Title (Units): **COMP7250 Machine Learning (3,2,1)**

Course Aims: To introduce the basic concepts, theories and techniques of machine learning from an application perspective. To give students practical insights into the current development of the field.

Prerequisite: Basic knowledge of probability theory and optimization techniques and COMP7015 Artificial Intelligence

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 machine learning techniques
2	Explain various machine learning algorithms and their applications
3	Describe learning models and algorithms
	Professional Skill
4	Apply machine learning algorithms to solve real-world problems, understand advanced machine learning algorithms, cultivate the ability to evaluate learning algorithms
5	Implement solutions to real-world problems using deep learning toolboxes such as PyTorch or Tensorflow

Calendar Description: This course aims to introduce the principles and techniques of machine learning. Students will learn the machine learning techniques and acquire practical insights into the current development of this field.

Teaching and Learning Activities (TLAs):

CILOs	Type of TLA
1-3	Students will learn the basic concepts and fundamental principles of machine learning, as well as the application examples, in lectures.
4-5	Students will work on assignments to enhance the understanding of learning principles, and acquire hands-on experience via laboratory classes and a mini project.

Assessment:

No.	Assessment Methods	Weighting	CILOs to be addressed	Description of Assessment Tasks
1	Continuous Assessment	40%	1-5	Assignments and a hands-on mini-project will be used to evaluate students' understanding of basic concepts and to assess their ability to apply learning theory to solve real-world problems
2	Examination	60%	1-5	Examination will be used to assess students' overall understanding of various machine learning algorithms, their applications, as well as their capabilities, strengths and limitations.

Assessment Rubrics:

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

Good (B)	• Achieve most of the first three CILOs, demonstrating a good understanding of the knowledge and skills associated with machine learning • Able to develop correct solutions to problems • Demonstrate a competent level of knowledge of machine learning • Able 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 machine learning • Able to provide acceptable solutions to problems • Demonstrate an adequate level of knowledge of machine learning • Able 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 machine learning 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	The Evolution of Machine Learning Algorithms and Applications	1,2,3,5
II	Deep Learning from Application Perspective	3,5
III	Foundation Models from Application Perspective	2,3,4,5
IV	Other Topics in Modern Machine Learning	2,3,4,5

References:

- Steven W. Knox, Machine Learning, Wiley, 2018
- Tom M. Mitchell, Machine Learning, McGraw-Hill International Editions, 1997
- Trevor Hastie, Robert Tibshirani and Jerome Friedman, The Elements of Statistical Learning: Data Mining, Inference, and Prediction, New York: Springer, 2nd Edition, 2009.
- Ian Goodfellow, Yoshua Bengio and Aaron Courville, Deep Learning, MIT Press, 2016.
- Hans Georg Schaathun, Machine Learning in Image Steganalysis, Wiley, 2012.
- Mehryar Mohri, Afshin Rostamizadeh and Ameet Talwalkar, Foundations of Machine Learning (Adaptive Computation and Machine Learning series) 2nd Edition, The MIT Press, December 2018.
- Francois Chollet, Deep Learning with Pytorch, Manning Publication, 2017.
- Aston Zhang, Zachary C. Lipton, Mu Li, and Alexander Smola, Dive into Deep Learning, 2021.

Course Content:

Topic

I. The Evolution of Machine Learning Algorithms and Applications
   A. Algorithm Level
   B. Application Level

II. Deep Learning from Application Perspective
   A. Models
   B. Learning Objectives
   C. Optimization and Regularizations
   D. Lab-1: ResNet Implementation
   E. Lab-2: LSTM Implementation

III. Foundation Models from Application Perspective
   A. Large Language Models
   B. Vision-language Models
   C. Learning Algorithms

D. Deployment Strategies
E. Lab-3: LoRA Implementation

IV. Other Topics in Modern Machine Learning
A. Trustworthy Machine Learning for Science
B. Lab-4: FGSM and PGD Implementation