Thinking about your Final Projects

Misha

Why Start Now?

- Mentor check-ins start early in the second half of the semester

- Opportunity to:
  - review literature
  - have ideas in your head
  - find partners

- 30% of your final grade

Types of Projects

- Model:
  - propose a new model, analyze it, and test it on several tasks for validation

- Task:
  - use available tools to try and do as well as possible on a certain task

- Analysis:
  - analyze existing models or tasks to understand what they tell us

Example of Model Project

- Yang, Lu, & Zheng. A Simple Regularization-based Algorithm for Learning Cross-Domain Word Embeddings. EMNLP 2017.

- Motivate and describe their model:
  - Learn small-corpus embeddings via transfer learning
  - Use existing large-corpus embeddings as regularization

- Evaluate the model's usefulness on several tasks:
  - Entity recognition
  - Sentiment analysis

Model Project Details

- Models should be well-founded: major design decisions should have some linguistic or mathematical motivation.

- Models will most-likely build upon existing work, and so must be placed in that context.

- Empirical verification should be conducted on model assumptions, not just downstream tasks.

- Models can help improve performance on downstream tasks by augmenting existing methods.

Example of Task Project

- Herbelot & Baroni. High-risk learning: Acquiring new word vectors from tiny data. EMNLP 2017.

- Motivate and describe the task:
  - Learn a word vector using only the word's definition and vectors of other words.
  - Simulates one-shot learning of word embeddings when we don't have enough data about a word but still want to represent it.

- Devise a method to do well on this task:
  - Modify word2vec algorithm to learn quickly from one example.
  - Compare method to baseline/previous methods.

Task Project Details

- If introducing a task, need to show that previous tasks are insufficient and come up with baseline methods.

- If approach performs better/worse than existing methods, give some reasons as to why.

- Be careful not to overfit - i.e. devise a method that only works well on your model because you have been evaluating using the test set.

Example of Analysis Project

- Chen, Bolton, & Manning. A Thorough Examination of the CNN/ Daily Mail Reading Comprehension Task. ACL 2016.

- Describe what they analyze and demonstrate its importance:
  - Reading comprehension is an important goal in NLP
  - Efforts often evaluated on the CNN/Daily Mail task, appealing due to its size and simplicity

- Perform the analysis and gain new insights:
  - Design a very simple system that does well on the task, indicating that it is perhaps too easy
  - Doing better than the state-of-the-art may be impossible due to annotation/co-reference errors in the original task

Analysis Project Details

- Choose an existing model/task for which you believe present understanding is insufficient.

- Sometimes goal is to ‘break’ the model or show that the task may not be that useful by demonstrating points of failure.

- If the outcome is negative for the subject of consideration, try to introduce alternatives.

Closing Thoughts

- Project categories have a lot of overlap and good reports will often have components of each.

- Negative results are okay, but hopefully they lead to better understanding.

- Attend the many NLP colloquia this spring (see course schedule for times).