Università
della
Svizzera
italiana

**Faculty
of
Informatics**

Master of Science in Artificial Intelligence

2017/18

Artificial Intelligence.

Artificial Intelligence (AI) is one of the most popular areas in computer science and engineering. AI deals with intelligent behavior, learning, and adaptation in machines, robots and body-less computer programmes. AI is everywhere: search engines use it to improve answers to queries, to recognize speech, to translate languages, email programmes use it to filter spam, banks use it to predict exchange rates and stock markets, doctors use it to recognize tumors, robots use it to localize themselves and obstacles, autonomous cars use it to drive, video games use it to enhance the player's experience, adaptive telescopes use it to improve image quality, smartphones use it to recognize objects / faces / gestures / voices / music, etc. People are discussing the possibility of super-intelligence and AI risks. Big players such as Google, Amazon, Baidu, Microsoft etc are investing billions in AI, and the AI-related job market is growing extremely rapidly.

Awarded Degree
Master of Science in Artificial Intelligence
Application Deadline
April 30th / June 30th depending on the nationality of the applicant.
Tuition fees per semester
Residents CHF 2'000.– / international CHF 4'000.–
Duration
4 semesters (2 years) - 120 ECTS
Scholarships
Fondazione per le Facoltà di Lugano
10 study grants for Faculty of Informatics, covers first year of tuition, renewable according to grade
Contacts/information
www.mai.usi.ch
studyadvisor@usi.ch

In this exciting context the first AI master in Switzerland is offered in Lugano, profiting from the competences of the Faculty of Informatics and the Swiss AI Lab, IDSIA, Dalle Molle Institute for Artificial Intelligence, a common institute with SUPSI and one of the world's leading research institutes in this field. For example, in 2016, IDSIA got the Swiss Special ICT award for its bio-inspired research activities and one of the ten NVIDIA "Pioneers in AI research" awards.

Goals and contents

Artificial Intelligence may not only be the most exciting field in computer science, but of science in general. In fact, the best scientists of the future might even be AIs themselves. Hardware soon will have more raw computational power (CP) than human brains, since CP per cent is still growing by a factor of 100-1000 per decade. And there is no reason to believe that general problem solving software similar to that of humans will be lacking: there already exist mathematically optimal (though not yet practical) universal problem solvers developed at IDSIA. And existing highly practical (but not quite as universal) AI already learn from experience, outperforming humans in more and more fields. For example, biologically plausible deep / recurrent artificial neural networks are learning to solve pattern recognition tasks that seemed infeasible only 10 years ago. Examples: images, handwriting, traffic signs, since 2011 even with superhuman performance - no end in sight. Even creativity has been formalized such that it can now be implemented on machines. The current developments in IS may soon lead to the end of history as we know it (more), and as an IS master student you can become part of this revolution. Artificial Intelligence systems have knowledge, beliefs, preferences and goals, and they have informational as well as motivational attitudes. They observe, learn, communicate, plan, anticipate and commit. They are able to reason about other systems and their own internal states, to simulate and optimize their performance. AI systems react to dynamic situations adapting their capabilities through learning mechanisms, with a high degree of autonomy.

Language

This programme is entirely held in English. Applicants who are not native English speaker or whose first degree was not taught in English, must supply an internationally recognised certificate to demonstrate a C1 level on the Common European Framework of Reference for language learning (CEFR).

Student profile and admission requirements

Bachelor's degree granted by a recognised university in the field of Computer Sciences or related disciplines.
Further information for applicants graduating from a University of Applied Sciences is available online:

www.mai.usi.ch/admission

Career opportunities

Students graduating from this programme will develop a taste for working on complex problems. In their future careers they will be able to apply their knowledge in many interdisciplinary areas including robotics, business forecasting, intelligent search, video games, music and entertainment, chat bots, medical diagnostics, self-driving cars, to name a few.

Contacts

USI Università della Svizzera italiana
Study Advisory Service
+41 58 666 4795
studyadvisor@usi.ch

Study programme

In this master programme a wide variety of techniques will be taught, including intelligent robotics, artificial deep neural networks, machine learning, meta-heuristics optimization techniques, data mining, data analytics, simulation and distributed algorithms. The main courses are integrated with laboratory works where students have the possibility to use real robots and to practice with state of the art tools and methodologies. After the first few lectures of the basic Machine Learning course, AI master students will already know how to train self-learning artificial neural networks to recognize the images and handwritings to the right better than any other known method.

Please be aware that slight changes in the study programme may occur.

First semester	Core Courses 18.0	Machine Learning	6.0
		Deep Learning lab	3.0
		Algorithms & Complexity	6.0
		Numerical Algorithms	3.0
	Electives 12.0	Advanced Networking	6.0
		Cyber-security	3.0
		High-Performance Computing	6.0
		Introduction to Partial Differential Equations	6.0
		Mobile Computing	6.0
		Simulation & Data Sciences Seminar	3.0
		User Experience Design	6.0
Second semester	Core Courses 24.0	Computer Vision & Pattern Recognition	6.0
		Data Analytics	6.0
		Stochastic Methods	6.0
		Robotics	6.0
	Electives 6.0	Advanced Computer Architectures	6.0
		Business Intelligence and Applications	6.0
		CPS-Intelligence	6.0
		Geometric Algorithms	6.0
		Multiscale Methods	6.0
		Quantum Computing	6.0
		Software Atelier: Simulation, Data Science & Supercomputing	6.0
Third semester	Core Courses 21.0	Artificial Intelligence	6.0
		Distributed Algorithms	6.0
		Master Thesis	9.0
	Electives 9.0	Choose from the electives of the 1st semester	
Fourth semester	Core Courses 24.0	Geometric Deep Learning	3.0
		Master Thesis	21.0
	Electives 6.0	Choose from the electives of the 2nd semester	

First semester

Core Courses

Machine Learning

Introductory Master's Course to Intelligent Systems (IS) or Artificial Intelligence (AI), taught by award-winning experts of the Swiss AI Lab IDSIA, and USI. The focus is on Machine Learning (ML). According to Computer World (2009), expertise in ML is the top skill sought by IT employers. Today ML is everywhere: search engines use it to improve answers to queries, email programmes use it to filter spam, banks use it to predict exchange rates and stock markets, doctors use it to recognize tumors, robots use it to localize themselves and obstacles, video games use it to enhance the player's experience, smartphones use it to recognize objects / faces / gestures / voices / music, etc. After the first few lectures of the basic IS course on ML, IS master students will already know how to train self-learning artificial neural networks to recognize images and handwriting better than any other known method. They will rapidly gain familiarity with state-of-the-art algorithms developed at IDSIA and other AI labs.

Deep Learning Lab

This course offers an opportunity to acquire practical experience in using TensorFlow to implement feedforward and recurrent neural networks (e.g., long short-term memory networks). Such networks currently achieve state-of-the-art results in many exciting tasks, such as object recognition, speech recognition, language translation, and learning to play games from experience. TensorFlow is an open source library developed by Google, and currently is perhaps the most popular choice among researchers and practitioners. The course will consist in practical assignments and a final course project, possibly proposed by the student.

Algorithms & Complexity

Algorithms are fundamental to computer science. They are the essence of computer programmes and lie at the core of any software system. This course will cover fundamental techniques for designing efficient computer algorithms, proving their correctness, and analyzing their performance. It will also cover a variety of application problems that use these techniques. The contents include greedy algorithms, divide and conquer algorithms, dynamic programming, network flow, NP completeness and computational intractability, approximation algorithms, and randomized algorithms. Techniques on algorithm design and analyisis will be developed by drawing on problems from across many areas of computer science and related fields.

Numerical Algorithms

This course is about the key numerical algorithms that you should really want to know about. How do TrueType fonts work? What is the secret of Google's success? Why is JPEG compression so efficient? The answers to these questions are clever numerical algorithms, based on Bézier curves, eigenvalues, and the discrete cosine transformation, respectively. We will be able to understand and discuss them once we have gone through some preliminary basics, including Newton's method for finding roots, polynomial interpolation, direct and iterative methods for solving linear systems of equations, and Gaussian quadrature. This course refreshes your basic math skills in calculus and linear algebra and shows how to utilize them for solving several real-world problems, like the ones mentioned earlier. We also provide references to the history of these solutions, going back to Newton, Leibniz, Euler, Gauss and others.

Electives

Advanced Networking

This course covers advanced topics in computer networks, with a blend of theoretical and practical topics. On the theoretical side, the syllabus will cover mathematical foundations of networking, including discussions of queuing theory, control theory, information theory, and optimization. On the practical side, the syllabus will cover concepts and designs related to modern network architectures and technologies (e.g., data-center networks, software-defined networks), protocols (e.g., SPDY, HTTP/2, IPSec), and services (e.g., Zookeeper, DHTs). Students will gain hands-on experience with topics discussed in class through a series of exercises using network simulators and emulators.

Cyber-security

The course provides an introduction to modern applied cryptography. Students will develop an understanding of the different types of cryptographic algorithms and the security services that can be realized with them. Representatives of widely used cryptographic algorithms will be introduced, and their implementation behavior will be discussed.

High-Performance Computing

Are you interested in using Europe's faster supercomputers (and getting ECTS credit points for doing so)? Would you like to learn how to write programmes for parallel supercomputers, such as a Cray or a cluster of Graphics Processing Units? The course is designed to teach students how to programme parallel computers to efficiently solve challenging problems in science and engineering, where very fast computers are required either to perform complex simulations or to analyze enormous datasets. It covers basic principles, architectures, and algorithms of parallel systems. The course is structured in four parts: (i) foundations of parallel systems, (ii) basic parallel algorithm, (iii) parallel programming, and (iv) parallel applications.

Introduction to Partial Differential Equations

Many phenomena occurring in real life applications (i.e. physics, finance, biology) are modeled by means of partial differential equations (PDEs) . These mathematical models are sets of differential equations, which describe the essential behavior of a natural or artificial system, in order to forecast and control its evolution The aim of the course is twofold. Firstly, we will give the students an overview on the construction of differential PDEs for basic physical applications. Then, focusing on the arising PDEs, their theoretical mathematical background will be discussed. As the understanding of PDEs is closely connected to understand their physical meaning and the qualitative and quantitative behavior of their solutions, the theoretical investigations will be accompanied by the introduction of numerical schemes, which will allow for the illustrative numerical investigation of PDEs. We will consider elliptic operators (Diffusion), parabolic (heat equation), and hyperbolic (fluid flow, advection).

Mobile Computing

Mobile devices such as mobile phones, smart watches, and other wearable devices can interact seamlessly by relying on available communication infrastructures. This course focuses on challenges and opportunities arising from the use of systems of mobile devices or "mobile sensing systems". Following an overview of applications enabled by mobile sensing systems the focus will be devoted to the most significant technologies, including hardware platforms, programming environments and tools. Relevant aspects related to the design and development of a mobile sensing system, including the handling of sensors, the design of user interfaces, the management of local and remote sensor data storage, privacy and security issues will be investigated and addressed. In order to gain practical hands-on experience, students will learn in the lab sessions how to design, implement, and demonstrate Android-based mobile sensing applications.

Simulation & Data Sciences Seminar

Data science is the profession of the future. During this course we focus on the advanced applications used to understand complex systems in broad areas including natural and physical sciences, social sciences, life sciences and management of Big Data. The students will have the opportunity to understand how to develop and apply high performance methods used to solve complex problems related to time series analysis, and modeling of real-life phenomena. A large part of the course is dedicated to data science and how to use data storage and data analysis in a smart way. Participants will learn various process discovery algorithms and the key analysis techniques in traditional model-based process analysis (e.g., simulation and other business process management techniques) and data-centric analysis techniques such as machine learning and data mining. A basic understanding of programming and statistics (undergraduate level) is assumed.

User Experience Design

This class aims at familiarising students with both the theory behind the discipline of Human Computer Interaction (HCI) and the practical process of User eXperience (UX) design. Students not only develop an awareness and appreciation of the crucial implications of good interfaces in terms of overall system performance and user satisfaction, but also learn core skills needed in order to identify user requirements, envision interfaces and processes, and evaluate competing design options. Students will work in small teams of 3-5 to drive a design project from start to finish. Core skills are introduced in hands-on classes, interspersed with lectures and discussions about the underlying theory.

Second semester

Core Courses

Computer Vision & Pattern Recognition

The purpose of the course is to introduce basic problems and notions in image processing, computer vision, and patter recognition though a common geometric framework and present some classical, industry-standard and state-of-the-art methods through this framework. The course uses tools from differential geometry, calculus of variations, and numerical optimization to address problems such as image recovery (denoising, impainting, deconvolution), filtering (adaptive diffusion, bilateral and non-local means filters), 3D structure reconstruction (shape from shading, stereo, photometric stereo); and rigid and non-rigid similarity and correspondence (iterative closest point methods, multidimensional scaling, Gromov-Hausdorff distance). The emphasis is made on both formulating a rigorous mathematical model of the problem and developing an efficient numerical method for its solution, with hands-on programming exercises that solve real-world problems.

Data Analytics

This is an applied statistics course focusing on data analysis. The course begins with an overview of how to organise, perform, and write-up data analyses. The course starts with a theoretical part on the how to mine very large datasets to get valuable data to analyse. Then it covers some of the most popular and widely used statistical methods to analyse the data, like linear regression, principal components analysis, cross-validation, and p-values. Instead of focusing on mathematical details, the lectures are designed to help you apply these techniques to real data using the R statistical programming language, interpret and visualise the results, and diagnose potential problems in your analysis.

Stochastic Methods

Many of the real-life applications (e.g., in banking/insurance, mechanics, medicine, etc.) can be only approached, modelled and computed as stochastic (or random) processes. The aim of this course is to introduce the most essential mathematical concepts and computational methods from the area of stochastic and random processes. Besides of gaining theoretical and practical backgrounds in the areas of stochastic calculus, random processes and uncertainty quantification, the participants will gain practical skills by doing supervised short research projects from real-life applications. The recurrent theme of the course is in establishing a joint stochastic/statistic perspective based on optimisation paradigm - for various computational methods and

algorithms from computational science, machine learning and informatics.

Robotics

Robotics addresses the design, construction, and automatic control of mechatronical systems. The course provides a general overview of robotics, focusing on autonomous mobile robots: autonomous systems which exist and move in the physical world, can sense their environment using multiple sensors, can reason about it to issue plans, and can act on it to achieve one or multiple goals. The fundamental concepts and models necessary to achieve such a view of a robotic system will be studied: Forward and Inverse Kinematics; Proprio- and Exteroceptive Sensing; Model-based and Model-free State Estimation; Feedback-based Control; Paradigms and Architectures for Robot Control; Localization and Mapping; Motion Planning; Navigation; Coordination and Cooperation in Multi-robots and Swarms. The course includes theory classes, hands-on classes, and homework. Students will learn how to use the ROS and the simulator Gazebo, and will apply the learned concepts through the programming of both simulated and real robots.

Electives

Advanced Computer Architectures

The course builds on previous knowledge in basic computer architecture, and visits the major techniques devised to get higher performance from a single processor, and, later on, from multiprocessors. It describes the concepts of pipelined CPUs, cache architecture and optimization, Instruction-Level parallelism (Superscalar and VLIW architectures), Thread-Level parallelism (fine-grained, coarse-grained, simultaneous multithreading), Data-level parallelism (Vector architectures), and shared-memory multi-processing. The course also includes a project where the Simplescalar and Watch simulation tools are used to perform design-space exploration, and to understand the tradeoffs that computer architects must consider between performance and cost.

Business Intelligence and Applications

The course develops a working knowledge of the principles, architectures, and tools for Enterprise Information Management and Business Intelligence. It addresses enterprise data integration and knowledge management, data mining and business intelligence. It gives an outlook on emerging data architectures, with focus on social network structures. It also presents agile and model-driven enterprise application development, using OMG's Model Driven Architecture. The notion of model is illustrated with different modeling languages

CPS-Intelligence

With ever-growing proliferation of cyber-physical systems in all walks of life, their properties like adaptation ability, dependability, security and timeliness are of utmost importance. The built-in intelligence in cyber-physical systems will have to face these challenges to ensure problem-free, continuous operation of such systems, maximum security, effectiveness and real time operability. The course will propose intelligent mechanisms to guarantee appropriate performance within an evolving, time invariant environment, optimal harvesting and management of the residual energy, to identify faults within a model-free framework as well as solve the compromise between output accuracy and computational complexity.

Geometric Algorithms

This course is an introduction to computational geometry and its applications. It covers techniques needed in designing and analyzing efficient algorithms for computational problems in discrete geometry such as convex hulls, triangulations, geometric intersections, Voronoi diagrams, Delaunay triangulations, arrangements of lines and hyperplanes, and range searching. Computational geometry is well related to diverse application domains, where geometric algorithms play a fundamental role, such as pattern recognition, image processing, computer graphics, robotics, geographic information systems (GIS), computer-aided design (CAD), information retrieval, computational science, and many others. The course covers general algorithmic techniques, such as plane sweep, divide and conquer, incremental construction, randomization, and approximation, through their application to basic geometric problems.

Multiscale Methods

In this course, we present the state of the art for linear as well as nonlinear multilevel and multigrid methods. The solution of large linear and nonlinear systems of equations is one of the most important tasks in numerical simulation. Since standard solution methods do not scale optimally, alternative solution strategies have been developed during the last decades. In particular multilevel or multiscale solution strategies have been developed, which are often employed due to their high efficiency. Prominent examples are multilevel or domain decomposition methods for linear elliptic problems, In this course, we start from well known subspace correction methods for linear problems and proceed to more recent developments as are nonlinear multigrid and monotone multigrid. Finally, we will consider (recursive) trust-region methods and their application to minimization problems in computational mechanics. For all methods, we will also discuss their parallelization.

Quantum Computing

Followed by an introduction to the basic principles of quantum physics, such as superposition, interference, or entanglement, a variety of subjects are treated: Quantum algorithms, teleportation, quantum communication complexity and "pseudo-telepathy", quantum cryptography, as well as the main concepts of quantum information theory.

Software Atelier: Simulation, Data Science & Supercomputing

The software atelier on simulation, data science and supercomputing presents advanced topics in parallel computing and numerical simulation for prospective computational/software engineers. There will be several programming assignments to acquaint students with basic issues in memory locality and parallelism needed for high performance. Most of the grade will be based on a final project (in which students are encouraged to work in small interdisciplinary teams), which could involve parallelizing an interesting application, or developing or evaluating a novel parallel computing tool. Students are expected to have identified a likely project by mid semester, so that they can begin working on it. We will provide many suggestions of possible projects as the class proceeds.

Third semester

Core Courses

Artificial Intelligence

AI deals with intelligent behavior, learning, and adaptation in machines, robots and body-less computer programmes. The goal of this course is to investigate advanced AI models and algorithms useful to solve real-world search, constraint satisfaction, games and machine learning problems also in case of unprecise and stochastic knowledge.

Distributed Algorithms

Distributed computing systems arise in a wide range of modern applications. This course surveys the foundations of many distributed computing systems, namely, the distributed algorithms that lie at their core. The course provides the basis for designing distributed algorithms and formally reasoning about their correctness. It addresses issues related to what distributed systems can and cannot do (i.e., impossibility results) in certain system models. The course focuses on three aspects of distributed computing: system models, fundamental problems in distributed computing, and application of distributed algorithms. System models include synchronous versus asynchronous systems, communication models, and failure models. Several fundamental problems are covered, including consensus, atomic broadcast, atomic multicast, atomic commit, and data consistency. Applications of distributed algorithms target various forms of replication techniques.

Master Thesis

The Master thesis is an academic piece of work, an original contribution to the body of knowledge in artificial intelligence. Such a contribution can be theoretical or experimental, but always builds on a solid research effort, and on the use of appropriate concepts, methods, and tools acquired during the Master. Faculty members advise students during their Master's thesis work.

Electives

See First semester's electives.

Fourth semester

Core Courses

Geometric Deep Learning

In the past decade, deep learning methods have achieved unprecedented performance on a broad range of problems in various fields from computer vision to speech recognition. However, so far research has mainly focused on developing deep learning methods for Euclidean-structured data. However, many important applications have to deal with non-Euclidean structured data, such as graphs and manifolds. Such geometric data are becoming increasingly important in computer graphics and 3D vision, sensor networks, drug design, biomedicine, recommendation systems, and web applications. The adoption of deep learning in these fields has been lagging behind until recently, primarily since the non-Euclidean nature of objects dealt with makes the very definition of basic operations used in deep networks rather elusive. The purpose of the course is to introduce the emerging field of geometric deep learning on graphs and manifolds, overview existing solutions and applications for this class of problems, as well as key difficulties and future research directions. The course will be held in the form of a seminar; following an introduction by the instructor, the students will present topics related to the field.

Master Thesis

The Master thesis is an academic piece of work, an original contribution to the body of knowledge in artificial intelligence. Such a contribution can be theoretical or experimental, but always builds on a solid research effort, and on the use of appropriate concepts, methods, and tools acquired during the Master. Faculty members advise students during their Master's thesis work.

Electives

See Second semester's electives.

Università
della
Svizzera
italiana

**Faculty
of
Informatics**

Master of Science in Artificial Intelligence

2017/18