Course Information

CS 6956: Deep Learning for NLP

What we will see

- A general overview of underlying concepts that pervade deep learning for NLP tasks

- A collection of successful design ideas to handle sparse, compositional varying sized inputs and outputs

Class focus

- We will see several example NLP tasks
  - Language modeling
  - Sequence prediction for semantic role labeling
  - Natural language inference and reading comprehension
  - Machine translation
  - And most importantly,

    Your favorite domain/problem…

- To understand underlying *concepts*
  - Defining models, training, prediction
  - Choosing the right architecture for a problem
  - Tricks and tips

- Also, applications
  - TensorFlow or PyTorch for programming homeworks

Course objectives

At the end of the course, you should be able to:

1. Define deep neural networks for new NLP problems,

2. Implement and train such models using off-the-shelf libraries, and

3. Be able to critically read, evaluate and perhaps replicate current literature in the field.

Course mechanics

Course website: https://svivek.com/teaching/deep-learning-nlp

- Course structure
  - Lectures by me initially and gradually, presentations by you
- No official text book
  - Many lectures will follow Yoav Goldberg's textbook
  - Useful background reading on course website
- Pre-requisites: Machine Learning and NLP
- Assignments *(due dates on schedule page of website)*
  1. 3-4 assignments (not hand written, please!)
  2. One class presentation
  3. One class project in groups of size at most two
  4. No midterm/final. Instead, project proposal, intermediate checkpoints, final report and presentation

Questions?

Assignments

Three kinds of assignments:

- Coding assignments:
  - We will use Google's Colaboratory
  - You will submit Jupyter notebooks for your assignments
- Theory:
  - Will be somewhat on the simpler side
- Paper review: You will pick a paper from a list and write a review for it

What assistance is available for you?

Course website: https://svivek.com/teaching/structured-prediction

We will use

Canvas for:

1. Announcements and communication
2. Discussion board
3. All submissions

Course website for:

1. Lecture slides
2. Notes and readings

Staff

Email: `svivek at cs.utah.edu`

Office hours:

Wed 2:00 PM, 3126 MEB,

or by appointment

Please prefix subjects of all emails with course number

Policies (see website for details)

- This class operates under the School of Computing policies and guidelines.
- Collaboration vs. Cheating
  - Collaboration is strongly encouraged, cheating will not be tolerated
  - The School of Computing policy on academic misconduct
  - Acknowledge sources and discussions in all deliverables
- Late policy
  - 10 % penalty if submitted one day late, no further extensions
- Access and assistance
  - If you need any assistance, please contact me as soon as possible

Questions?

Course expectations

This is an advanced course aimed at helping you navigate recent research.

I expect you to

- Participate in the class
- Complete the readings for the lectures
- And most importantly, demonstrate independence and mathematical rigor in your work

- No readings for next lecture
- For questions about registration, please meet me now