



# Introduction to AI: Thinking Analytically about Creative Machines

Prof. Gaya Nadarajan, (PhD, University of Edinburgh)

## SHORT COURSE DESCRIPTION

---

This course aims to teach the fundamentals of artificial intelligence starting with the concepts of intelligence, **rationality** and intelligent agents. Next, it will probe into problem solving, introducing the notion of **search** by drawing examples from puzzles and games amongst others. Then, the basics of knowledge **representation** and reasoning, such as logic and planning will be explored. Machine **learning**, a fast growing subfield of AI will also be covered focusing on technologies and real-world applications such as games, biomedical applications, social networks and smart technologies. Further topics (time-permitting) include the **impact** of major AI areas such as robotics and computer vision, natural language and speech processing in our society today. This is an introductory course and would be suitable for anyone interested to delve deeper into AI in the near future. Students will be given assignments that do not require any programming.

## READING MATERIALS

---

Teaching material (slides and/or notes) will be provided during the course. Here are some recommended reading should you be interested:

1. **Textbook:** Stuart Russell and Peter Norvig. 2009. *Artificial Intelligence: A Modern Approach (3rd ed.)*. Prentice Hall Press, Upper Saddle River, NJ, USA. (<https://bit.ly/2wpXrNS>)
2. Andrew Ng. (On-going). *Machine Learning Yearning*. (<http://www.mlyearning.org/>)
3. Ian Goodfellow, Yoshua Bengio, Aaron Courville. 2016. *Deep Learning*. MIT Press (<http://www.deeplearningbook.org/>)

## COURSE REQUIREMENTS AND GRADING

---

Students are not required to have any background in AI, computer science or maths. No prior knowledge of programming is required. The core search algorithms will be taught in conceptual notation to promote analytical thinking and reasoning. Other algorithms will be covered broadly, rather than requiring deep mathematical or statistical knowledge.

### Learning outcomes:

- Articulate what artificial intelligence is and the progress that has been made to date
- Able to describe a rational agent and its environment (knowledge representation)
- Grasp basic search strategies, which are the crux of all computing algorithms
- Understand the technologies behind machine learning, an AI branch that has succeeded to pervade our daily lives
- Demonstrate basic knowledge of AI applications such as robotics and computer vision, natural language and speech processing and their impact on various domains from classical games to more critical ones such as medicine and security
- Have awareness and thoughts on ethical and philosophical issues related to the rapid development of AI today

### Grading:

All ISS classes are pass/fail based on the student academic achievement evaluated by grades on a scale of 100 points (grade of 60 or above is Pass). **SKKU regulations require students to attend at least 80% of all classes.**

### Breakdown:

30% Class attendance and participation

30% Quizzes & Presentations

40% Assignments

*Note:* In assignments and presentations, students should always acknowledge the source where a particular piece of information is obtained (e.g. books, papers, credible websites), if it is not their own. Failure to do so could lead to suspicion of plagiarism.

Grounds for failing the class include failure to reach at least 60% in grades, failure to show up for 80% of classes and/or academic dishonesty.

## **COURSE SCHEDULE**

---

Subject to minor changes, e.g. guest speaker, presentations, quizzes and discussions

<u>Thursday (27 June)</u>	Intro	<b>– WEEK I –</b>
<u>Friday (28 June)</u>		<i>What is AI?:</i> history, rationality, strong & weak AI
		<i>Intelligent agents:</i> perception, actions, goals & environment
		<b>– WEEK II –</b>
<u>Monday (1 July)</u>	Core Concepts	<i>Problem solving and search:</i> initial states, operators, goal test, path cost function. Classical search problems
<u>Tuesday (2 July)</u>		<i>Uninformed search algorithms:</i> depth-first, breadth-first, iterative deepening
<u>Wednesday (3 July)</u>		<i>Informed search algorithms:</i> best-first, greedy, A*. Heuristics
<u>Thursday (4 July)</u>		<i>Propositional and first order logic:</i> syntax & semantics, connectives, premises and conclusions, inference rules, quantifiers, sentences, axioms, theorems, indexing, retrieval
		<b>– WEEK III –</b>
<u>Monday (8 July)</u>	Technologies	<i>Planning and reasoning:</i> initial state, goal, operators (action, precondition, effect), scheduling, reinforcement learning
<u>Tuesday (9 July)</u>		<i>Machine learning I:</i> supervised & unsupervised learning, linear regression, logistic regression (classification), clustering, overfitting, validation and testing
<u>Wednesday (10 July)</u>		<i>Machine learning II:</i> decision trees, k-nearest neighbours, naïve Bayes, random forests, SVMs, dimensionality reduction, boosting, bagging, bootstrapping

Thursday (11 July)

*Neural networks and deep learning:* NN basics, cost function, activation functions, vanishing gradients, backpropagation, gradient descent, CNNs, RNNs, GANs

– WEEK IV –

Monday (15 July)

*Application 1 – Robotics & computer vision:* robot-assisted surgeries, Robocup, ASIMO, Sofia, robots assisting the disabled and elderly, autonomous vehicles, medical imaging

Tuesday (16 July)

*Application 2 – Text and speech processing:* automatic text/voice translators, handwriting recognition, sentiment analysis, virtual digital assistants, recommendation systems (music, news), chatbots

Wednesday (17 July)

*Application 3 – Games and others:* Chess (Deep Blue), Go (AlphaGo, AlphaZero), drug discovery, personalised advertisements and content, image recognition

Thursday (18 July)

Philosophical & ethical issues, Review

Friday (19 July)

Presentations & conclusions

AI Impact