ISOM4000B

Natural Language Processing in Business Using Deep Learning

Fall 2022

Instructor: **Dr. Jia JIA**

(justinjia@ust.hk)

Office: LSK 5045

Office Hours: By appointment

Teaching Assistant: **Ray PANG (**imncpang@ust.hk**)**

Office: LSK 4063

Class Schedule

Section	Date	Time	Venue
Lecture	Thur.	03:00PM - 05:50PM	Rm 6602, Lift 31-32
Lab	Wed.	.06:00PM - 07:20PM	G005, LSK Bldg

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

Natural language processing (NLP) is one of the most important technologies of the information age. Applications of NLP are everywhere because people communicate almost everything in language: web search, advertising, emails, customer service, language translation, virtual agents, medical reports, politics, etc. In the last decade, deep learning (or neural network) approaches have obtained very high performance across many different NLP tasks, using single end-to-end neural models that do not require traditional, task-specific feature engineering.

In this course, students will gain a thorough introduction to cutting-edge approaches in Deep Learning for NLP. Python will be the primary implementation tool. Through lectures, labs, assignments and a final project, students will learn the necessary skills to design, implement, and understand their own neural network models, using the Tensorflow framework.

After taking this course, students should:

- Know the basics of NLP techniques and deep learning models well enough to interact with business analysts, data scientists, engineers, and machine learning developers.
- Be able to apply NLP techniques and deep learning approaches to analyze text for business intelligence.
- Gain hands-on experience in using popular NLP and deep learning tools and get ready for the job positions that require familiarities with the NLP and deep learning tools.

LEARNING OBJECTIVES

1. Model business problems logically (PILO: 2, 3, 5)
2. Develop programming solutions to solve business problems (PILO: 2, 3, 5, 7)
3. Write computer programs with common programming practices (PILO: 2, 3, 5)
4. Identify and fix logical and runtime errors in computer programs (PILO: 2, 3, 5).

PREREQUISITES

- Official prequisites: ISOM 3360 or ISOM 3400

- Proficiency in Python
  All labs and assignments will be in Python. If you need to remind yourself of Python, or you're not very familiar with NumPy, you can come to the first several lab sessions.

- College Calculus and Linear Algebra (e.g. MATH 1003 or equivalent)
  You should be comfortable taking derivatives and understanding matrix/vector notation and operations.

- Basic Probability and Statistics (e.g. ISOM 2500 or equivalent)
  You should know the basics of probabilities, gaussian 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.
   - Natural Language Processing in Action by Hobson Lane et al. 2019
   - A Primer on Neural Network Models for Natural Language Processing by Yoav Goldberg

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 Machining 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	15%
B. Homework Assignments (× 3)	30% (10% × 3)
C. Group Project	25%
D. Final Exam	30%
TOTAL:	**100%**

A. Lab Submissions (15%)

Lab sessions are an essential part of the learning process. You are NOT going to learn without practicing data analysis and modeling yourselves. During lab sessions, I 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 and the group project. Lab submissions will be graded and count towards your final grade.

B. Homework Assignments (30%)

3 assignments will be given in the course. 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. Group Project (25%)

There will be one group project, with no more than 4 individuals in each group. Students are required to build a deep learning neural network model tackling a real-word NLP problems. A video demonstration is required as one of project deliverables. Details of the group project will be provided later in the semester.

D. Final Exam (30%)

The final examl is on Nov. 24, in classroom. More details on the final exam will be provided later.

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 **[ISOM4000B] 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	Assignments
1	Sep 1	Course Overview	
2	Sep 8	Machine Learning Basics I: Regression	
3	Sep 15	Machine Learning Basics II: Classification	
4	Sep 22	Machine Learning Basics III: Model Generalization and Regularization	HW1 release
5	Sep. 29	Deep Learning I: Architecture and Learning	
6	Oct. 6	Deep Learning II: Regularization and Optimization	HW1 due; HW2 release
7	Oct. 13	Recurrent Neural Networks	
8	Oct 20	Word Embeddings	HW2 due
9	Oct. 27	RNN in NLP: Sentiment Analysis and Language Models	HW3 release
10	Nov 3	RNN in NLP: Machine Transliation	
11	Nov 10	Attention-based Models and Transformers	HW3 due
12	Nov 17	Question Answering and BERT	
13	Nov 24	**Final Examination**	