

EECS 598: Special Topics on Natural Language Processing with Deep Learning

Professor Lu Wang

Fall 2020

E-mail: wangluxy@umich.edu

Web: <https://web.eecs.umich.edu/~wangluxy>

Office Hours: Wednesdays 5pm - 6pm

online via Zoom

Class Hours: Mondays and Wednesdays 12pm - 1:30pm

online via Zoom

Course Description

Deep learning models have made impressive progress in natural language understanding and generation problems, such as automatic text summarization and dialogue systems. In this **research-oriented** seminar course, we will focus on the discussion of recent advances of using deep learning models for solving **natural language generation (NLG)** problems. Topics include text summarization, long-form text generation, dialogue systems, question generation, NLG evaluation, and bias and ethics issues in generation systems. The class takes the following format: Students will present and discuss papers, and work on course projects in small groups. After taking this course, students are expected to gain knowledge in major tasks and the state-of-the-arts in NLG, how to design and evaluate NLG systems, how to evaluate research papers in the area, and how to implement and evaluate their own idea.

Course and Discussion Websites

Course Website: https://web.eecs.umich.edu/~wangluxy/courses/eecs598_fa2020/eecs598_fa2020.html

Discussion Forum: <https://piazza.com/umich/fall2020/eecs598017/home>

Required Materials

This course does not have any required textbook, but you can refer to the following resources.

- Jacob Eisenstein, "Introduction to Natural Language Processing", The MIT Press, 2019
- Dan Jurafsky and James H. Martin, "Speech and Language Processing, 2nd Edition", Prentice Hall, 2009. Third edition draft is available at web.stanford.edu/~jurafsky/slp3/.

Prerequisites/Corequisites

This course is designed for graduate students and senior undergraduate students majoring computer science, computational linguistics, and other related areas. Concretely,

- Programming: Students are expected to be proficient in some programming languages (Python is encouraged).
- Natural language processing/Machine learning knowledge: Students should have familiarity with natural language processing concepts and machine learning fundamentals, e.g., have done projects with machine learning tools to train and evaluate computational and statistical models.

Course Structure

Grading Policy

Each assignment or report is due by the end of day on the corresponding due date (i.e. 11:59pm, EST). Canvas is used for electronic submission. Assignment or report turned in late will be charged 20 points (out of 100 points) off for each late day (i.e. every 24 hours). Each student has a budget of 5 days throughout the semester before a late penalty is applied. You may want to use it wisely, e.g. save for emergencies. Each 24 hours or part thereof that a submission is late uses up one full late day. Late days are not applicable to final presentation. Each group member is charged with the same number of late days, if any, for their submission. There is no need to inform the instructors if late days are used; timestamp of the last submission on Canvas will be used for automatic grade calculation.

Grades will be determined based on paper critiques and presentation, project, and participation:

- Paper critiques (38%): write a short critique (using this [template]) for each paper, each submission worth 2%
- Paper presentation and discussion leading (10%): each student will present papers and lead corresponding discussions twice (with another student or individually), each of 5%
- Project (40%): team of 2 to 3 students, proposal (5%), reports (10%+15%, mid-term and final), presentations (5%+5%, mid-term and final)
- Project feedback (6%): each student will write short feedback to other groups' projects after mid-term presentation
- Participation (6%): attendance, participating in-class discussions, etc

Schedule and Weekly Learning Goals

The schedule is tentative and subject to change. Please check the course website for most recent schedule with readings.

Week 01, 08/31 - 09/04: Topic: Introduction, basics of neural models in NLP, neural abstractive summarization

Week 02, 09/07 - 09/11: Topic: reinforcement learning for summarization

Week 03, 09/14 - 09/18: Topic: unsupervised summarization

Week 04, 09/21 - 09/25: Topic: errors in summaries and evaluation, large language models

Week 05, 09/28 - 10/02: Topic: knowledge-enhanced generation, content planning for long-text generation

Week 06, 10/05 - 10/09: Topic: non-monotonic decoding, issues in neural generation

Week 07, 10/12 - 10/16: Project Mid-term Presentation

Week 08, 10/19 - 10/23: Topic: goal-oriented dialogues with knowledge augmentation, personalized dialogue agents

Week 09, 10/26 - 10/30: Topic: chitchat, factoid question generation

Week 10, 11/02 - 11/06: Topic: open-ended/complex question generation, issues in automatic metrics

Week 11, 11/09 - 11/13: Topic: reference-based evaluation, reference-free evaluation

Week 12, 11/16 - 11/20: Topic: can human tell machine-generated text, ethics in neural generation

Week 13, 11/23 - 11/27: Thanksgiving week, no class

Week 14, 11/30 - 12/04: Topic: Group presentation

Week 15, 12/07 - 12/11: Topic: Group presentation