

ISOM4030 Deep Learning for Business Applications

Fall 2026

Instructor: Dr. Jia JIA (justinjia@ust.hk)

Office: LSK 5045

Office Hours: By appointment

Teaching Assistant: Ray Pang (imncpang@ust.hk)

Office: LSK 4065

Class Schedule

Section	Date & Time	Venue
Lecture	Mon. 03:00PM - 04:20PM Fri. 10:30AM - 11:50AM	LSK1001
Lab	Mon. 04:30PM - 05:50PM	G021

Course Website: <https://canvas.ust.hk>. All course materials and announcements will be posted on this site. You are advised to check it regularly throughout the course.

COURSE DESCRIPTION

The past few years have marked a pivotal era in Artificial Intelligence (AI), driven by rapid advancements in deep learning technologies and transformative innovations such as Large Language Models (LLMs). These developments have catalyzed a profound shift in how AI is integrated into various sectors, enabling breakthroughs in natural language processing, computer vision, speech recognition, recommendation systems, and beyond. Deep learning applications are now ubiquitous, revolutionizing how we work, communicate, and make decisions across

This course offers a comprehensive introduction to state-of-the-art deep learning techniques and their wide-ranging practical applications. Using Python as the primary programming tool, students will acquire the skills to design, implement, and fine-tune deep learning models that address complex challenges across various business functions. The course covers applications in natural language processing, computer vision, speech recognition, and more, preparing students to leverage AI for innovation in areas such as web search, advertising, customer engagement, financial analytics, and automation.

Upon completing this course, students will be able to:

- Understand core deep learning concepts, architectures, and techniques applicable to a variety of data modalities (text, images, audio, and structured data).
- Apply deep learning methods effectively to solve real-world business problems, including natural language processing, computer vision, and predictive analytics.
- Gain practical experience with popular deep learning frameworks and tools to rapidly prototype and deploy AI models.
- Collaborate productively with cross-functional teams — such as data scientists, engineers, and business analysts — by clearly communicating AI concepts and solutions.

PREREQUISITES

- Official prerequisites: ISOM 3360
- All labs and assignments will be in Python. If you need a refresher on Python, or are not very familiar with NumPy, I highly recommend attending the first several lab sessions. These will help you catch up and ensure you have a solid foundation in these essential tools.
- College Calculus and Linear Algebra (e.g. MATH 1003 or equivalent)
You should be comfortable taking derivatives and understanding matrix/vector notation and operations. We will have labs specifically designed to reinforce these foundational concepts.
- Basic Probability and Statistics (e.g. ISOM 2500 or equivalent)
You should know the basics of probabilities, gaussian/normal distributions, mean, standard deviation, etc.

MATERIALS

1. CLASS WEBSITE

All course relevant materials, including but not limited to, reading materials, handouts, and programming files, will be uploaded to CANVAS. You are advised to check this site regularly throughout this course.

2. MAIN READING

The course will not follow a specific book, but will draw from a number of sources. We list relevant books below. We will also put up links to relevant reading material for each class.

- [Neural Networks and Deep Learning](#) by Michael A. Nielsen.
- [Dive Into Deep Learning](#) by Aston Zhang, et al.
- [Understanding Deep Learning](#) by Simon J. D. Prince

3. SOFTWARE REQUIREMENT

- Google CoLab – Most deep learning models will be built and run using Google CoLab programming environment. You must have a Google account, and have access to Google and Google Drive, in order to use Google CoLab.
- Jupyter Notebook – For Python and Machine Learning fundamentals, you may also install Anaconda and use Jupyter Notebook for faster and easier programming in Python.

EVALUATION

Components	Percentage of the grade
A. Lab submissions	10%
B. Assignment 1	15%
C. Assignment 2	15%
D. Assignment 3	15%
E. Final exam – Concepts	25%
F. Final exam – Applications	20%
TOTAL:	100%

A. Lab Submissions

Labs are an essential part of the learning process. You are NOT going to learn without practicing data processing and modeling yourselves. During lab sessions, we expect you to be able to actively link the empirical results to the concepts covered in the lectures, and learn necessary skills for completing the homework assignments. Lab submissions will be graded and count towards your final grade.

B. Assignments

3 individual assignments will be given throughout the semester. Each assignment has a specific due date and time. Late submission **within 24 hours** after the specified due date and time will be accepted with a **25% penalty**. Late submission **beyond 24 hours** will NOT be accepted.

C. Final Exam

There will be ONE open-book final exam scheduled in the final examination period. It comprises two parts labeled as *Final exam – Concepts* (25%) and *Final exam – Applications* (20%) in the evaluation section. Details of the exam will be provided later in the semester.

ACADEMIC INTEGRITY

Academic integrity is a critical value of the university community. Integrity violations destroy the fabric of a learning community and the spirit of inquiry that is vital to the effectiveness of the University. I have absolutely no tolerance for cheating and there are no acceptable excuses. Anyone caught cheating, plagiarizing, and any other form of academic dishonesty will have their course grade lowered by at least one letter grade. Any unethical behavior or evidence of dishonesty in this course will be reported to the University. Please remember the current university rule: "If a student is discovered cheating however minor the offence, the course grade will appear on the student's record with an X, to show that the grade resulted from cheating. This X grade stays

on the record until graduation. If the student cheats again and “earns” another X grade, the student will be dismissed from the University."

Plagiarism is copying anything (text or ideas) from another source without citing that source. If you use another person's idea you must cite it, even if you rewrite the idea in your own words. Extreme care must be taken to avoid passing of other's work as one's own. You are required to provide appropriate citations when you use ideas and arguments or otherwise draw on others' work. If you use research from another source or from the Web you **MUST** cite the source. This is true even if you use only the general idea and not the exact words.

OTHERS

Email Policy

Please put **[ISOM4030]** at the beginning of the subject line of your email along with your email subject. Failure to do so may result in a longer response time.

As expected, there will be numerous emails when it is closer to the due dates. If you need any assistance, raise them **as early as possible**, and/or take advantage of the office hours of the instructor and the TA. Note that **neither the instructor nor the TAs will provide direct answers to the assignments.**

Learning Environment

We welcome feedbacks on our teaching throughout the semester. You are encouraged to contact any of them at any time when you have any questions, suggestions, concerns, or would like to ask for advice. Please remember, we are here to help you learn. Therefore, please do NOT hesitate to contact us at any time, so we could do our jobs better!

TENTATIVE LECTURE SCHEDULE

(Will be adjusted as needed. Please visit Canvas site for the latest schedule, readings, and assignments)

Week	Date	Topics	Remarks
1	Sep. 4	Course overview	
2	Sep. 7	Introduction to deep learning	
	Sep. 11	Machine learning basics I: Loss function and gradient descent	
3	Sep. 14	Machine learning basics II: Multi-class classification	
	Sep. 18	Machine learning basics II: Model generalization and regularization	
4	Sep. 21	Feedforward networks I	
	Sep. 25	Feedforward networks II	
5	Sep. 28	Autoencoders & applications (Fraud detection, Denoising, Profiling, etc.)	HW1 release
	Oct. 2	Model training techniques and optimization algorithms	
6	Oct. 5	Convolutional neural networks (CNNs)	
	Oct. 9	Advanced CNN models	

7	Oct. 12	CNN applications: Semantic segmentation, object detection, & face landmarking	HW1 due
	Oct. 16	Language modeling & recurrent neural networks (RNNs)	
8	Oct. 19	<i>Public Holiday</i>	HW2 release
	Oct. 23	LSTMs, GRUs, and Bidirectional RNNs	
9	Oct. 26	Sequence-to-sequence models for machine translation	
	Oct. 30	Attention-based sequence-to-sequence models	
10	Nov. 2	Transformers	HW2 due
	Nov. 6	Applications of Transformer-based models	
11	Nov. 9	Advanced Transformers & model Serving	
	Nov. 13	Large language models (LLMs): BERT, T5, GPTs, and more	HW3 release
12	Nov. 16	Turn LLMs to intelligent assistants: Instruction fine-tuning	
	Nov. 20	Turn LLMs to intelligent assistants: Reinforcement learning	
13	Nov. 23	LLM agents	
	Nov. 27	Multi-modal models	HW3 due
14	Nov. 30	Revision	