Next Classes and Conclusion

Robin Jia
USC CSCI 467, Spring 2024
April 25, 2024

Announcements/Reminders

- HW4 due today at 11:59pm
  - Using up to 3 late days is allowed
- Section tomorrow: Final review
- Final exam Tuesday, May 7, 2-4pm
  - Please use pen! (For scanning/grading purposes)
  - Do not write on the back of the exam, we added overflow space at the end (backs are OK for ungraded scratch work)
- My OH tomorrow moved to 4-5pm (on the calendar)
- Office hours next week: Check the calendar
- Final project May 3

Next classes to take next semester

- Natural Language Processing
- Computer Vision
- Robotics
- Theory of Machine Learning

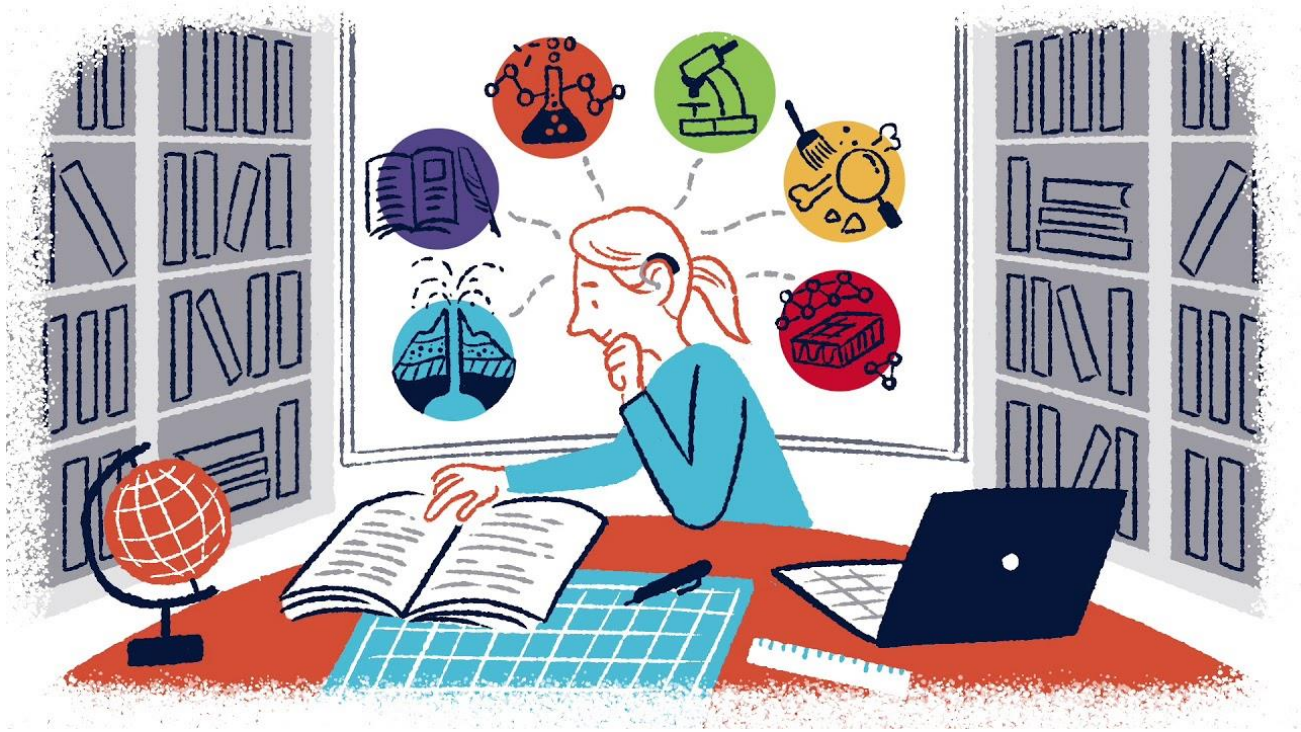

Natural Language Processing

- **CSCI 499: Natural Language Processing (Jesse Thomason)**
  - Undergraduate NLP class, focused on language models
- **CSCI 544: Applied NLP (Swabha Swayamdipta)**
  - General NLP class, recently updated to focus more on language models
- **CSCI 626: Text as Data (Morteza Dehghani)**
  - Applications of natural language processing to psychology research
- **CSCI 662: Advanced NLP (Jon May)**
  - Research-focused class, covering machine translation, dialogue, question answering, information extraction, etc.

Computer Vision

- **CSCI 677: Advanced Computer Vision (Yue Wang)**
  - Foundational computer vision topics (geometry)
  - Deep learning for computer vision
  - Standard tasks (object detection, semantic segmentation, motion analysis, activity recognition, visual question answering)

Robotics

- (Not exactly ML topics but highly related)
- **CSCI 445L: Introduction to Robotics (Heather Culbertson)**
  - Hands-on introduction to robotics, will work with real physical robots
- **CSCI 545: Robotics (Stefanos Nikolaidis)**
  - More advanced course covering control theory, kinematics, dynamics, sensor processing
  - Seems to get more into the math, which involves a lot of linear algebra

Math & Machine Learning

- **CSCI 678: Theoretical Machine Learning (Haipeng Luo)**
  - How can you prove a learning algorithm will have good test accuracy?
  - How can you prove that regularization improves test accuracy?
  - How can you prove that UCB/Q-learning learn optimal policies for bandits/RL?
- **Math 447: Mathematics of Machine Learning**
  - Math-focused class on machine learning
  - More detail on some topics like kernels/SVMs
  - Also will discuss ways to prove good test accuracy
- **CSCI 458: Numerical Methods**
  - More on methods like Newton's method, computing eigenvectors/eigenvalues, etc.
- **CSCI 612: Optimization for the Information and Data Sciences**
  - Class on advanced optimization methods for convex problems
  - Goes beyond gradient descent to methods with faster convergence (e.g., how SVM's actually work)
  - How to handle constraints during optimization

Classes that may be offered later…

- **CSCI 461: AI for Sustainable Development (Bistra Dilkina)**
  - Project-based class focusing on AI for social good (e.g., sustainability, poverty, homelessness, health)
  - Includes discussions of research papers
- **CSCI 499: Foundations of Multi-Agent Systems (Sven Koenig)**
  - In class: Reinforcement learning involves 1 agent interacting with an environment
  - Often times, there are many agents interacting simultaneously with an environment + each other (e.g., multiple robots)
  - Agents have to learn, communicate, reason about other agents (game theory), etc.

Other classes of note

- **CSCI 566: Deep Learning and Its Applications**
  - Some overlap with 467 but will focus exclusively on deep learning, go into some more detail
- **CSCI 567: Machine Learning**
  - Significant overlap with 467, focusing on non-deep learning
  - Likely to be more theoretical
- Could take these but other classes are recommended more, due to the degree of overlap

That's it!

- Thank you for a wonderful semester!