

Deep Learning for Natural Language Processing

Course Objective

This course covers the basic theory and applications of language-based and multimodal interactive intelligent system using deep learning. Topics include distributional semantics with word embeddings, text classification, emotion recognition, sentiment analysis from text and speech, language modeling using various deep learning architecture such as neural networks and backpropagation, convolutional neural network, and recurrent neural network. We will also cover multi-linguality and multi-modality, end-to-end chatbot and task-oriented dialogue systems.

Students will learn about various natural language processing (NLP) tasks related to building interactive intelligent systems, how to design and build a deep learning model using PyTorch and Python, and the basics of software engineering. At the end of the course, students will work alone or in pairs to implement a deep learning model for various NLP research tasks, and describe their methods and results in a conference paper format. Successful groups will be able to submit their papers to an international conference.

Professor Pascale Fung

Course Instructor

Tel: (852) 2358 8537
Office: Room 2454 (Lift 25/26)
Email: pascale@ece.ust.hk

This course covers the basic theories and applications of natural language processing including speech processing using deep learning. It also introduces the topic of ethics in Artificial Intelligence. Topics include distributional semantics with word embeddings, text classification, emotion recognition, sentiment analysis from text and speech, language modeling, machine translation, and question answering. We will also cover multi-linguality and multi-modality, end-to-end chatbots, and task-oriented dialogue systems. The topics in Deep Learning ranges from feed forward neural networks, convolutional neural networks, recurrent neural networks, sequence-to-sequence neural networks, and the latest architecture in DL. Students will learn about various natural language processing (NLP) topics, how to design and build a deep learning model using PyTorch and Python, and the basics of software engineering. At the end of the course, students will work alone or in pairs to implement a deep learning model for various NLP research tasks, and describe their methods and results in a conference paper format. Successful groups will be able to submit their papers to an international conference.

Prof. Fung is a member of the Global Future Council on Artificial Intelligence and Robotics of the World Economic Forum, the current Vice-President and 2017 President of the Special Interest Group on Data-driven Methods of the Association for Computational Linguistics, and editor of the journal Computer Speech and Language, among other responsibilities.

Her full profile can be found at <http://www.ece.ust.hk/~pascale>

Time and Venue

Lectures

Tuesday 12:00 - 13:20
Thursday 12:00 - 13:20
On Zoom (link on Canvas)

Tutorials

Monday 9:30 - 10:20
On Zoom (link on Canvas)

Teaching Assistants

Wenliang Dai
Tiezhang Yu
Ziwei Ji
Office: Room 2602 (Lift 29/30)
TA Hour: 17:00 - 18:00, Wednesday
Email (For HW Submission and Inquiries): hkust.hlte.courses@gmail.com

References

- Aston Zhang, Zack Lipton, Mu Li and Alex Smola "Dive into Deep Learning"
- Ian Goodfellow, Yoshua Bengio and Aaron Courville, "Deep Learning" MIT press, 2016
- Chris Manning and Hinrich Schütze, "Foundations of Statistical Natural Language Processing" MIT Press, 1999
- Dan Jurafsky and James Martin, "Speech and Language Processing" Online Draft
- Yoav Goldberg, "Neural Network Methods for Natural Language Processing" Morgan & Claypool, 2017

Examples of Final Project

- Learning Knowledge Bases with Parameters for Task-Oriented Dialogue Systems
- Knowledge Embedded Framework for Sensible Question Answering
- Question Rewrites in Conversational QA
- Situational Conversation Generation in Movie Script
- A Combined BM25 and BERT model for CLEF Task-2 Claim Verification

Final Project 2021


[Final Project Instructions](#)

Schedule

Week	Dates	Topic
1	2 Feb / 4 Feb	Lecture 1: Artificial Intelligence and Natural Language Processing Lecture Recording (Part 1) (Part 2)
2	9 Feb / 11 Feb	Lecture 2: Regression and Neural Networks Lecture Recording (Part 1) (Part 2) Homework 1 (Due on 25 Feb, 23:59)
3	16 Feb / 18 Feb	Lecture 3: Linguistic Essentials Lecture Recording (Part 1) (Part 2) Tutorial 1: Basics of Python and Numpy + A Brief Intro to HW1
4	23 Feb / 25 Feb	Lecture 4: Lexical Semantics and Word Embeddings Lecture Recording (Part 1) (Part 2) Tutorial 2: Back Propagation Recap and PyTorch Basics Homework 2 (Due on 11 Mar, 23:59)
5	2 Mar / 4 Mar	Lecture 5: Convolutional Neural Networks and Applications Lecture Recording (Part 1) (Part 2) Tutorial 3: A Brief Intro to HW2
6	9 Mar / 11 Mar	Guest Lecture (Federated Learning by Prof. Yang Qiang) Tutorial 4: A Review on CNN Homework 3 (Due on 29 Mar, 23:59)
7	16 Mar / 18 Mar	Lecture 6: Language Modeling and Recurrent Neural Networks Tutorial 5: A Brief Intro to HW3
8	23 Mar / 25 Mar	Midterm Review Midterm Exam Tutorial 6: Midterm Instructions
9	30 Mar / 1 Apr	Online Webinar: Moving Away from One-Size-Fits-All Natural Language Processing by Prof. Rada Mihalcea (30 Mar) (Thu, 1 April is in the midterm break, so no lecture on that day) Tutorial 7: Project Discussion & A Review on Language Models, RNN Homework 4 (1) - Dialogue Annotation (Due on 15 Apr, 23:59) Homework 4 (2) - RNN (Optional, No Due Date)
10	6 Apr / 8 Apr	(no tutorial, and no lecture on Tue, 6 Apr, due to the midterm break) Lecture 7: Sequence-to-Sequence and Attention
11	13 Apr / 15 Apr	Lecture 8: Transformers and BERT Tutorial 8: Q&A about Midterm
12	20 Apr / 22 Apr	Lecture 9: Two-Minute Presentation of Final Project and Feedback Giving Lecture 10: Beneficial NLP Tutorial 9: A review of the Transformer and BERT
13	27 Apr / 29 Apr	Lecture 11: Neural Conversational AI Systems Tutorial 10: Q&A
14	4 May / 6 May	Final Project Presentations (normal lecture time, please check the Canvas announcement for the schedule of your group) Slides submission deadline: 7th May, 11:59 PM Report submission deadline: 18th May, 11:59 PM

Grading Scheme


- 10% Class and tutorial participation
- 40% Assignments 1-4 (10% each)
- 30% Midterm Exam
- 20% Final Project



Talking to Machines - Prof. Pascale Fung



Empathetic Machines - Prof. Pascale Fung



Moodbox - Prof. Pascale Fung

