ECSE 6965: Advanced Topics in Probabilistic Deep Learning

Department of Electrical, Computer, and Systems Engineering
Rensselaer Polytechnic Institute

Course Information

Course Title: Advanced Topics in Probabilistic Deep Learning
Transcript Title: Probabilistic Deep Learning
Course Number: ECSE 6965
Semester and Year: Fall, 2021
Credit Hours: 2
Meeting Times: 5-6:50 pm, Tuesdays.
Classroom: JONSSN 6309

Instructor

Prof. Qiang Ji
Office: JEC 7002
Office Phone: x6440
Email: jiq@rpi.edu
Office Hours: TBD
TA(s): TBD

Course Description

Current deterministic deep learning models tend to overfit, cannot effectively quantify their prediction uncertainty, and involve a complex heuristic training process. Probabilistic deep learning represents the cutting-edge research in deep learning that can effectively address these limitations.

The goal of this course is to introduce probabilistic deep learning and learn about the latest developments in probabilistic deep learning and it's applications. Specifically, the course will focus on three types of probabilistic deep learning models: probabilistic deep neural networks (PDNNs), Bayesian deep neural networks (BDNNs) and deep probabilistic graphical models (DPGMs). The course will start with an introduction of the theories to the three probabilistic deep learning models, the associated learning and inference methods, and various challenges facing these models. The second part of this class will be devoted to student presentations, whereby students review the top AI/ML conferences to identify the latest developments in the three types of models, and their applications. Each student will select 3 papers and give a detailed discussion of the work. Besides student presentations, we will also invite guest speakers from industry and academics to give talks on the related topics.

Pre-requisites

ECSE 6850/4850 or equivalent

Textbooks

Lecture notes distributed during class

Other References or Supplements

- Probabilistic graphical models for computer vision, Qiang Ji, Elsevier, 2020.
- Conference on Uncertainty in Artificial Intelligence
- Annual Conference on Neural Information Processing Systems
- International conference on machine learning

Student Learning Outcomes: Students who successfully complete this course will be able to

1. Understand the fundamental theories for probabilistic deep learning (PDL)
2. Understand the computational challenges with PDL and the associated solutions
3. Learn the latest developments in PDL and their applications
4. Be able to apply PDL algorithms to their own projects

Course Assessment Measures

Course assessment will be done through class participation and presentations.

Presentations: each student will be required to select papers from the top AI/ML conferences and give a detailed and professional presentation of these work, including identifying the purposes of these papers, their technical contributions, their performances, and possible improvements.

Grading policy for 6xxx:

1. Class attendance and active participation — 25% (earn credits by actively engaging in class discussions; lose credits for passive, not paying attention, and missing classes.)
2. 3 presentations — 75% (presentation quality, knowledge in the topics, and ability to field questions)

Course Topics and Schedule (Subject to change)

Week	Topic
Week 1	Introduction
Week 2	Fundamentals
Week 2	Probabilistic Deep Neural Networks : discriminative deep neural works
Week 3	Probabilistic Deep Neural Networks : generative deep neural works
Week 3	Bayesian deep neural networks
Week 3	Bayesian deep neural networks
Week 4	Bayesian deep neural networks
Week 5	Deep probabilistic graphical models
Week 6	Deep probabilistic graphical models
Week 7	Deep probabilistic graphical models
Week 8	Student presentations
Week 9	Student presentations
Week 10	Student presentations
Week 11	Student presentations
Week 12	Student presentations
Week 13	Student presentations
Week 14	Student presentations

Academic Integrity

Student-teacher relationships are built on trust. For example, students must trust that teachers have made appropriate decisions about the structure and content of the courses they teach, and teachers must trust that the assignments that students turn in are their own. Acts that violate this trust undermine the educational process. The Rensselaer Handbook of Student Rights and Responsibilities and The Graduate Student Supplement define various forms of Academic Dishonesty and you should make yourself familiar with these. In this class, all assignments that are turned in for a grade must represent the student's own work. In cases where help was received, or teamwork was allowed, a notation on the assignment should indicate your collaboration.

In this class, all presentations given must represent the student's own work. In cases where materials from others are used, a citation/credit must be clearly indicated to identify the source of the materials. Any violation of this policy will result in a *penalty of removing the entire grade of the presentation*. Repeat violations will be reported to the office of the Dean of Students.