Course Content

Course Code	Core Courses
AI6101	Introduction to AI and AI Ethics
AI6102	Machine Learning: Methodologies and Applications
AI6103	Deep Learning and Applications
AI6104	Mathematics for AI
	Bridging Course
AI6120	Python Programming
	Elective Courses
AI6121	Computer Vision
AI6122	Text Data Management and Processing
AI6123	Time Series Analysis
AI6124	Neuro Evolution and Fuzzy Intelligence
AI6125	Multi-Agent System
AI6126	Advanced Computer Vision
AI6127	Deep Neural Networks for Natural Language Processing
AI6128	Urban Computing
AI6129	AI Master Project

*Note: Not all courses listed in the curriculum will be offered in a semester. Courses offered are subject to availability of instructors and resources.

AI6101 Introduction to AI and AI Ethics

This course introduces the basics of artificial intelligence (AI), paying special attention to the ethical issues involved. It offers an interdisciplinary overview of AI to help students start their learning journey in this field. Well established paradigms pertaining to the formal representations of artificial intelligence and knowledge engineering for computational problem solving will be introduced.

The main topics covered include:
1. The foundation of AI, history, and recent progress
2. Search and intelligent agents
3. Reinforcement Learning
4. Computational game theory
5. The philosophical foundations of AI ethics
6. Techniques for incorporating ethical considerations into AI systems
7. Example applications of ethical AI – A

AI6102 Machine Learning: Methodologies and Applications

** Prerequisite: pass AI6104 (except those with who have Math degrees / taken probability and linear algebra courses (full semester courses)*

This course aims to provide an introductory but broad perspective of machine learning fundamental methodologies, and show how to apply machine learning techniques to real-world applications. It is relevant for anyone pursuing a career in AI or Data Science.

The teaching content of this course includes different machine learning methodologies in various machine learning paradigms, such as supervised learning, semi-supervised learning, unsupervised learning, etc., as well as applications of machine learning.

AI6103 Deep Learning and Applications

Deep learning has recently introduced a paradigm shift from human-design features to end-to-end systems, and has revolutionized several fields including computer vision, speech recognition, and natural language processing. Top IT companies like Google, Facebook, Microsoft, Apple, Amazon have been actively redesigned their products with deep learning techniques, and the impacts in the coming decades will go beyond self-driving cars, strategic games like Go, and MRI cancer detection.

The main objective of this course is to introduce the mathematical foundations, the

state-of-the-art architectures, and a professional library of deep learning architectures. Students will learn how to design their own artificial neural network to solve their data analysis task. They will also learn how to code efficiently these new algorithms using PyTorch, one of the most powerful libraries in this field.

AI6104 Mathematics for Artificial Intelligence

This course aims to provide the appropriate mathematical background to students who will study other courses in the Master of Science in Artificial Intelligence programme. Upon completion of this subject, the student should be able to:

- Solve systems of linear equations
- Use linear independence basis to decompose vectors.
- Compute partial derivatives
- Perform approximation based on Taylor series
- Search maximum and minimum in multivariable functions.

AI6120 Python Programming

This course aims to provide appropriate computing background to students who will study other courses in Data Science and AI. Upon completion of this subject, the student should be able to:

- Use the Jupyter Notebook environment for Python.
- Solve standard algorithmic problems in Python.
- Perform standard numeric computations in Python.
- Handle data in Pandas, and graphics in Matplotlib.

AI6121 Computer Vision

Computer vision has been attracting increasing interests thanks to the recent advances in deep learning especially in convolutional neural networks, recurrent neural networks, etc. Top IT companies around the world such as Google, Apple, Alibaba, Amazon, Tencent and Baidu have invested and will continue to invest heavily in various computer vision technologies due to demands in a wide spectrum of applications and domains such as robotics, autonomous driving, surveillance and security, computer-aided medical diagnosis, etc.

The main objective of this course is to introduce the relevant mathematical foundations for computer vision and machine learning, the basic image analytics and machine learning technology in computer vision, and the computer vision technologies in detection, recognition and classification tasks. A series of real-world problems and challenges will be presented throughout the course via case studies and projects. Students will have a good understanding of various machine learning and computer vision technologies at the end of

this course. They will also learn how to design their own machine learning and computer vision systems for various real-world problems.

AI6122: Text Data Management and Processing

This course covers fundamental techniques to manage and process text data. The main topics include: 1) text indexing and search: inverted index, query processing, ranking, and evaluation, 2) word-level, sentence-level, document-level, and collection-level processing: morphological analysis, part-of-speech tagging, parsing, summarization, classification and clustering, and topic modeling, and 3) case studies and applications: social media text, sentiment analysis, and information extraction.

AI6123: Time Series Analysis

Many of the complex systems are dynamic systems in which their states change over time. This course introduces time series models and the corresponding methods for data analysis and inference. Topics include regression models, autoregressive (AR), moving average (MA), ARMA, and ARIMA processes, stationary and non- stationary processes, seasonal processes, identification of models, estimation of parameters, diagnostic checking of fitted models, rare event detection, forecasting, spectral analysis and time series models of heteroscedasticity . Real world applications for understanding characteristics, modelling and evaluating forecasts of time series data in economics, finance and industries are elaborated with lab on using R.

AI6124 Neuro Evolution and Fuzzy Intelligence

We begin with a short introduction of learning, adaptation andcognition: examining: learning forms, supervised, unsupervised, reinforcement; Adaptation mechanism - error driven, performance objective; cognitive knowledge representation and reasoning for rational decision making. The tools for performing such learning via evolutionary and neural computations are examined. Cognitive organisation of knowledge and the integration of such structures with soft computing models are also discussed. Introduction to knowledge abstraction and expression are discussed in relation to the evolving deep neural structures that can be used to evolve and to directly learn such cognitive knowledge structure from data. Symbolic and semantic (Fuzzy representation) knowledge are considered as sound deductive inference systems. Their logical representation and philosophy within the cognitive knowledge organisation are considered. Fuzzy logic and fuzzy inferences are covered. Their implementations within a deep learning neural structure is discussed. Different form of learning neural network that can be dynamically evolve or construct the fuzzy rules on the basis of data is covered. We will also cover model cross validation and verification for the assessment of the quality of the constructed/evolved fuzzy neural networks. Issues of concept formation, rule association, data drifts and shifts, forgetting, positive and negative

potentiation, evolutionary trade-off between expression and abstraction will be covered.

Next, we examine related cognitive intelligence applications such as those: in computational finance - risk based trading system, asset assessment, fuzzy trading models; Computational Attention focus - emotion, sentiment and motivation modelling.

AI6125 Multi-Agent System

This course covers an introduction to multi-agent systems. It will begin by introducing multi-agent systems and a number of interesting application areas, as well as explain the complexity of a definition. Decision making and planning in these multi-agent systems are then addressed. Multi-agent systems are also explained through the concept of conflicting interactions, game theory and the equilibrium, categorized as the systems of benevolent agents and those of self-interested agents. In each topic, students are given several state-of-the-art papers in the literature to read and present. The course will have a large project component. Students are asked to implement their own multi-agent system and present their project outcome in the end.

AI6126 Advanced Computer Vision

The rise of deep learning leads to a wave of breakthroughs in computer vision, opening up a number of possibilities that are unimaginable before. A lot of new applications emerge, including face recognition, image super-resolution, and machine vision in self-driving cars.

This course aims to develop specific knowledge in advanced computer vision by using deep learning techniques. The course will cover both image and video recognition, including image classification, various object detection techniques, image and video segmentation, image super-resolution, and video captioning. Focus will be on both conceptual understanding of techniques, as well as on practical implementation of these techniques using modern deep learning frameworks. Students will also gain experience in critically reviewing scientific literature relevant to the area. Course project will encourage students to solve real-world research problems, and also give them experience in written and oral scientific communication. Overall, this course will provide good exposure to students who intend to conduct research & development in computer vision.

AI6127: Deep Neural Networks for Natural Language Processing

Natural language processing (NLP) is one of the most important fields in artificial intelligence (AI). It has become very crucial in the information age because most of the information is in

the form of unstructured text. NLP technologies are applied everywhere as people communicate mostly in language; language translation, web search, customer support, emails, forums, advertisement, radiology reports, to name a few. There are a number of core NLP tasks and machine learning models behind NLP applications.

Deep learning has recently brought a paradigm shift from traditional task-specific feature engineering to end- to-end systems, and has obtained high performance across many different NLP tasks and downstream applications. Tech companies like Google, Alibaba, Apple, Amazon, Facebook, Tencent, and Microsoft are now actively working on deep learning methods to improve their products. For example, Google replaced their traditional statistical machine translation and speech recognition systems with systems based on deep learning. The course provides a thorough introduction to cutting-edge methods in deep learning applied to NLP. On the model side, we will cover word representations, window-based neural networks for NLP, recurrent neural networks for NLP, long-short-term-memory models, recursive neural networks for parsing, convolutional neural networks, as well as more advanced topics like encoder-decoder models and attention mechanisms.

AI6128 Urban Computing

Urban computing is an emerging subject that aims to deal with the challenges in the rapid urbanization progress such as pollution, energy consumption, and traffic congestion, etc. Based on the Internet of Things (IoT) infrastructure and artificial intelligence, urban computing acquires, integrates, and analyzes big and heterogeneous data generated by a diversity of sources in urban spaces, such as sensors, devices, vehicles, buildings, and human, to tackle the major urban challenges.

This course will introduce 1) the concept and applications of urban computing, 2) urban sensing technologies including the acquisition of spatio-temporal urban data, 3) cloud computing infrastructure to support urban computing applications, 4) urban data management including spatial data management and spatio-temporal data management, 5) ubran data mining including spatial data mining and spatio-temporal data mining, 6) urban data learning including spatio-temporal prediction and reinforcement learning and 7) advanced topics such as edge computing and cross-domain urban data fusion.

AI6129 AI Master Project

Course Coordinator: Asst Prof Yu Han han.yu@ntu.edu.sg

This project, to be performed over two consecutive semesters, provides students an opportunity to work with faculty members in SCSE and to learn state-of-the-art AI techniques for a particular problem. Scope of project may be based on staff research interest or industrial collaboration. Students are expected to document their work and report their findings in formal reports, and give oral presentations together with demonstration (if any) as the conclusion of their projects. The master project of AI is a one year project with 6 AUs.

The master project of AI is a one-year project with 6 AUs.

Interested students are to register for AI6129 during the Course Registration Period first, project is chosen in consultation with a supervisor (via the on-line MSAI project application available via GSLink) during the first two weeks of the semester in which the course will be taken. The project topic requires agreement by the proposed supervisor.

Project Timeline
- Faculty Members propose Projects
- Students choose AI Master Projects
- Students work on the Projects
- Students submit the AI Master Project Report
- Supervisor and Examiner Evaluation
- Finalization of Marks and Marks Submission