stream, and 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). It also covers a number of topics from classical algorithm design that have undiminished relevance in the era of big data, such as approximation algorithms (obtaining a suboptimal solution with a mathematically rigorous guarantee of proximity to optimality), on-line algorithms, multivariate algorithmic analysis (where the running time can be described by parameters of the input other than purely its size) and multi-core computation. After completion of this course, students will be familiar with the complexities and difficulties of dealing with very large datasets and will be able to address these issues with a variety for suited tools.

Prerequisites

none

KEN4254

Period 4

4 Feb 2019

5 Apr 2019

[Print course description](#)

ECTS credits:

6.0

Coordinators:

- [S.M. Kelk](#)
- [M. Mihalak](#)

Teaching methods:

PBL

Data Science & Knowledge Engineering

Dynamic Game Theory

Full course description

Dynamic Game Theory introduces the students to non-cooperative games and dynamic games. The games treated in this course are games in which the players are acting as strategic decision makers, who cannot make binding agreements to achieve their goals. Therefore each player may apply treats to establish a stable outcome. Such games have strong relations with population dynamics, which will be used as an example in this course. The following games will be, among others, treated in this course: matrix and bi-matrix games, repeated games, Stackelberg games, differential games, specific models of stochastic games and evolutionary games. Students will learned the essential solution concepts like “value” and “optimal strategies” for zero sum games, and the concept of “equilibrium” for non-zero sum games. In the evolutionary models the “evolutionary stable strategy” and the “replicator dynamics” will be examined. After completing this course the student will be able to apply their knowledge of games in a wide variety of domains. The student will be able to judge which model is best suited to express a certain strategic situation, while keeping in mind the (im) possibilities and accuracies for deriving solutions.

Prerequisites

None.

Recommended reading

None.

KEN4251

Period 4

4 Feb 2019

5 Apr 2019

[Print course description](#)

ECTS credits:

6.0

Instruction language:

English

Coordinators:

- [F. Thuijsman](#)
- [K. Stankova](#)

Teaching methods:

PBL

Assessment methods:

Written exam

Data Science & Knowledge Engineering

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.

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

4 Feb 2019

5 Apr 2019

[Print course description](#)

ECTS credits:

6.0

Coordinator:

- [M.J. Dumontier](#)

Data Science & Knowledge Engineering

Information Retrieval and Text Mining

Full course description

This course will provide an in-depth overview of the broad fields of Information Retrieval and Text Mining. Information retrieval is about finding material (such as documents) of an unstructured nature that satisfies an information need from within large collections. Information retrieval is about words, phrases, search terms and natural language with operators and qualifiers. 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 pattern recognition, neural networks, natural language processing, information retrieval and machine learning. An important difference with standard information retrieval 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, clinical research or and due diligence investigations. After completing this course the student will be able to adequately apply information retrieval and text-mining techniques to relevant problems.

Prerequisites

None

Recommended reading

None

KEN4153

Period 5

8 Apr 2019

7 Jun 2019

[Print course description](#)

ECTS credits:

6.0

Instruction language:

English

Coordinator:

- [J.C. Scholtes](#)

Teaching methods:

PBL

Data Science & Knowledge Engineering

Deep Learning

Full course description

Conventional machine learning techniques were limited in processing data in their raw forms and a lot of 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. Tensorflow, an open-source machine learning platform, will be introduced. In this course we will also study deep kernel based models and their connections to artificial neural network based models. 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 in Machine Learning

- [Advanced Concepts in Machine Learning](#)

Recommended reading

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

KEN4257

Period 5

8 Apr 2019

7 Jun 2019

[Print course description](#)

ECTS credits:

6.0

Data Science & Knowledge Engineering

Planning and Scheduling

Full course description

Planning and scheduling problems occur in many real-world processes. Particularly in industrial processes there are often constraints on the order in which tasks can be performed. This leads to natural, industrially-critical optimization problems. For example, a company might choose to buy many machines to process tasks 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, and machines might spend a significant amount of time standing idle, waiting for the output of other machines. In this course the various mathematical techniques for optimizing planning and scheduling problems are investigated. Various models will be investigated such as single-machine models, parallel-machine models, job-shop models. The complexity of algorithms and models will be emphasized. After completing this course, students will understand the mathematics and algorithms associated with modelling and solving planning/scheduling problems. Students will be able to understand under which circumstances different planning/scheduling problems are computationally tractable. The students will be able to solve planning and scheduling problems, using the right technique, in practice.

Prerequisites

Prerequisites: none
Desired Prior Knowledge: Some experience with optimization (e.g. linear programming) and/or design and analysis of efficient algorithms.

KEN4253

Period 5

8 Apr 2019

7 Jun 2019

[Print course description](#)

ECTS credits:

6.0

Instruction language:

English

Coordinator:

Master Artificial Intelligence

- [M. Mihalak](#)

Teaching methods:

PBL

Data Science & Knowledge Engineering

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? Computer Vision 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.

Prerequisites

none

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

8 Apr 2019

7 Jun 2019

[Print course description](#)

ECTS credits:

6.0

Coordinator:

- [S. Asteriadis](#)

Teaching methods:

Lecture(s)

MAI YR1 Core Courses

Data Science & Knowledge Engineering

Foundations of Agents

Full course description

Foundations of Agents introduces the student to the formal foundation of agents. An agent is a computational being, such as a software program, robot or human. Agents operate in some environment, which they can observe and in which they can realize objectives through the execution of actions. Examples of such environments are computer game environments, the internet and the physical world in case of robots and humans. This course addresses the problem of how an agent can act optimally in order to realize its objectives. It is investigated how the agent's environment, its objectives, actions and observations can be formalized to solve this problem. Initially Markov decision processes will be used to model the agent and its environment. Other models that will be investigated include logic-based models, such as epistemic logic, doxastic logic, dynamic logic, temporal logic, and BDI logics. The extension of these models towards multi-agent systems will be emphasized as often agents do not operate alone in an environment. After completing this course the student will be familiar with the formal models of describing agents and how these models can be used to determine and agent's behaviour. The student will be able to judge which model is applicable for specific problem domains.

Prerequisites

Knowledge of Propositional and Predicate Logic, Calculus

Recommended reading

None

KEN4115

Period 1

3 Sep 2018

26 Oct 2018

[Print course description](#)

ECTS credits:

6.0

Instruction language:

English

Coordinator:

- [N. Roos](#)

Teaching methods:

PBL

Assessment methods:

Written exam

Data Science & Knowledge Engineering

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) Game design. Game quality metrics; self-play evaluation; procedural content generation (PCG); evolutionary game design; novelty search.

Prerequisites

Good programming skills are required.

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

3 Sep 2018

26 Oct 2018

[Print course description](#)

ECTS credits:

6.0

Instruction language:

English

Coordinators:

- [J.W.H.M. Uiterwijk](#)
- [M.H.M. Winands](#)

Master Artificial Intelligence

Teaching methods:

PBL

Assessment methods:

Written exam, Assignment

Data Science & Knowledge Engineering

Multi-Agent Systems

Full course description

Multi-Agent Systems introduces the student to systems composed of multiple interacting intelligent agents. An agent is a computational being, such as a software program, robot or human. Agents operate in some environment, which they can observe and in which they can realize objectives through the execution of actions. Multi-agent systems are an enabling technology for applications that rely on distributed and parallel processing of data, information and knowledge in complex computing environments. Due to advances 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. This course covers the key conceptual, theoretical and practical foundations of multi-agent systems. The following topics, among others, will be discussed: agent-agent communication, automated negotiation and argumentation in cooperative and competitive settings, multi-agent learning and planning, automated decision making based on voting and auctioning, and development and engineering of agent-based systems. After completing this course student will be familiar with the underlying theory of agents, and multi-agent systems. The student will be able to judge whether a multi-agent approach is beneficial to use over other approaches for handling the same problem.

Prerequisites

Basic logic, basic mathematics, probability theory and Java programming.

Recommended reading

Multiagent Systems Algorithmic, Game-Theoretic, and Logical Foundations” by Yoav Shoham and Kevin Leyton-Brown. Cambridge University Press, 2009. This book is also available as ebook: <http://www.masfoundations.org/>.

KEN4111

Period 2

29 Oct 2018

21 Dec 2018

[Print course description](#)

ECTS credits:

6.0

Instruction language:

English

Coordinator:

- [G.B. Weiss](#)

Teaching methods:

Master Artificial Intelligence

PBL

Assessment methods:

Written exam

Data Science & Knowledge Engineering

Advanced Concepts in Machine Learning

Full course description

In Advanced Concepts in Machine Learning a selected number of recent developments in the field are presented and experimented with. Machine learning deals with the prediction of labels or real values for unseen objects, based on a set of previously encountered examples, or automatically adapting behaviour according to previous experience. In the past years, topics such as Deep Neural Networks, Recommender Systems, Relational Learning, Reinforcement Learning, Support Vector Machines and Gaussian Processes have made their appearance in the course. Besides an overview of recent Machine Learning techniques, the course also highlights the importance of representations in successful applications of machine learning. In this context, propositional representations are contrasted with multi-instance and relational representations, but also automatically generated representations through Sparse Coding, Auto encoders, Deep Belief Nets and indirect representations such as distances and kernel receive substantial attention. After completion of this course, the students are able to select most suited representations and best-fits of learning techniques for a given machine learning problem and reason about the limitations of the suggested selections.

Prerequisites

Desired Prior Knowledge: Familiarity with the basics of machine learning through a Machine Learning or Data Mining course

Recommended reading

Recommended literature: 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

29 Oct 2018

21 Dec 2018

[Print course description](#)

ECTS credits:

6.0

Coordinator:

- [K. Driessens](#)

Teaching methods:

PBL

Data Science & Knowledge Engineering

Research Project MAI 1

Full course description

The research project will take place in the three blocks of the first semester of the first year of the Master programme. The emphasis in the first two blocks is on exploration and modelling, and in block 3 on implementation and experimenting. The subject of the project will be closely related to the courses in the semester. The research project will be performed in small groups (3 - 5 students). During the project the student will practice team-based research and a range of scientific skills. The project will result in a project presentation, a project report, and possibly a product. At the end of each block, a project presentation will take place. Examination: Project presentation, report, and product.

Prerequisites

No Prerequisites.

Recommended reading

None.

KEN4130

Semester 1

3 Sep 2018

1 Feb 2019

[Print course description](#)

ECTS credits:

6.0

Instruction language:

English

Coordinators:

- [J.J.M. Derks](#)
- [M.H.M. Winands](#)

Teaching methods:

PBL

Assessment methods:

Written exam

Data Science & Knowledge Engineering

Autonomous Robotic Systems

Full course description

Autonomous Systems introduces the students to the foundation of situated autonomous robots. The course will start with an introduction to the field of mobile robots. Concepts such as recursive state estimation and Kalman filtering will be introduced. Additionally basic questions of how to effectively control a mobile robot will be addressed. The core of this course will address the problems of

Master Artificial Intelligence

localization, planning and control, perception and robot motion and navigation. Planning and control will be approached from a probabilistic perspective in the form of Markov Decision Processes, and Partial Observable Markov Decision processes. Robot navigation on the other hand will be approached from an evolutionary perspective where the focus will be on swarm intelligence, genetic algorithms, and neural networks. This course will be accompanied by a large practical part in which students have the opportunity to apply the learned material in practice. After completing this course, students will have a good understanding of the major concepts in autonomous systems such as localization, planning and control. The student will be able to apply the learned concepts to real autonomous systems.

Prerequisites

Discrete Mathematics, Linear Algebra, Probabilities and Statistics, Data Structures and Algorithms

Recommended reading

Some papers to be announced at Eleum.

KEN4114

Period 4

4 Feb 2019

5 Apr 2019

[Print course description](#)

ECTS credits:

6.0

Instruction language:

English

Coordinators:

- K.P. Tuyls
- [R. Möckel](#)

Teaching methods:

PBL

Assessment methods:

Written exam

Data Science & Knowledge Engineering

Research Project MAI 2

Full course description

The research project will take place in the three blocks of the second semester of the first year of the Master programme. The emphasis in the first two blocks is on exploration and modelling, and in block 6 on implementation and experimenting. The subject of the project will be closely related to the courses in the semester. The research project will be performed in small groups (3 - 5 students). During the project the student will practice team-based research and a range of scientific skills. The project will result in a project presentation, a project report, and possibly a product. At the end of each block, a project presentation will take place. Examination: Project presentation, report, and product.

Master Artificial Intelligence

Prerequisites

No Prerequisites.

Recommended reading

None.

KEN4131

Semester 2

4 Feb 2019

5 Jul 2019

[Print course description](#)

ECTS credits:

6.0

Instruction language:

English

Coordinators:

- [N. Roos](#)
- [E.N. Smirnov](#)

Teaching methods:

PBL

Data Science & Knowledge Engineering

Research Project Master's Programme

Full course description

Research Project Master's Programme

KEN4330

[Print course description](#)

ECTS credits:

6.0

Coordinator:

- [K. Driessens](#)

Second year courses

MAI YR2

Data Science & Knowledge Engineering

Electives

KEN4170

Semester 1

3 Sep 2018

1 Feb 2019

Master Artificial Intelligence

[Print course description](#)

ECTS credits:

6.0

Coordinator:

- [K. Driessens](#)

Data Science & Knowledge Engineering

Master Internship

KEN4176

Year

1 Sep 2018

31 Aug 2019

[Print course description](#)

ECTS credits:

30.0

Instruction language:

English

Coordinator:

- [K. Driessens](#)

Data Science & Knowledge Engineering

Master's Thesis AI

Full course description

The Master Artificial Intelligence will be completed by writing a master thesis. The thesis is produced individually and is the result of a master research project that runs during the second semester of year 2 of the master 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 thesis. The master project is completed by a presentation of the results. The master project will be supervised by one of the senior researchers. Examination: Master thesis and presentation.

KEN4160

Year

1 Sep 2018

31 Aug 2019

[Print course description](#)

ECTS credits:

30.0

Instruction language:

English

Coordinators:

- [J.J.M. Derks](#)
- [K. Driessens](#)

Teaching methods:

Master Artificial Intelligence

PBL

Assessment methods:

Written exam

MAI YR2 Electives

Data Science & Knowledge Engineering

Signal and Image Processing

Full course description

Signal and Image processing offers the student a hands-on introduction into the area of digital signal and image processing. The course is started with the fundamental concepts and mathematical foundation of signal and image processing. This includes a brief review of Fourier analysis, z-transforms and digital filters. After this review, classical filtering from a linear systems perspective is discussed. This is followed by wavelet transforms and principal component analysis. For each of these techniques it is discussed how they relate to a variety of objectives, such as detection, noise removal, compression, prediction, reconstruction and feature extraction. The techniques in this course are explained for both signal and image processing. During the lectures several practical cases from biomedical engineering are highlighted. The course is accompanied by practical classes in which the students will apply the techniques discussed in the lectures using the software package Matlab. After completing this course the student will be familiar with the fundamental concepts of signal and image processing and their mathematical foundation. The student will be familiar with various types of filters and their properties. The student will be able to apply the various techniques discussed in the lectures to real life problems.

Prerequisites

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

Recommended reading

Wavelets and Filter Banks, Gilbert Strang and Truong Nguyen, Wellesley College; 2nd edition, ISBN-13: 978-0961408879 A Wavelet Tour of Signal Processing, Third Edition: The Sparse Way, Stephane Mallat, Academic Press; 3rd edition, ISBN-13: 978-0123743701 Wavelet Theory: An Elementary Approach with Applications, David K. Ruch, Patrick J. Van Fleet, Wiley, ISBN: 978-0-470-38840-2.

KEN4222

Period 1

3 Sep 2018

26 Oct 2018

[Print course description](#)

ECTS credits:

6.0

Instruction language:

English

Master Artificial Intelligence

Coordinators:

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

Teaching methods:

PBL

Assessment methods:

Written exam

Data Science & Knowledge Engineering

Mathematical Optimization

Full course description

Mathematical Simulation is concerned with the study of processes and systems. When modelling a process or system there is often an uncertainty factor present. Such uncertainty is often caused by the random nature of certain factors that affect the process/system. In order to properly analyse a model it is important to correctly model any uncertainty that is present. Once the right model is in place various scenarios can be simulated, using Monte Carlo simulation, to gain insight. The results of such analyses need to be properly interpreted and uncertainty has to be reduced. The modelling, implementation, analysis and technical aspects will all be discussed in this course. The emphasis will be on discrete event simulation. After completing this course the students will be familiar with the essentials of simulation, such as the model cycle, discrete event simulation, output analysis and experimental design. Students will be able to employ simulation as a tool for evaluation.

Prerequisites

Calculus (1&2), 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

3 Sep 2018

26 Oct 2018

[Print course description](#)

ECTS credits:

6.0

Instruction language:

English

Coordinator:

- [P.J. Collins](#)

Teaching methods:

Stochastic Decision Making

Full course description

Stochastic Decision making introduces the student to modelling dynamic processes that involve randomness. Any realistic model of a real-world phenomenon must take into account the possibility of randomness. That is, the quantities one is interested in will not be predictable in advance but instead will exhibit an inherent variation that should be taken into account by the model. This is usually accomplished by allowing the model to be probabilistic in nature. Such a model is referred to as a probability model. In this course the following topics, among others, are discussed: basic concepts of probability theory, probability distribution functions, conditional probability, expectation and probability conditioning, Markov chains, Markov decision problems, Poisson processes and continuous time Markov chains. These topics are accompanied by a discussion on their mathematical framework. After completing this course the student will have obtained knowledge of modelling dynamic processes that involve randomness. This includes knowledge about appropriate probability distributions, analysis tools and knowledge of the most relevant and applicable processes. The student will be able to model and analyse all kind of real life practical situations involving stochastic uncertainty.

Prerequisites

None.

Recommended reading

None.

KEN4221

Period 1

3 Sep 2018

26 Oct 2018

[Print course description](#)

ECTS credits:

6.0

Instruction language:

English

Coordinator:

- [G.M. Schoenmakers](#)

Teaching methods:

PBL

Assessment methods:

Written exam

Data Science & Knowledge Engineering

Data Mining

Full course description

Data mining is a major frontier field of computer science that studies the extraction of useful and interesting patterns from large collections of data. This course consists of a step-by-step introduction to data-mining systems. This includes a discussion of the process of acquiring raw data, as well as several pre-processing techniques. Several data-mining techniques are discussed, varying from basic models to state-of-the-art techniques. For each technique various characteristics will be highlighted which help one decide which technique to use. Several evaluation criteria will be discussed which help one decide whether the data-mining system is producing useful patterns. The lectures and labs will emphasize the practical use of the presented techniques and the problems of developing real data-mining applications. A number of real data sets will be analysed and discussed. After completing this course students will have obtained a preliminary methodological and theoretical bases for studying and applying data mining techniques to large collections of data.

Prerequisites

None.

Recommended reading

I.H. Witten and E. Frank (2005). Data Mining: Practical Machine Learning Tools and Techniques (Second Edition), Morgan Kaufmann, June 2005, ISBN 0-12-088407-0 T. Mitchell (1997). Machine Learning, McGraw-Hill, ISBN 0-07-042807-7.

KEN4113

Period 1

3 Sep 2018

26 Oct 2018

[Print course description](#)

ECTS credits:

6.0

Instruction language:

English

Coordinator:

- [E.N. Smirnov](#)

Teaching methods:

PBL

Data Science & Knowledge Engineering

Model Identification and Data Fitting

Full course description

Model Identification and Data Fitting is centred around the estimation of a mathematical model based on previous observations. This course is devoted to the various practical and theoretical

Master Artificial Intelligence

aspects of such estimations (identifications) of mathematical models from several model classes. The course starts by addressing distance measures, norms, and criterion functions. After this the prediction error identification of linear regression models will be discussed. The emphasis will be on the various interpretations of these models such as deterministic, stochastic with Gaussian white noise and maximum likelihood estimation, stochastic in a Bayesian estimation context. Additionally several numerical implementation aspects such as: recursion, numerical complexity, numerical conditioning, and square root filtering will be highlighted. Next the focus will be on identification within the class of auto-regressive dynamical models, to which the Levinson algorithm applies. Other topics that will be discussed include identifiability, model reduction and model approximation. Several of the techniques that will be discussed are illustrated in Matlab. After completing this course the student will have obtained insight into the various aspects that play a key role in building a mathematical model from measurement data. The student will be able to apply the techniques learned to real world problems to construct models from observed data. The student will be able to judge and predict the quality of such models.

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

29 Oct 2018

21 Dec 2018

[Print course description](#)

ECTS credits:

6.0

Instruction language:

English

Coordinators:

- [R.L.M. Peeters](#)
- [P. Bonizzi](#)

Teaching methods:

PBL

Assessment methods:

Written exam

Data Science & Knowledge Engineering

Applications of Image Video Processing

Full course description

Applications of Image Video Processing

KEN4244

Master Artificial Intelligence

Period 2

29 Oct 2018

21 Dec 2018

[Print course description](#)

ECTS credits:

6.0

Coordinator:

- [A. Briassouli](#)

Data Science & Knowledge Engineering

Algorithms for Big Data

Full course description

Algorithms for Big Data presents an algorithmic toolkit to efficiently deal with the challenges that the ever growing amount of data pose. For example, the data may not fit into 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 inputs. 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. Standard algorithmic solutions do not address these challenges, and new algorithmic techniques are needed. This course looks at a number of algorithmic responses to these problems, such as algorithms with linear or sublinear running time, algorithms where the data arrives in a stream, and 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). It also covers a number of topics from classical algorithm design that have undiminished relevance in the era of big data, such as approximation algorithms (obtaining a suboptimal solution with a mathematically rigorous guarantee of proximity to optimality), on-line algorithms, multivariate algorithmic analysis (where the running time can be described by parameters of the input other than purely its size) and multi-core computation. After completion of this course, students will be familiar with the complexities and difficulties of dealing with very large datasets and will be able to address these issues with a variety of suited tools.

Prerequisites

none

KEN4254

Period 4

4 Feb 2019

5 Apr 2019

[Print course description](#)

ECTS credits:

6.0

Coordinators:

- [S.M. Kelk](#)
- [M. Mihalak](#)

Dynamic Game Theory

Full course description

Dynamic Game Theory introduces the students to non-cooperative games and dynamic games. The games treated in this course are games in which the players are acting as strategic decision makers, who cannot make binding agreements to achieve their goals. Therefore each player may apply treats to establish a stable outcome. Such games have strong relations with population dynamics, which will be used as an example in this course. The following games will be, among others, treated in this course: matrix and bi-matrix games, repeated games, Stackelberg games, differential games, specific models of stochastic games and evolutionary games. Students will learned the essential solution concepts like “value” and “optimal strategies” for zero sum games, and the concept of “equilibrium” for non-zero sum games. In the evolutionary models the “evolutionary stable strategy” and the “replicator dynamics” will be examined. After completing this course the student will be able to apply their knowledge of games in a wide variety of domains. The student will be able to judge which model is best suited to express a certain strategic situation, while keeping in mind the (im) possibilities and accuracies for deriving solutions.

Prerequisites

None.

Recommended reading

None.

KEN4251

Period 4

4 Feb 2019

5 Apr 2019

[Print course description](#)

ECTS credits:

6.0

Instruction language:

English

Coordinators:

- [F. Thuijsman](#)
- [K. Stankova](#)

Teaching methods:

PBL

Assessment methods:

Written exam

Data Science & Knowledge Engineering

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.

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

4 Feb 2019

5 Apr 2019

[Print course description](#)

ECTS credits:

6.0

Coordinator:

- [M.J. Dumontier](#)

Data Science & Knowledge Engineering

Information Retrieval and Text Mining

Full course description

This course will provide an in-depth overview of the broad fields of Information Retrieval and Text Mining. Information retrieval is about finding material (such as documents) of an unstructured nature that satisfies an information need from within large collections. Information retrieval is about words, phrases, search terms and natural language with operators and qualifiers. 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 pattern recognition, neural networks, natural language processing, information retrieval and machine learning. An important difference with standard information retrieval 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, clinical research or and due diligence investigations. After completing this course the

Master Artificial Intelligence

student will be able to adequately apply information retrieval and text-mining techniques to relevant problems.

Prerequisites

None

Recommended reading

None

KEN4153

Period 5

8 Apr 2019

7 Jun 2019

[Print course description](#)

ECTS credits:

6.0

Instruction language:

English

Coordinator:

- [J.C. Scholtes](#)

Teaching methods:

PBL

Data Science & Knowledge Engineering

Deep Learning

Full course description

Conventional machine learning techniques were limited in processing data in their raw forms and a lot of 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. Tensorflow, an open-source machine learning platform, will be introduced. In this course we will also study deep kernel based models and their connections to artificial neural network based models. 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 in Machine Learning

- [Advanced Concepts in Machine Learning](#)

Recommended reading

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

KEN4257

Period 5

8 Apr 2019

7 Jun 2019

[Print course description](#)

ECTS credits:

6.0

Data Science & Knowledge Engineering

Planning and Scheduling

Full course description

Planning and scheduling problems occur in many real-world processes. Particularly in industrial processes there are often constraints on the order in which tasks can be performed. This leads to natural, industrially-critical optimization problems. For example, a company might choose to buy many machines to process tasks 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, and machines might spend a significant amount of time standing idle, waiting for the output of other machines. In this course the various mathematical techniques for optimizing planning and scheduling problems are investigated. Various models will be investigated such as single-machine models, parallel-machine models, job-shop models. The complexity of algorithms and models will be emphasized. After completing this course, students will understand the mathematics and algorithms associated with modelling and solving planning/scheduling problems. Students will be able to understand under which circumstances different planning/scheduling problems are computationally tractable. The students will be able to solve planning and scheduling problems, using the right technique, in practice.

Prerequisites

Prerequisites: none
Desired Prior Knowledge: Some experience with optimization (e.g. linear programming) and/or design and analysis of efficient algorithms.

KEN4253

Period 5

8 Apr 2019

7 Jun 2019

[Print course description](#)

ECTS credits:

Master Artificial Intelligence

6.0

Instruction language:

English

Coordinator:

- [M. Mihalak](#)

Teaching methods:

PBL

Data Science & Knowledge Engineering

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? Computer Vision 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.

Prerequisites

none

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

8 Apr 2019

7 Jun 2019

[Print course description](#)

ECTS credits:

6.0

Coordinator:

- [S. Asteriadis](#)

Teaching methods:

Lecture(s)

Data Science & Knowledge Engineering

Symbolic Computation and Control

Full course description

Computation and Control consists of two interrelated parts. In the first part the basic techniques for the digital control of linear dynamical systems using feedback will be discussed. For such systems stability will be addressed as well as the technique of pole placement by state feedback to solve the regulation problem. This basic framework will be extended by introducing state observers in order to solve the regulation problem by output feedback. Next the course will focus on tracking problems. Here the aim is to have a system follow a trajectory, rather than just stabilize, which is achieved by extending the earlier set-up for output feedback. The first part of the course is concluded by a discussion on optimal control, which make use of a utility function to determine optimality, and which can widely be applied in practice. The second part of this course focusses on computational issues that are related to the field of systems and control. Using modern software packages such as Maple and Mathematica, the area of symbolic computation will be explored. It will be demonstrated how speed, efficiency and memory usage considerations often lead to surprising and fundamentally different algorithmic solutions in a symbolic or exact context. Examples of applications involve: stability of linear systems, model approximations. After completing this course, students will be able to design state feedback and output feedback control laws for regulation and tracking problems. Students will be able to solve computation problems involving free parameters with adequate algorithms, avoiding pitfalls from approximate numerical computation.

Prerequisites

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 2019

7 Jun 2019

[Print course description](#)

ECTS credits:

6.0

Instruction language:

English

Coordinator:

- [R.L.M. Peeters](#)

Teaching methods:

PBL

Assessment methods:

Written exam