Lingnan University
Course Syllabus

Course Title	:	Deep Learning
Course Code	:	CDS4006
Recommended Study Year	:	4
No. of Credits/Term	:	3
Mode of Tuition	:	Sectional Approach
Class Contact Hours	:	3 hours per week
Category in Major Prog.	:	Elective
Prerequisite(s)	:	CDS3003 Machine Learning
Co-requisite	:	Nil
Exclusion	:	Nil
Exemption Requirement	:	Nil

Brief Course Description:

Deep learning is one of the bleeding edge technologies of machine learning. It has been used in the community and industry to solve a number of big data problems in computer vision and natural language processing. Deep learning drives improvement in the precision of medical diagnosis, enhancing digital painting and the development of self-driving vehicles. This course teaches students to understand the theoretical and technical details of deep learning and to design intelligent systems that learn from complex and large-scale datasets.

Aims:

This course aims at providing an intensive understanding of the existing deep learning approaches. The topics will cover large-scale deep learning, multi-task deep learning, and transferred deep learning. Related techniques and models will be introduced: how to train and optimize basic neural networks, deep belief nets, deep autoencoders, convolutional neural networks, deep residual networks, and long short term memory networks. TensorFlow or other state-of-the-art deep learning tools will be introduced and applied to solve different classes of problems with huge datasets in Science, Social Science, Arts, and Business domains.

Learning Outcomes (LOs):

Upon the successful completion of this course, the student will be able to:
1. Identify, explore, describe, and test various deep learning models; (PLO5)
2. Apply deep learning models to solve data science problems; (PLO5)
3. Compare and evaluate various deep learning networks under different scenarios; (PLO7)
4. Develop and deploy deep learning systems to solve some sophisticated real-world problems. (PLO9)

Indicative Contents:

Deep Feedforward Networks
Basic concepts, overfitting, over generalization, stochastic gradient descent, hidden units, architecture design, back-propagation and other differentiation algorithms

Regularization for Deep Learning
Parameter norm penalties, regularization and under-constrained problems, semi-supervised learning, multitask learning, early stopping, sparse representations, dropout, adversarial training

Convolutional Neural Networks
Convolution layers, pooling layers, batch normalization, deep residual networks, deconvolution, transferred learning

Recurrent Neural Networks
Recurrent neural networks, bidirectional recurrent neural networks, long-short term memory models

Other Deep Learning Methods
Variational autoencoders, stacked denoising autoencoders, deep belief network, deep reinforcement learning, generative models

Ethical Deep learning
Ethical awareness, data privacy, ethical case studies and dillemmas in deep learning

Teaching Method:

There are different class activities including lectures, laboratories, group project presentations, and tutorials. The concepts of deep forward networks, regularization, convolutional neural networks, and recurrent neural networks will be discussed in lectures. The implementations of deep neural networks in TensorFlow, Theano, Caffe2, PyTorch, or IBM PowerAI Enterprise will be discussed in laboratories. Students are required to perform a group project to apply the concepts and principles covered in this course to induce deep learning models for handling challenging data science problems with big datasets. They are required to present their models and the outcomes obtained in carrying out the project. The detailed elaborations and computations of some learning algorithms will be discussed in tutorials.

Assessment:

Class Attendance and Participation	5%
Assignments	40%
Group Project & Presentation	20%
Examination	35%
Total	100%

Measurement of Learning Outcomes:

	Class Attendance and Participation	Assignments	Group Project & Presentation	Examination
1. Identify, explore, describe, and test various deep learning models	x			x
2. Apply deep learning models to solve data science problems		x	x	
3. Compare and evaluate various deep learning networks under different scenarios		x	x	
4. Develop and deploy deep learning systems to solve some sophisticated real-world problems			x	
5. Ethical Awareness	x			

1. There are a number of classroom activities to evaluate students' comprehensive understanding on the theory of deep learning methods. Students are expected to identify, explore, describe, and test various deep learning models (LO1 and PLO5).
2. Assignments require students to implement existing deep neural network architecture(s) and apply them to solve data science problems. Students are expected to compare and evaluate their deep neural network models (LO2, LO3, PLO5, and PLO7).
3. A Group Project requires students to apply the deep learning technologies to solve a (big) data mining problem. They need to consider and apply different deep learning models. Experiments should be performed to compare and evaluate their deep learning models. Students need to develop and deploy their deep learning models in some systems (LO2-4, PLO5, PLO7, PLO9).
4. The Group Project presentation measures student's presentation skills by requiring each student group to present their projects orally as well as answering questions from audiences. The Group Project and presentation are assessed for logicality, strength of arguments and feasibility of recommendations (LO2-4, PLO5, PLO7, PLO9).
5. The understanding of deep learning is evaluated in the examination. Students are expected to identify, explore, describe, and test various deep learning models. Programming and testing on some software packages will be performed during the examination (LO1, PLO5).

Required/Essential Readings:

1. Goodfellow, Ian, Bengio, Yoshua, and Courville, Aaron, *Deep learning*, MIT press, 2016.

Recommended/Supplementary Readings:

1. Buduma, N., *Fundamentals of Deep Learning: Designing Next-Generation Machine Intelligence Algorithms*, O'Reilly Media, 2017.

2. Geron, Aurelien, *Hands-On Machine Learning with Scikit-Learn and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems*, O'Reilly Media, 2017.
3. Patterson, Josh and Gibson, Adam, *Deep Learning: A Practitioner's Approach*, O'Reilly Media, 2017.
4. Poole, David L. and Mackworth, Alan K. *Artificial Intelligence: Foundations of Computational Agents, 2nd Edition*, Cambridge University Press, 2017.
5. Raschka, Sebasian and Mirjalili, Vahid, *Python Machine Learning: Machine Learning and Deep Learning with Python, scikit-learn, and TensorFlow, 2nd Edition*. Packt Publishing, 2017.

Recommended/Supplementary Paper Readings:

Vaswani A, Shazeer N, Parmar N, et al. Attention is all you need[J]. Advances in neural information processing systems, 2017, 30.

He K, Zhang X, Ren S, et al. Deep residual learning for image recognition[C]//Proceedings of the IEEE conference on computer vision and pattern recognition. 2016: 770-778.

Li Q, Wen Z, Wu Z, et al. A survey on federated learning systems: Vision, hype and reality for data privacy and protection[J]. IEEE Transactions on Knowledge and Data Engineering, 2021, 35(4): 3347-3366.

Ramesh A, Dhariwal P, Nichol A, et al. Hierarchical text-conditional image generation with clip latents[J]. arXiv preprint arXiv:2204.06125, 2022, 1(2): 3.

Important Notes:

(1) Students are expected to spend a total of 9 hours (i.e. 3 hours of class contact and 6 hours of personal study) per week to achieve the course learning outcomes.
(2) Students shall be aware of the University regulations about dishonest practice in course work, tests and examinations, and the possible consequences as stipulated in the Regulations Governing University Examinations and Coursework. In particular, plagiarism, being a kind of dishonest practice, is "the presentation of another person's work without proper acknowledgement of the source, including exact phrases, or summarised ideas, or even footnotes/citations, whether protected by copyright or not, as the student's own work". Students are required to strictly follow university regulations governing academic integrity and honesty.
(3) Students are required to submit writing assignment(s) using Turnitin.
(4) To enhance students' understanding of plagiarism, a mini-course "Online Tutorial on Plagiarism Awareness" is available on https://pla.ln.edu.hk/.
(5) Students must adhere to the University's guidelines and practices when using Generative Artificial Intelligence (GAI) tools. The official documents *"Guidelines for Using GAI Tools at Lingnan University" and "Best Practices for Ethical and Responsible Use of GAI Tools in Course Assessments" can be found here: https://www.ln.edu.hk/tlc/generative-artificial-intelligence/gai-guidelines-best-practices* .

Rubric for Final Examination of CDS4006 – Deep Learning

Criteria	Very good (4-6)	Satisfactory (2-4)	Unsatisfactory (0-2)
Identify different deep learning models for various applications (LO1)	The student can correctly identify most deep learning models for different problems. Explanations have been given and most of them are correct.	The student can correctly identify many deep learning models for different problems. Explanations have been given but some of them are incorrect.	The student can correctly identify very few deep learning models for different problems. Explanations have not been provided or the explanations are basically incorrect.
Describe different deep learning models (LO1)	Nearly all deep learning models can be described correctly.	The student can describe many deep learning models correctly.	The student can only describe very few deep learning models correctly.
Explore different deep learning models (LO1)	The student can explore and implement different deep learning models for some problems. Most of the implementations are correct.	The student can explore and implement different deep learning models for some problems. But many implementations have issues.	The student cannot explore and implement different deep learning models for some problems, or the implementations are basically incorrect.
Test different deep learning models (LO1)	**All** of the following points are achieved ▪ Employ the right testing procedure(s) ▪ Employ the right testing metrics ▪ Calculate all testing values correctly.	**Two** of the following points are achieved ▪ Employ the right testing procedure(s) ▪ Employ the right testing metrics ▪ Calculate all testing values correctly.	**None or One** of the following tasks are achieved ▪ Employ the right testing procedure(s) ▪ Employ the right testing metrics ▪ Calculate all testing values correctly.

Rubric for Individual Assignments of CDS4006 – Deep Learning

Criteria	Very good (4-6)	Satisfactory (2-4)	Unsatisfactory (0-2)
Able to load the datasets, encode and feed input features, and handle missing and noisy data. (LO2)	**All** of the following points are achieved ▪ Load large datasets ▪ Encode all input features in the correct form ▪ Feed all input features to the deep neural network appropriately ▪ Handle all missing and noisy data correctly	**Two** of the following points are achieved ▪ Load large datasets ▪ Encode all input features in the correct form ▪ Feed all input features to the deep neural network appropriately ▪ Handle all missing and noisy data correctly	**None or One** of the following points are achieved ▪ Load large datasets ▪ Encode all input features in the correct form ▪ Feed all input features to the deep neural network appropriately ▪ Handle all missing and noisy data correctly
Able to construct various deep learning networks. (LO2)	Can construct various deep learning networks according to the specifications.	Can construct various deep learning networks but some modules and/or connections are misplaced according to the specifications.	Cannot construct any deep learning network.
Can employ different technologies to perform data augmentation and dropout to handle overfitting. (LO2)	Appropriate data augmentation and dropout methods have been performed. Suitable new data sets have been generated. The settings of the methods are completely correct.	Appropriate data augmentation and dropout methods have been performed, although the transformed data sets have some minor issues, because the settings of the methods are not completely correct.	Data augmentation and dropout methods have not been performed or most procedures have been performed incorrectly.
Able to learn models from data. (LO2)	Appropriate deep learning methods have been used. The procedure of using these methods and the settings of these methods are completely correct. The learnt models are accurate.	Appropriate deep learning methods have been used. But the procedure of using these methods and the settings of these methods are not completely correct. Thus the learnt models have some issues.	Inappropriate deep learning methods have been used. The learnt models are incorrect.
Able to evaluate the performance of the learnt models. (LO3)	Appropriate methods (e.g. cross-validation, nested cross-validation) have been applied to divide the dataset into training, testing and/or validation data sets. Validates the models using the training and/or the validation datasets. Applies suitable methods (e.g. loss functions) to evaluate the performance of the model on the testing dataset.	Performance evaluation has been done. The evaluation strategy is generally correct, but the evaluation results are biased or unstable.	Performance evaluation has not been done or the evaluation strategy is incorrect.
Able to deliver the learnt models and employ existing pretrained neural networks. (LO2)	Applies the learnt models and combines pretrained neural networks for decision making on new examples. The decisions are sound and reasonable.	Applies the learnt models and combines pretrained neural networks for decision making on new examples. The decisions are not completely sound and reasonable.	Cannot apply the learnt models and combines pretrained neural networks for decision make on new examples.

Rubric for Group Project Presentation of CDS4006 – Deep Learning

Criteria	Very good (4-6)	Satisfactory (2-4)	Unsatisfactory (0-2)
Appropriate time allocation and pace.	Allocate time appropriately, and manage time effectively, with smooth progression. Appropriate pace. Start presentation punctually.	Marginally long or marginally short but uses time reasonably effectively. Reasonable pace. Start presentation relatively punctually.	Significantly too short or too long and did not use time effectively. Pace is significantly too fast or too slow. Don't start presentation punctually.
Clear, logically organized and relevant content.	Information included is always relevant. Clearly stated and developed points. Material flows extremely well and is well organized. No ambiguities are left unexplained.	Information included is generally relevant. Key points are relatively clear. Most information presented in logical sequence; sufficiently well-organized with generally satisfactory flow. Some ambiguities are left unexplained.	Much of the information included is not relevant and even key points are not clear. Presentation is choppy or disjointed, does not flow well, and has no apparent logical order.
Effective use of presentation tools.	Balanced and proper use of presentation tools with little or no distraction (e.g., appropriate animation/pictures, appropriate information on one slide, good color combination, clear titles, etc.)	Generally good use of presentation tools. Some distractions but they are not overwhelming (e.g., reasonable animation/pictures, fair information on one slide, fair color combination, fair titles, etc.)	Poor use of presentation tools and/or many distractions (e.g., too much animation/pictures, too much information on one slide, poor color combination, absence of titles, etc.)
Uses good body language, eye contact, appropriate voice tone.	Shows poise and composure; makes good eye contact with audience; balanced posture; shows enthusiasm and confidence; uses voice tone effectively.	Fairly poised and composed; makes fairly good eye contact with audience; balanced posture; shows some enthusiasm and confidence; uses voice tone relatively effectively.	Little poise and composure; makes little or no eye contact with audience; poor posture; shows little or no enthusiasm and confidence; uses voice tone ineffectively or too monotonously.
Gains/holds attention	Provides good motivation to engage the audience's interest. Presents the content in a manner that captivates the audience's attention.	Provides reasonable motivation to engage the audience's interest. Audience is reasonably engaged but there are instances where the presentation is otherwise dull.	Provides insufficient motivation to engage the audience's interest. Dull presentation of content that does not engage the audience.
Uses instructor defined role appropriate dress	Professionally dressed as expected by the instructor.	Minor deviations from instructor's expectations.	Do not dress in a manner expected by the instructor.
Clarity of speech/Accuracy of grammar & pronunciation	Voice is consistently comprehensible; grammar and pronunciation are accurate.	Voice is generally comprehensible; grammar and pronunciation are adequate but with some mistakes.	Voice is incomprehensible on several occasions; many mistakes in terms of grammar and pronunciation.

Rubric for Group Project of CDS4006 – Deep Learning

Criteria	Very good (4-6)	Satisfactory (2-4)	Unsatisfactory (0-2)
Problem definition (demonstrate the understanding of the problem and consider alternative solutions)	Clearly state the problem, list related constrains, and are able to identify alternative solutions.	The problem is stated but related constrains and alternative solutions are not considered thoroughly.	The problem is marginally defined and with no consideration of constrains and alternative solutions.
Solution design	The design of the solution is strongly related to the problem and the approaches and techniques involved are clearly explained. Moreover, the design of the solution is very innovative.	The design of the solution is related to the problem but the approach and techniques involved are not clearly explained. However, the design of the solution is not very innovative.	The design of the solution is weakly related to the problem and no explanation of the approach and techniques involved. The design of the solution is not innovative.
Able to load the datasets, encode and feed input features, and handle missing and noisy data (LO2)	**All** of the following points are achieved ▪ Load large datasets ▪ Encode all input features in the correct form ▪ Feed all input features to the deep neural network appropriately ▪ Handle all missing and noisy data correctly	**Two** of the following points are achieved ▪ Load large datasets ▪ Encode all input features in the correct form ▪ Feed all input features to the deep neural network appropriately ▪ Handle all missing and noisy data correctly	**None or One** of the following points are achieved ▪ Load large datasets ▪ Encode all input features in the correct form ▪ Feed all input features to the deep neural network appropriately ▪ Handle all missing and noisy data correctly
Able to construct various deep learning networks (LO2)	Can construct various deep learning networks according to the design of the solution.	Can construct some deep learning networks but some modules and/or connections are misplaced according to the design of the solution.	Can construct only one or just a few deep learning network.
Can employ different technologies to perform data augmentation and dropout to handle overfitting (LO2)	Appropriate data augmentation and dropout methods have been performed. Suitable new data sets have been generated. The settings of the methods are completely correct.	Appropriate data augmentation and dropout methods have been performed, although the transformed data sets have some minor issues, because the settings of the methods are not completely correct.	Data augmentation and dropout methods have not been performed or most procedures have been performed incorrectly.
Able to learn models from data (LO2)	Appropriate deep learning methods have been used. The procedure of using these methods and the settings of these methods are completely correct. The learnt models are accurate.	Appropriate deep learning methods have been used. However the procedure of using these methods and the settings of these methods are not completely correct. Thus the learnt models have some issues.	Inappropriate deep learning methods have been used. The learnt models are incorrect.
Able to evaluate the performance of the learnt models (LO3)	Appropriate methods (e.g. cross-validation, nested cross-validation) have been applied to divide the dataset into training, testing and/or validation data sets. Validate the models using the	Performance evaluation has been done. The evaluation strategy is generally correct. But the evaluation results are biased or unstable.	Performance evaluation has not been done or the evaluation strategy is incorrect.

Criteria	Very good (4-6)	Satisfactory (2-4)	Unsatisfactory (0-2)
	training and/or the validation datasets. Suitable methods (e.g. loss functions) have been applied to evaluate the performance of the model on the testing dataset.		
Able to develop and deploy the learnt models in some intelligent systems (LO4)	The learnt models are implemented and deployed correctly. The intelligent system performs decision making on new examples. The decisions are sound and reasonable.	The learnt models are implemented and deployed correctly. The intelligent system performs decision making on new examples. The decisions are not completely sound and reasonable.	The learnt models cannot be implemented and deployed correctly.

Rubric for Class Attendance and Participation of CDS4006 – Deep Learning

Criteria	Very good (4-6)	Satisfactory (2-4)	Unsatisfactory (0-2)
Attendance	Full, punctual attendance in class and mandatory seminars	Occasional absences or lateness from class or mandatory seminars	Frequent or recurring absence or lateness from class or mandatory seminars
Class Participation	Active class participation and leadership in group activities.	Passive class participation and active in group activities.	Lack of participation and active disruption of class and group activities.
Identify different deep learning models for various applications (LO1)	The student can correctly identify most deep learning models for different problems. Explanations have been given and most of them are correct.	The student can correctly identify many deep learning models for different problems. Explanations have been given but some of them are incorrect.	The student can correctly identify very few deep learning models for different problems. Explanations have not been provided or the explanations are basically incorrect.
Describe different deep learning models (LO1)	Nearly all deep learning models can be described correctly.	The student can describe some deep learning models correctly.	The student can only describe very few deep learning models correctly.