Spring 2020

IS 698-002: Special Topic - Emerging Topics in Deep Learning - Artificial Intelligence

Hai Phan

IS:698 EMERGING TOPICS IN DEEP LEARNING - ARTIFICIAL INTELLIGENCE

Course Description: The last decade has seen a resurgence of deep neural networks as the machine learning tool of choice for many key application areas including software engineering, program analysis, computer vision, reinforcement learning, and natural language processing (NLP). This course offers a holistic and hands-on theoretical foundation to deep networks, their many varieties and applications, as well as the algorithms used to train them. Using real datasets and basic Python libraries (TensorFlow) for data manipulation, vector/matrix algebra, and automatic differentiation students will code up well-known examples of deep neural networks and apply them to applications in software engineering, program analysis, computer vision, NLP, privacy and security applications. Emergent areas of AI and deep learning, e.g., adversarial learning, interpretable machine learning, generative adversarial networks, etc., will be covered as well.

Learning Goals:

1. Obtain a good foundation of recent developments in deep learning

2. Learn and implement state-of-the-art deep learning models, addressing emergent challenges in real-world applications, i.e., software engineering, program analysis, image processing, NLP, privacy and security.

3. Gain hands-on multidisciplinary machine learning research experiences through open source toolkit, which is either publicly available or developed by our team.

4. Design and communicate project deliverables to solve specific deep learning-based problems and applications, i.e., adversarial learning, generative adversarial networks, interpretable deep learning, privacy-preserving machine learning.

5. Scientific paper reading, critique, presenting, and writing, based on the term project results.

Term: Spring 2020

Pre-requisites: IS 688

Instructors: Hai Phan and Shaohua Wang

Office: GITC 5115 and GITC 5118

Email: phan@njit.edu, davidsw@njit.edu

Wegpage: https://sites.google.com/site/ihaiphan/, https://davidshaohuawang.wordpress.com

Office Hours: Monday 8:30-6:30, Thursday 12-1pm

Canvas: Additional material and resources will be found on the class website on Canvas, (https://canvas.njit.edu). It will be modified and updated as the course progresses and will contain the most recent information.

Schedule: The following is a tentative schedule and subject to change. Refer to class web page for most recent information.

Course outline (15-week schedule) (14 actually week, as one week has no class due to spring break)

Weeks

1. Introduction to deep learning and TensorFlow + Neural Network Foundations (Hai) (01/27)
2. Shallow and Deep Neural Network + Recurrent Neural Networks + CNN (Hai) (02/03)
3. Adversarial Learning (project) + Interpretable Deep Learning (project) + Generative Adversarial Networks (project) (Hai) (02/10)
4. Overview-AI-based Software Engineering, NLP, and Big Code (project) (David) (02/17)

(in the following week, they give any thoughts/progress/updates on their project)

5. Project Continue (David) + Project Kick-off (4 from Hai and 4 from David): (02/24)
6. Students presentation started. Hai's related topics (Hai) (03/02)
7. David's topic. (03/09) (here, they need to submit a 1-page proposal about what they do. See the details below)
8. No Class (3/16) Spring break.
9. bug detection topics: student presenting (David)(3/23)
10. testing topics: student presenting (David) (3/30)
11. Hai's related topic topics (Hai) (4/6)
12. Hai's related topic (Hai) (4/13)
13. auto-fix topics (David) (4/20)
14. Hai's related topics (Hai) (4/27)
15. Final projects presentations or updates (Hai and David) (5/4)

One paper (presenter + discussant) and (discussant update projects) [3 papers / week, 1 topic]

Laboratory Sessions: This course does not have a separate laboratory session. However, some class meeting time throughout the semester will be dedicated to hands-on laboratory assignments. This work will be done using the computers in the classroom (if not, please bring your laptop). If necessary, laboratory assignments should be worked on outside the class time.

Credit: 3

Grade: Final Grades will be based on:

60% project

30% paper presentations, summaries, and participation in class (they need to be involved in the paper presentation each week)

10% pear evaluation

The final letter grades for the semester are based solely on the points you earn according to Table 2.

Grade	Points
A	90-
B+	86-90
B	80-85
C+	70-79
C	60-69
D	50-59
F	0-49

Table 2: The final letter grade converting table

Paper presentation and discussion (10%+5%+5%): Each paper will be assigned to one group of students who will act as a presenter (10%), another one group of students who will act as discussant (5%). The presentation will last 40-50 mins for presenter and 10 minutes for the discussant ***strict***. After the presentation from the presenter, one student will be picked randomly to use 1 minutes summarizing the paper. Afterwards, the discussion will last 10-20 mins (5% for discussion participation). Each group should upload the slides to Moodle before each class.

- **Role of presenter:** As a presenter you should not simply repeat the paper's content (remember you only have 40-50 mins), instead you should point out the main important findings of the work. You should highlight any novel contributions, any surprises, and other possible applications of the proposed techniques. You should check the authors' other work related to the presented paper. Finally, you should discuss how the presented work relate to other papers covered in the course (especially the papers covered in that particular week).
- **Role of discussant:** As a discussant, you should take an adversarial position by pointing out weak and controversial positions in the paper. You should present a short rebuttal of the paper. You should come prepared with problems and counterexamples for the presented work.

Your presentations should have (at least):

- describes the main technique that is used in the paper.
- lists the main contributions of the paper.
- places the paper relative to any recent work done by the authors of the paper.
- links places the paper relative to other papers presented that week.
- a listing of at least three technical points that you liked and three areas that should be improved.

Weekly critique and summary (10%): Each week, each student should pick one of the papers for that week and submit on easy chair a one-page critique of the paper before the start of class. The critique

should offer a brief summary of the paper, points in favor, points against, and comments for improvement. You do not need to submit a critique if you act as presenter in that week. Additional advice for critiquing papers is here. In addition to the critique, each student should also submit summaries for all the papers that will be presented in that week.

Project (60%=10%+50%): One original project (10 pages ACM format) done in a group of 4 (or 3) students. The project will explore one or more of the themes covered in the course. You need to submit a project proposal (2 pages ACM format). The proposal should provide a brief motivation of the project, a detailed discussion of the data and systems that will be used in the project, along with a timeline of milestones, and expected outcome. Make sure that you have cited at least 3 papers in your proposal. Additional advice for project proposals will be discussed in class. The proposal does not worth any grade. The goal is to ensure the feasibility of the project.

POLICIES:

Academic Integrity Policy

The NJIT academic honor code is located at: http://integrity.njit.edu/index.html. This honor code applies in its entirety to this class. Violations will not be tolerated. In addition, students should familiarize themselves with NJIT's "Best Practices related to Academic Integrity" which is developed and published on the Provost's website (on the policies page).

Disabilities

If you have a disability that may require some modification of seating, testing, or any other class requirement; please let the Professor know so that appropriate arrangements can be made. Similarly let the Professor know if you have any emergency medical information about which to be aware, or if you need special arrangements in the event of building evacuation. See the Professor after class hours or schedule an appointment. Assistance is available from the Office of Student Disability Services (205 Campbell Hall; 973-596-3420). Be sure and fill out appropriate paperwork with this office during the first week of class.