

REGULATIONS FOR THE DEGREE OF MASTER OF SCIENCE IN ARTIFICIAL INTELLIGENCE (MSc[AI])

For students admitted in 2022-2023 and thereafter

(See also General Regulations and Regulations for Taught Postgraduate Curricula)

Any publication based on work approved for a higher degree should contain a reference to the effect that the work was submitted to the University of Hong Kong for the award of the degree.

Admission requirements

MAI1 To be eligible for admission to the courses leading to the degree of Master of Science in Artificial Intelligence, a candidate

- (a) shall comply with the General Regulations and the Regulations for Taught Postgraduate Curricula;
 - (b) shall hold
 - (i) a Bachelor's degree with honours of this University, or
 - (ii) another qualification of equivalent standard from this University or another University or comparable institution acceptable for this purpose; and
 - (c) shall pass a qualifying examination if so required; and
 - (d) shall possess knowledge of linear algebra, calculus, probability theory, introductory statistics and computer programming.
-

Qualifying examination

- MAI2**
- (a) A qualifying examination may be set to test the candidate's formal academic ability or his/her ability to follow the courses of study prescribed. It shall consist of one or more written papers or their equivalent and may include a project report.
 - (b) A candidate who is required to satisfy the examiners in a qualifying examination shall not be permitted to register until he/she has satisfied the examiners in the examination.
-

Period of study

- MAI3**
- (a) The curriculum shall normally extend over one and a half academic years of full-time study or two and a half academic years of part-time study. Candidates shall not be permitted to extend their studies beyond the maximum period of registration of three academic years of full-time study or four academic years of part-time study, unless otherwise permitted or required by the Board of the Faculty.
 - (b) Candidates of full-time study may be permitted to complete the curriculum in one academic year, subject to the approval of the Programme Director and the Board of the Faculty. The candidate should write formally to apply for shortening the normative period of study via the Department within one month after admission to the curriculum. For such candidates, it is recommended that the capstone course will start in the second semester of first year with expectation to be completed in the summer semester.
-

Course Exemption and advanced standing

- MAI4**
- (a) In recognition of studies completed successfully before admission to the curriculum, advanced standing of up to 12 credits may be granted to a candidate with appropriate qualification and professional experiences, on production of appropriate certification, subject to the approval of the Board of the Faculty. The candidate should write formally to apply for advanced standing via the Department within two weeks after admission to the curriculum.
 - (b) For cases of having satisfactorily completed more than 12 credits of another course or courses equivalent in content to any of the compulsory courses as specified in the syllabuses, candidates may, on production of appropriate certification, be exempted from the compulsory course(s), subject to approval of the Board of the Faculty. Candidates so exempted must replace the number of exempted credits with electives course(s) in the curriculum of the same credit value.
-

Award of degree

- MAI5** To be eligible for the award of the degree of Master of Science in Artificial Intelligence, a candidate shall
- (a) comply with the General Regulations and the Regulations for Taught Postgraduate Curricula; and
 - (b) successfully complete the curriculum in accordance with the regulations set out below.

A candidate who fails to fulfill the requirements within the maximum (i) three academic years for full-time mode of study or (ii) four academic years for part-time mode of study shall be recommended for discontinuation under the provisions of General Regulation G12, except that a candidate is granted permission to extend period of study by the Board of the Faculty in accordance with Regulation MAI3.

Completion of curriculum

MAI6 To successfully complete the curriculum, a candidate shall satisfy the requirements prescribed in TPG 6 of the Regulations for Taught Postgraduate Curricula; follow courses of instruction; and satisfy the examiners in the prescribed courses and in any prescribed form of examination in accordance with the regulations set out below.

Assessments

- MAI7**
- (a) In any course where so prescribed in the syllabus, coursework or a project report may constitute part or whole of the examination for the course.
 - (b) The written examination for each module shall be held after the completion of the prescribed course of study for that module, and not later than January, May or August immediately following the completion of the course of study for that module.
 - (c) There shall be no appeal against the results of examinations and all other forms of assessment.

MAI8 If during any academic year a candidate has failed at his/her first attempt in a course or courses, but is not required to discontinue his/her studies by Regulation MAI10, the candidate may be permitted to make up for the failed courses in the following manner:

- (a) undergoing re-assessment/re-examination in the failed course or courses to be held before the next academic year; or
- (b) for repeating the course and re-examination in the failed course or courses in the next academic year; or
- (c) for elective courses, taking another course in lieu and satisfying the assessment requirements.

MAI9 Failure to undertake the examination of a course as scheduled shall normally result in automatic failure in that course. A candidate who, because of illness, is unable to be present at the written examination of any course may apply for permission to present himself/herself at a supplementary examination of the same course to be held before the beginning of the following academic year. Any such application shall be made on the form prescribed within seven calendar days of the examination concerned.

- MAI10** A candidate may be required to discontinue his/her studies if he/she
- (a) during any academic year has failed in half or more than half the number of credits of all the courses to be examined in that academic year; or
 - (b) has failed at a repeated attempt in any course; or
 - (c) has exceeded the maximum period of registration.

Grading

MAI11 Individual courses shall be graded according the letter grading system as determined by the Board of Examiners. The standards and the grade points for assessment are as follows:

| Grade | Standard | Grade Point |
|-------|---------------------|-------------|
| A+ | Excellent | 4.3 |
| A | | 4.0 |
| A- | | 3.7 |
| B+ | Good | 3.3 |
| B | | 3.0 |
| B- | | 2.7 |
| C+ | Satisfactory | 2.3 |
| C | | 2.0 |
| C- | | 1.7 |
| D+ | Pass | 1.3 |
| D | | 1.0 |
| F | Fail | 0 |

MAI12 On successful completion of the curriculum, candidates who have shown exceptional merit at the whole examination may be awarded a mark of distinction, and this mark shall be recorded in the candidates' degree diploma.

**SYLLABUSES FOR THE DEGREE OF
MASTER OF SCIENCE IN ARTIFICIAL INTELLIGENCE
(MSc[AI])**

The Department of Mathematics, the Department of Statistics and Actuarial Science, and the Department of Computer Science jointly offer a postgraduate curriculum leading to the degree of Master of Science in Artificial Intelligence, with two study modes: the one and a half academic years' full-time mode and the two and a half academic years' part-time mode. The curriculum is designed to provide graduates with training in the principles and practice of artificial intelligence. Candidates should possess knowledge of linear algebra, calculus, probability theory, introductory statistics and computer programming.

STRUCTURE AND EVALUATION

Each student must complete at least 72 credits of courses, split into 42 credits of core courses, 18 credits of disciplinary electives, and 12 credits of a capstone project. If a student selects a course whose contents are similar to a course (or courses) which he/she has taken in his/her previous study, the Department may not approve the selection in question.

CURRICULUM

(applicable for both full-time and part-time modes)

Compulsory Courses (42 credits)

| | |
|----------|---|
| ARIN7001 | Foundations of artificial intelligence |
| ARIN7011 | Optimization in artificial intelligence |
| ARIN7013 | Numerical methods in artificial intelligence |
| ARIN7101 | Statistics in artificial intelligence |
| ARIN7102 | Applied data mining and text analytics |
| COMP7404 | Computational intelligence and machine learning |
| DASC7606 | Deep learning |

Disciplinary Electives (18 credits)*

with at least 6 credits from each of the following lists

List A:

| | |
|----------|--|
| ARIN7014 | Topics in advanced numerical analysis |
| ARIN7015 | Topics in artificial intelligence and machine learning |
| MATH7224 | Topics in advanced probability theory |
| MATH7502 | Topics in applied discrete mathematics |
| MATH7503 | Topics in advanced optimization |

List B:

| | |
|----------|---|
| STAT6011 | Computational statistics and Bayesian learning |
| STAT7008 | Programming for data science |
| STAT8020 | Quantitative strategies and algorithmic trading |
| STAT8021 | Big data analytics |

List C:

| | |
|----------|--|
| COMP7308 | Introduction to unmanned systems |
| COMP7309 | Quantum computing and artificial intelligence |
| COMP7409 | Machine learning in trading and finance |
| COMP7502 | Image processing and computer vision |
| ARIN7017 | Legal issues in artificial intelligence and data science |

* Students who have completed the same or similar courses in their previous studies may, on production of relevant transcripts, be permitted to select up to 18 credits of disciplinary electives from the other two lists if they are not able to find any untaken options from any one of the lists of disciplinary electives.

Capstone Project (12 credits)

| | |
|----------|--|
| ARIN7600 | Artificial intelligence project (12 credits) |
|----------|--|

All courses should be 6-credit bearing unless otherwise stated.

COURSE DESCRIPTION

Compulsory Courses

ARIN7001 Foundations of artificial intelligence (6 credits)

This course introduces foundational knowledge, methods and tools in mathematics, statistics and computer science for the purpose of studying and applying artificial intelligence.

Prerequisites: Nil

Assessment: coursework (50%) and examination (50%)

ARIN7011 Optimization in artificial intelligence (6 credits)

This course introduces students to the topics in theory and algorithms of optimization that play important roles in artificial intelligence and machine learning. Topics include: 1) Fundamental optimization models in AI (linear programming models, integer programming models, network models, reinforcement learning and deep learning models, etc.); 2) Optimization theory in AI (optimality conditions, constraint qualification, global landscape analysis of deep neural networks, P- and NP-hard problems, approximation algorithms, preliminary graph theory, etc.), 3) Optimization algorithms in AI: (a) Classic algorithms (simplex method, interior point method, branch and bound method, cutting plane method, representative algorithms, gradient type methods, CG methods, projection methods, penalty method, Lagrange methods, quasi-Newton methods, Newton type methods), (b) Stochastic algorithms (stochastic gradient descent (SGD), stochastic coordinate descent methods, subsampled Newton, stochastic quasi-Newton), (c) Algorithms for large-scale optimization problems (Operator splitting algorithms (BCD type algorithms, ADMM, primal-dual type algorithms, etc.), centralized/decentralized algorithms, etc.). (d) Algorithms for nonconvex optimization and training deep neural networks.

Prerequisites: Nil

Assessment: coursework (50%) and examination (50%)

ARIN7013 Numerical methods in artificial intelligence (6 credits)

This course introduces students to the numerical methods that are instrumental in artificial intelligence and machine learning. Topics include: 1) Notions and concepts in numerical analysis (convolution matrix (related to CNN), kernel methods, pattern analysis, direct methods for sparse matrices). 2) Numerical method for solving linear systems (Jacobi Method, Gauss-Seidel method, Cholesky decomposition, singular value decomposition (SVD), low-rank matrix approximation, with applications in artificial intelligence and machine learning). 3) Principal component analysis, tensor decomposition and their applications to computer vision, image processing and artificial intelligence and machine learning in general. 4) Compute eigenvalues and eigenvectors (Rayleigh quotient, with applications in artificial intelligence and machine learning). 5) Numerical methods for ordinary differential equations (stability, convergence analysis, relation between the SGD and Euler method, using DNN to compute ODEs).

Prerequisites: Nil

Assessment: coursework (50%) and examination (50%)

ARIN7101 Statistics in artificial intelligence (6 credits)

The development of artificial intelligence has revolutionized the theory and practice of statistical learning, while novel statistical learning approaches are becoming an integral part of artificial intelligence. By focusing on the interplay between statistical learning and artificial intelligence, this course reviews the main concepts underpinning classical statistical learning, studies computer-intensive methods for conducting statistical learning, and examines important issues concerning statistical learning drawn upon modern artificial intelligence technologies. Contents include classical frequentist and Bayesian inferences, resampling methods, large-scale hypothesis testing, regularization, and high-dimensional modeling.

Prerequisites: Nil

Assessment: coursework (40%) and examination (60%)

ARIN7102 Applied data mining and text analytics (6 credits)

With the rapid developments in computer and data storage technologies, the fundamental paradigms of classical data analysis are mature for change. Data mining aims at automated discovery of underlying structure and patterns in large amounts of data, especially text data. This course takes a practical approach to acquaint students with the new generation of data mining tools and techniques, and show how to use them to make informed decisions. Topics include data preparation, feature selection, association rules, decision trees, bagging, random forests and gradient boosting, cluster analysis, neural networks, introduction to text mining.

Prerequisites: Nil

Assessment: coursework (100%)

COMP7404 Computational intelligence and machine learning (6 credits)

This course will teach a broad set of principles and tools that will provide the mathematical, algorithmic and philosophical framework for tackling problems using Artificial Intelligence (AI) and Machine Learning (ML). AI and ML are highly interdisciplinary fields with impact in different applications, such as, biology, robotics, language, economics, and computer science. AI is the science and engineering of making intelligent machines, especially intelligent computer programs, while ML refers to the changes in systems that perform tasks associated with AI. Ethical issues in advanced AI and how to prevent learning algorithms from acquiring morally undesirable biases will be covered.

Topics may include a subset of the following: problem solving by search, heuristic (informed) search, constraint satisfaction, games, knowledge-based agents, supervised learning, unsupervised learning; learning theory, reinforcement learning and adaptive control and ethical challenges of AI and ML.

Pre-requisites: Nil, but knowledge of data structures and algorithms, probability, linear algebra, and programming would be an advantage.

Assessment: coursework (50%) and examination (50%)

DASC7606 Deep learning (6 credits)

Machine learning is a fast-growing field in computer science and deep learning is the cutting edge technology that enables machines to learn from large-scale and complex datasets. Ethical implications of deep learning and its applications will be covered and the course will focus on how deep neural networks are applied to solve a wide range of problems in areas such as natural language processing, and image processing. Other applications such as financial predictions, game playing and robotics may also be covered. Topics covered include linear and logistic regression, artificial neural networks and how to train them, recurrent neural networks, convolutional neural networks, generative models, deep reinforcement learning, and unsupervised feature learning.

Prerequisites: Basic programming skills, e.g., Python is required.

Assessment: coursework (40%) and examination (60%)

Disciplinary Electives

ARIN7014 Topics in advanced numerical analysis (6 credits)

This course covers a selection of topics in advanced numerical analysis which may include: 1) Krylov subspace, generalized minimal residual method (GMRES); 2) numerical (partial) differential equations; 3) stochastic methods and their applications to artificial intelligence and machine learning; 4) approximation theory, high-dimensional approximation (MC, QMC, sparse grid method); 5) Fourier analysis, wavelet analysis; 6) robust PCA and dimensional reduction methods. The selected topics may vary from year to year.

Prerequisites: Students should have basic knowledge in numerical analysis and scientific computing; pass in ARIN7013 Numerical methods in artificial intelligence or equivalent.

Assessment: coursework (50%) and examination (50%)

ARIN7015 Topics in artificial intelligence and machine learning (6 credits)

Selected topics in artificial intelligence that are of current interest will be discussed in this course.

Assessment: coursework (50%) and examination (50%)

MATH7224 Topics in advanced probability theory (6 credits)

Selected topics in probability theory will be discussed in this course.

Assessment: coursework (50%) and examination (50%)

MATH7502 Topics in applied discrete mathematics (6 credits)

This course aims to provide students with the opportunity to study some further topics in applied discrete mathematics. A selection of topics in discrete mathematics applied in combinatorics and

optimization (such as algebraic coding theory, cryptography, discrete optimization, etc.) The selected topics may vary from year to year.

Pre-requisites: Knowledge in introductory discrete mathematics. Students may be asked to present appropriate evidence of having met the pre-requisites for enrolling in this course.

Assessment: coursework (50%) and examination (50%)

MATH7503 Topics in advanced optimization (6 credits)

A study in greater depth of some special topics in mathematical programming or optimization. It is mainly intended for students in Operations Research or related subject areas. This course covers a selection of topics which may include convex programming, nonconvex programming, saddle point problems, variational inequalities, optimization theory and algorithms suitable for applications in various areas such as machine learning, artificial intelligence, imaging and computer vision. The selected topics may vary from year to year.

Pre-requisites: Knowledge in introductory mathematical programming and optimization. Students may be asked to present appropriate evidence of having met the pre-requisites for enrolling in this course.

Assessment: coursework (100%)

STAT6011 Computational statistics and Bayesian learning (6 credits)

This course aims to give undergraduate and postgraduate students an introduction on modern computationally intensive methods in statistics. It emphasizes the role of computation as a fundamental tool of discovery in data analysis and statistical inference, and for development of statistical theory and methods. Contents include: Bayesian statistics, Markov chain Monte Carlo methods such as Gibbs sampler, Metropolis-Hastings algorithm, and data augmentation; generation of random variables using the inversion methods, rejection sampling, the sampling/importance resampling method; optimization techniques including Newton's method, expectation-maximization (EM) algorithm and its variants, and minorization-maximization (MM) algorithm; integration including Laplace approximation, Gaussian quadrature, the importance sampling method, Monte Carlo integration, and other topics such as hidden Markov models, and Bootstrap methods. More advanced Bayesian learning methods cover approximate Bayesian computation, the Hamiltonian Monte Carlo algorithm, hierarchical models and nonparametric Bayes.

Pre-requisites: Nil

Assessment: coursework (50%) and examination (50%)

STAT7008 Programming for data science (6 credits)

In the big data era, it is very easy to collect huge amounts of data. Capturing and exploiting the important information contained within such datasets poses a number of statistical challenges. This course aims to provide students with a strong foundation in computing skills necessary to use R or Python to tackle some of these challenges. Possible topics to be covered may include exploratory data analysis and visualization, collecting data from a variety of sources (e.g. Excel, web-scraping, APIs

and others), object-oriented programming concepts and scientific computation tools. Students will learn to create their own R packages or Python libraries.

Assessment: coursework (100%)

STAT8020 Quantitative strategies and algorithmic trading (6 credits)

Quantitative trading is a systematic investment approach that consists of identification of trading opportunities via statistical data analysis and implementation via computer algorithms. This course introduces various methodologies that are commonly employed in quantitative trading.

The first half of the course focuses at strategies and methodologies derived from the data snapshotted at daily or minute frequency. Some specific topics are: (1) techniques for trading trending and mean-reverting instruments, (2) statistical arbitrage and pairs trading, (3) detection of “time-series” mean reversion or stationarity, (4) cross-sectional momentum and contrarian strategies, (5) back-testing methodologies and corresponding performance measures, and (6) Kelly formula, money and risk management. The second half of the course discusses statistical models of high frequency data and related trading strategies. Topics that planned to be covered are: (7) introduction of market microstructure, (8) stylized features and models of high frequency transaction prices, (9) limit order book models, (10) optimal execution and smart order routing algorithms, and (11) regulation and compliance issues in algorithmic trading.

Pre-requisites: Students should have basic knowledge and experience in financial data analysis.

Assessment: coursework (50%) and examination (50%)

STAT8021 Big data analytics (6 credits)

The recent explosion of social media and the computerization of every aspect of life resulted in the creation of volumes of mostly unstructured data (big data): web logs, e-mails, videos, speech recordings, photographs, tweets and others. This course aims to provide students with knowledge and skills of some advanced analytics and statistical modelling for solving big data problems. Topics include recommender system, deep learning: CNN, RNN, LSTM, GRU, natural language processing, sentiment analysis and topic modeling. Students are required to possess basic understanding of Python language.

Pre-requisites: Pass in ARIN7102 Applied data mining and text analytics or equivalent

Assessment: coursework (100%)

COMP7308 Introduction to unmanned systems (6 credits)

To study the theory and algorithms in unmanned systems. Topics include vehicle modelling, vehicle control, state estimation, perception and mapping, motion planning, and deep learning related techniques.

Assessment: coursework (50%) and examination (50%)

COMP7309 Quantum computing and artificial intelligence (6 credits)

This course offers a theoretical overview of selected topics from the interdisciplinary fields of quantum computation and quantum AI. The scope of the lectures encompasses an accessible introduction to the fundamental concepts of quantum computation. Importantly, the introduction takes the angle of computer science and logic, such that no preliminary knowledge of quantum theory is required. Thereupon, detailed comparisons of computational principles and related phenomena in the classical and quantum domain outline the stark potential and challenges of quantum theory for fundamentally novel algorithms which are more powerful than possible with conventional computers. Thereupon, the theoretical capability of quantum computers is illustrated by analyzing a selection of milestone algorithms of quantum computation, and their potential applications to artificial intelligence.

Assessment: coursework (60%) and examination (40%)

COMP7409 Machine learning in trading and finance (6 credits)

The course introduces our students to the field of Machine Learning, and help them develop skills of applying Machine Learning, or more precisely, applying supervised learning, unsupervised learning and reinforcement learning to solve problems in Trading and Finance.

This course will cover the following topics. (1) Overview of Machine Learning and Artificial Intelligence, (2) Supervised Learning, Unsupervised Learning and Reinforcement Learning, (3) Major algorithms for Supervised Learning and Unsupervised Learning with applications to Trading and Finance, (4) Basic algorithms for Reinforcement Learning with applications to optimal trading, asset management, and portfolio optimization, (5) Advanced methods of Reinforcement Learning with applications to high-frequency trading, cryptocurrency trading and peer-to-peer lending.

Assessment: coursework (65%) and examination (35%)

COMP7502 Image processing and computer vision (6 credits)

To study the theory and algorithms in image processing and computer vision. Topics include image representation; image enhancement; image restoration; mathematical morphology; image compression; scene understanding and motion analysis.

Assessment: coursework (50%) and examination (50%)

ARIN7017 Legal issues in artificial intelligence and data science (6 credits)

Modern information systems have had unprecedented impact on privacy while building dependency on them. The immense social benefits of such systems as data mining and cloud computing must be weighed against potential dangers with consideration of methods of mitigation of risk.

This course examines the growing legal, administrative, policy and technical issues associated with the use of artificial intelligence and information security and assurance. In particular, the relationship of data mining to information assurance and privacy are analyzed and legislative responses studied.

Assessment: coursework (100%)

Capstone Project

ARIN7600 Artificial intelligence project (12 credits)

The students will be required to attend an artificial intelligence ethics workshop and then carry out independent work on a major project under the supervision of staff members. A research report as well as an oral presentation on the research work and related ethics issues are required.

Assessment: research report (75%) and oral presentation (25%)
