Title (Units): **COMP7015 Artificial Intelligence (3,3,0)**

Course Aims: To describe the fundamentals concepts, learning models, and techniques in artificial intelligence (AI). To give students practical insights into the current development of the field.

Prerequisite: Nil

Course Intended Learning Outcomes (CILOs):
Upon successful completion of this course, students should be able to:

No.	Course Intended Learning Outcomes (CILOs)
	Knowledge
1	Explain the capabilities, strengths and limitations of various AI techniques
2	Explain various AI algorithms and their applications
3	Describe learning models and algorithms
	Professional Skill
4	Apply selected AI algorithms to solve real world problems
5	Understand complex ideas and relate them to specific real-world applications

Calendar Description: This course aims to introduce the AI principles and the associated techniques. Students will learn the AI ideas and algorithms, and acquire practical insights into the current development of this field.

Teaching and Learning Activities (TLAs):

CILOs	Type of TLA
1-3,5	Students will learn the AI principles and techniques in lectures.
3-5	Students will work on assignments to enhance the understanding of AI ideas and acquire hands-on experience through a course project.

Assessment:

No.	Assessment Methods	Weighting	CILOs to be addressed	Description of Assessment Tasks
1	Continuous Assessment	50%	1-5	Written and programming assignments, quiz/test and hands-on course project will be used to evaluate students' understanding of basic concepts and to assess their ability to apply AI ideas and learning algorithms to solve real world problems
2	Examination	50%	1-5	Examination will be used to assess students' overall understanding of various AI algorithms, their applications, as well as their capabilities, strengths and limitations.

Assessment Rubrics:

Excellent (A)	• Achieve the first three CILOs, demonstrating an excellent mastery of both the theoretical and practical aspects of the knowledge and skills associated with AI • Able to develop correct solutions to problems • Demonstrate a thorough understanding and solid knowledge of AI • Able to apply a variety of techniques and relevant knowledge for solving problems in AI
Good (B)	• Achieve most of the first three CILOs, demonstrating a good understanding of the knowledge and skills associated with AI • Able to develop correct solutions to problems • Demonstrate a competent level of knowledge of AI

	• Ability to make use of appropriate techniques and knowledge and apply them to familiar problems
Satisfactory (C)	• Achieve some of the first three CILOs, demonstrating a basic level of understanding of the knowledge and skills associated with AI • Able to provide acceptable solutions to problems • Demonstrate an adequate level of knowledge of AI • Ability to make use of some techniques and knowledge and apply them to familiar situations
Fail (F)	• Achieve none of the first three CILOs, with little understanding of the associated concepts and underlying methodologies • Unable to provide solutions to simple problems • Knowledge of AI falling below the basic minimum level • Unable to apply techniques and knowledge to situations or problems

Course Content and CILOs Mapping:

Content		**CILO No.**
I	Introduction to AI	1,3,5
II	Search	2,3
III	Knowledge Representations and Reasoning	3,4
IV	Basics of Statistical Machine Learning	1,2,3,4,5
V	Artificial Neural Networks and Deep Learning	1,3,5
VI	Bayesian Learning	2,3,4,5
VII	Evolutionary Computation	1,2,3
VIII	Generative AI Methods and Applications	1,2,4
IX	Reinforcement Learning	2,3,5

References:

- Stuart Russell and Peter Norvig, Artificial Intelligence, A Modern Approach, Prentice Hall, 4th Edition, 2020.
- Michael Negnevitsky, Artificial Intelligence: A Guide to Intelligent Systems, Addison Wesley, 3rd Edition, 2011.
- Vladimir N. Vapnik, Statistical Learning Theory, Wiley, 1998.
- Ian Goodfellow, Yoshua Bengio and Aaron Courville, Deep Learning, MIT, 2016.
- Tom M. Mitchell, Machine Learning, McGraw-Hill 1997.
- Christopher M. Bishop, Pattern Recognition and Machine Learning, Springer, 2006.
- Kevin P. Murphy, Machine Learning: A Probabilistic Perspective, MIT, 2012.
- Trevor Hastie, Robert Tibshirani and Jerome Friedman, The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Springer, 2009.
- Aston Zhang, Zachary C. Lipton and Mu Li, Dive into Deep Learning, Cambridge University Press, 2024.

Course Content:

Topic

I. Introduction to AI
   A. History
   B. Applications
   C. Prospect

II. Search
   A. Uninformed search
   B. Heuristic search
   C. Constraint satisfaction search

III. Knowledge Representations and Reasoning
   A. Prepositional and predicate logic

B. Other representation techniques
C. Uncertainty knowledge and reasoning

IV. Basics of Statistical Machine Learning
A. Decision tree learning algorithms
B. Logistic regression and loss minimization framework
C. Overfitting and regularization techniques
D. Evaluating hypotheses

V. Artificial Neural Networks and Deep Learning
A. Deep feedforward networks
B. Regularization for deep learning
C. Convolutional networks
D. Sequence modeling
E. Using pretrained models

VI. Bayesian Learning
A. Maximum likelihood and least-squared error hypotheses
B. Minimum description length principle
C. Bayes optimal classifier and Gibbs algorithm
D. Bayesian belief networks

VII. Evolutionary Computation
A. Genetic algorithms
B. Hypothesis space search
C. Genetic programming
D. Models of evolution and learning

VIII. Generative AI Methods and Applications
A. Diffusion Models and Their Applications
B. Generative Models with Large Language Models

IX. Reinforcement Learning
A. Q-learning
B. Temporal difference learning
C. Relationship to dynamic programming