Course Proposal: CS 260C Deep Learning

Course number: CS 260C
Course Catalog Title: Deep Learning
Short Title: Deep Learning
Units: 4 (4 is the standard)
Grading basis: Letter
Format: lecture with discussions
GE or Major/Minor requirement: N/A
Requisites: N/A
Recommended Requisites: CS 180, CS 260
Description:

In this course, we will teach the basics of deep neural networks and their applications, including but not limited to computer vision, natural language processing and graph mining. The course will cover topics including the foundation of deep learning, how to train a neural network (optimization), architecture designs for various tasks, and some other advanced topics. By the end of the course, the students are expected to be familiar with deep learning and be able to apply deep learning algorithms to a variety of tasks.

Justification:
Deep learning has become one of the most important topics in artificial intelligence and machine learning. However, the computer science department currently does not have a course specifically developed for deep learning. Due to this emergent need, the instructor (Dr. Hsieh) previously has offered deep learning lectures in his CS 260 class (Winter 2019). However, due to the original design of CS 260, which is a general-purpose machine learning course, the instructor can only allocate about 4 weeks to teaching deep learning. This is not enough for covering deep learning algorithms, and many students want to learn more about this topic both inside and outside the CS department. Therefore, we think it is worthwhile creating a deep learning course.

Supplemental Information:

Grading structure:
70% homework and 30% final project. A series of homework will guide students to implement their own deep learning models step-by-step on various tasks (computer vision, natural language processing and graph mining), and will be based on popular deep learning tools (PyTorch and TensorFlow). The final project will let students explore a particular topic that they are interested in.

Effective Date:

Winter

Syllabus:

Title of the course: Deep Learning

Course objective: In this course, we will learn about deep neural networks and their applications to various tasks, including but not limited to computer vision, natural language processing and graph mining. By the end of the course, the students are expected to be familiar with deep learning and be able to apply deep learning algorithms to a variety of tasks.

Detailed course description:

The course will cover the following topics:

Part I: Deep Learning Foundation (week 1 -- week 3.5)

We will first discuss the general concept of deep neural networks and why it can extract important knowledge from data. We will then discuss how to train a neural network in detail, including the detailed derivation of back-propagation and the foundation of stochastic optimization algorithms for training neural networks.

Lecture 1-2: introduction and basic definitions
Lecture 3-5: training algorithms and back-propagation
Lecture 6-7: regularizations and normalizations

Part II: Deep Learning in Various Tasks (Week 3.5 -- Week 8)
This will cover deep learning architectures in various domains, including computer vision, natural language processing, and graph mining. For computer vision tasks, we will introduce convolutional neural networks (CNN), residual network (ResNet), and generative adversarial network (GAN) for image generation. For natural language processing, we will introduce recurrent neural networks (RNN) and Transformer architectures, and discuss how to apply them in a variety of NLP tasks. For graph mining, we will introduce unsupervised and supervised graph neural networks. Finally, networks with multi-modality such as vision+text or text+graph will be discussed.

Lecture 8-10: Neural networks for computer vision (CNN, ResNet, GAN)
Lecture 11-13: Neural networks for NLP (RNN, Transformer)
Lecture 14-15: Neural networks for graph
Lecture 16: Multi-modal networks (text+vision or text+graph)

Part III: Advanced Topics in Deep Learning (Week 9 --Week 10)

We will talk about how several advanced topics including neural architecture search, meta-learning, and discuss the current limitations of deep learning (robustness, fairness, interpretability, scalability, reproducibility).

Lecture 17: Neural architecture search
Lecture 18: Meta-learning
Lecture 19-20: Limitations of deep learning (robustness, fairness, interpretability, scalability, reproducibility).

Grading structure:

Homework: 70%
Final project: 30%

There will be roughly 4 homework assignments to guide students to implement and train neural networks step-by-step. All the homework will be based on Pytorch or Tensorflow. The first 3 homeworks have been developed in the Winter 2021 offer of CS 260, including 1) implementation of single-layer network, 2) implementation of multi-layer fully connected and convolutional networks for image classification, 3) recurrent neural networks for transformers. We plan to design another homework on graph neural networks or meta learning for CS 260C.
At the end of the course, a final project will allow students to explore deeper for a particular area. They will survey and compare state-of-the-art methods or try to solve a real problem with deep learning.

Enrollment and student evaluation numbers from past course offering:
The instructor has delivered most of the material listed above in CS 260 (Winter 2019 and Winter 2021). There were more than 65 students enrolled for both quarters and the overall rating of the instructor/course is 7.93/7.67 (Winter 2019) and 8.28/8.39 (Winter 2021).