CSC 483: Applied Deep Learning

Tianxiang (Adam) Gao

School of Computing
DePaul University

Course Information

- **Sections**: **801** (Loop; in-class), **841** (Hybrid; Flex)
- **Lectures**:
  - Meeting Time: **Thursdays 5:45 PM – 9:00 PM**
  - Location: **CDM Center 224**
- **Instructor**: Tianxiang (Adam) Gao
  - Research: Deep learning theory, generative AI, graph representation learning
  - Office: CDM 712
  - Zoom: https://depaul.zoom.us/my/gaotx
  - Email: tgao9@depaul.edu
  - Office Hours: Mondays, 10:00–11:00 AM via Zoom, or by appointment (contact via email).
- **Course Website & Materials**:
  - D2L: https://d2l.depaul.edu/
  - Public: https://gaotx-cs.github.io/teaching/csc483/

Prerequisites

- **Prerequisites Courses**:
  - CSC 483: (CSC 412 and DSC 430 and DSC 441) or (CSC 412 and CSC 480)
- **Required Knowledge**:
  - Calculus: derivatives of multivariate functions, chain rules
  - Linear Algebra: vectors, matrices, matrix-vector computations,
  - Probability: mean, variance, Gaussian distribution, conditional probability, Bayes' rule
  - Programming (Python): list, loops, functions, Numpy, Matplotlib, Jupyter
- **Textbooks**:
  - Dive into Deep Learning, by Aston Zhang, Zack C. Lipton, Mu Li, Alex J. Smola.
  - Deep Learning Book, by Goodfellow, Bengio, and Courville.
  - The course is **self-contained** and complemented with research papers.

Course Overview

The course provides an essential introduction to neural networks and deep learning

- **Foundation of Deep Learning**: Multilayer perceptions, backpropagation, SGD
- **Advanced Optimization**: Momentum, RMSprop, Adam.
- **Generalization & Regularization**: Model complexity, overparameterization, weight decay.
- **Neural Network Architectures**: CNNs, RNNs, Transformers (e.g., GPT and BERT)
- **Applications**: CV, NLP, and Bio, such as object detection, Face recognition, and language models.

Course Outline

Week	Topic
Week 1	Introduction to Neural Networks
Week 2	Training Neural Networks
Week 3	Advanced Optimizers
Week 4	Generalization and Regularization
Week 5	Convolutional Neural Networks (CNNs)
Week 6	Learning in CNNs
Week 7	Recurrent Neural Networks (RNNs)
Week 8	Sequence-to-Sequence Models
Week 9	Attention Mechanisms and Transformers
Week 10	Large Language Models (LLM)
Week 11	No meeting

Course Objective

By the end of this course, you should be able to:

- Understand the fundamentals of deep neural networks
- Explore advanced topics such as convolutional neural networks (CNNs), recurrent neural networks (RNNs), and transformer architectures.
- Gain practical experience through assignments and projects implementing deep learning models in Python using popular frameworks (e.g., Keras or PyTorch).
- Learn how to apply deep learning techniques to real-world problems such as computer vision, natural language processing, and biomedical data science.

Coursework

- **Quiz (25%)**:
  - Quizzes are generally released on lecture days and will be due by the next lecture.
  - Approximately 8 quizzes will be given
  - We will count your best 5 out of them.
  - Each quiz will consist of around 10 multiple-choice questions ($\sim$0.5-1h).

- **Assignments (35%)**
  - These assignments are usually released on lecture days and are due on the next lecture.
  - Students should fill in the correct code for implementation rather than starting from scratch ($\sim$1-2h).
  - Download the Jupyter Notebook from D2L and fill in only the designated code sections.
  - Submit the notebook in its original format (.ipynb); do not change other parts of the file.
  - Any violation, such as altering the notebook or submitting a different format, will result in a **zero**.
  - Approximately 6 programming assignments will be given, and we will count your best 4 scores toward the final grade.

Coursework (Continued)

- **Student-Designed Assignment (25%):**
  - Team size: **2–3**
  - Design a deep learning assignment based on course **topics**
  - Must include: a clear learning task, appropriate dataset, justified model choice, evaluation criteria, instructions, and a grading rubric
  - Submitted assignments will be assigned to other students for completion and review

- **Peer-Reviewed Evaluation (20%):**
  - Each student reviews **2** student-designed assignments
  - Students must complete assigned assignments before submitting reviews
  - Reviews assess clarity, difficulty, topic alignment, evaluation quality, and educational value
  - Submission includes **3** documents: two completed assignments and their reviews

- **Bonus Points (5%, Optional):** Bonus points may be offered to encourage learning and may help improve borderline final grades (e.g., A– to A). Bonus points may be withheld for late, incomplete, or low-quality submissions.

Late Submissions

- All submissions are due by 11:59 PM on the posted due date.
- Standard assignments may be submitted up to **3 days late** with a **10% penalty per day**.
- Late submissions are **not accepted** for quizzes, student-designed assignments, peer-reviewed evaluations, or bonus-point assignments.

Grading Scale

Grade	Range
A	93-100
A-	90-92
B+	87-89
B	83-86
B-	80-82
C+	77-79
C	73-76
C-	70-72
D+	67-69
D	63-66
D-	60-62
F	0-59

University and Course Policies

- **Academic Integrity**
  - Cheating, plagiarism, and other forms of academic dishonesty are strictly prohibited.
  - Violations may result in failing the assignment, the course, or referral to the academic conduct committee.
  - Please ensure that all work submitted is your own and properly cited.
  - For detailed guidelines, refer to the university's academic integrity policy.
- **Other Policies**
  - We adhere to all university and course policies, including those on attendance, accommodations, and student conduct.
  - Detailed information on these policies can be found on the university policies website, or on the Teaching Commons page.

Note: For a comprehensive list of policies, please refer to the course **syllabus**.

Introduce Yourself (Bonus Point 1%)

Please introduce yourself by sharing your responses to the following prompts and post your introduction in the online discussion so we can all get to know each other better:

- **Name (and Preferred Pronouns)**
- **Major and Year**
- **Background (and Previous Experience)**
- **Research Interests**
- **Why You're Taking the Course**
- **Expectations for the Course**
- **Fun Fact or Hobby**

Discussion Questions

Questions to Discuss:

- What is Artificial Intelligence (AI)?
- What is Machine Learning (ML)?
- What is Deep Learning (DL)?
- What are the pros and cons of DL compared to conventional machine learning?

Instructions: Discuss these questions in small groups of 2-3 members.