

### Course Description

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. *Exclusion(s)*: COMP 4901I *Prerequisite(s)*: (COMP 2012 OR COMP 2012H) AND (ELEC 2600 OR ELEC 2600H)

### List of Topics

- Introduction to Interactive Intelligent Systems
- General Intro to Natural Language Processing
- Regression and Neural Networks
- Linguistic Essentials
- Lexical Semantics and Word Embeddings
- Convolutional Neural Networks Lecture
- Convolutional NN Applications in NLP, Vision and Art
- Language Modeling and Neural Networks in NLP
- Language Modeling and Recurrent Neural Networks in NLP
- Plug and Play Language Models (PPLM)
- Emo2Vec; Cross-lingual; Meta-Embedding; GlobalTrait
- Sequence-to-Sequence and Attention in NLP
- Transformers and BERT
- Introduction to Task-oriented Dialogue Systems
- Towards Neural Speech Recognition

### Statement of Objectives/Outcomes:

On completion of this course, students will be able to:

CO1 - Implement Deep Learning models for Natural Language Processing (NLP)

CO2 - Analyze related problems in Natural Language Processing and their possible solutions.

CO3 - Describe and understand concepts of Natural Language Processing and Deep Learning.

CO4 - Understand the basic theory behind Deep Learning and Machine Learning

CO5 - Have the basic skills for solving AI problems using Deep Learning

CO6 - Understand the broader topic of ethics in AI and beneficial AI

Reference Books/Materials:

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

Relationship of Course to Program Outcomes:

Please refer to the Report Section 4.3.2 (iii).

Grading Scheme:

Class and tutorial participation	10%
Assignments	40%
Midterm Exam	25%
Final Project	25%