

Final projects

CS 685, Fall 2020

Advanced Natural Language Processing
<http://people.cs.umass.edu/~miyyer/cs685/>

Mohit Iyyer

College of Information and Computer Sciences
University of Massachusetts Amherst

Timeline

- All groups will be formed by *Sep 7*
- Only two deliverables:
 - Project proposal: 2-4 pages, due *Sep 21*
 - Final report: 12+ pages, due *Dec 4*
- Almost completely open-ended!
 - All projects must involve natural language data
 - All projects should include at least some degree of model implementation

Project

- Either *build* natural language processing systems, or *apply* them for some task.
- Use or develop a dataset. Report empirical results or analyses with it.
- Different possible areas of focus
 - Implementation & development of algorithms
 - Defining a new task or applying a linguistic formalism
 - Exploring a dataset or task

Formulating a proposal

- What is the **research question**?
- What's been done before?
- What experiments will you do?
- How will you know whether it worked?
 - If data: held-out accuracy
 - If no data: manual evaluation of system output.
Or, annotate new data

Feel free to be ambitious (in fact, we explicitly encourage creative ideas)! Your project doesn't necessarily have to "work" to get a good grade.

The Heilmeier Catechism

- What are you trying to do? Articulate your objectives using absolutely no jargon.
- How is it done today, and what are the limits of current practice?
- What is new in your approach and why do you think it will be successful?
- Who cares? If you are successful, what difference will it make?
- What are the risks?
- How much will it cost?
- How long will it take?
- What are the mid-term and final “exams” to check for success?

An example proposal

- Introduction / problem statement
- Motivation (why should we care? why is this problem interesting?)
- Literature review (what has prev. been done?)
- Possible datasets
- Evaluation
- Tools and resources
- Project milestones / tentative schedule

NLP Research

- All the best publications in NLP are open access!
 - Conference proceedings: ACL, EMNLP, NAACL (EACL, LREC...)
 - Journals: TACL, CL
 - “aclweb”: ACL Anthology-hosted papers
<http://aclweb.org/anthology/>
 - NLP-related work appears in other journals/conferences too: data mining (KDD), machine learning (ICML, NIPS), AI (AAAI), information retrieval (SIGIR, CIKM), social sciences (Text as Data), etc.
- Reading tips
 - Google Scholar
 - Find papers
 - See paper’s number of citations (imperfect but useful correlate of paper quality) and what later papers cite it
 - [... or SemanticScholar...]
 - For topic X: search e.g. [[nlp X]], [[aclweb X]], [[acl X]], [[X research]]...
 - Authors’ webpages
find researchers who are good at writing and whose work you like
 - Misc. NLP research reading tips:
<http://idibon.com/top-nlp-conferences-journals/>

A few examples

We will post some sample project reports from previous semesters after getting student permission

A few examples

We will post some sample project reports from previous semesters after getting student permission

- Detection tasks
 - Sentiment detection
 - Sarcasm and humor detection
 - Emoticon detection / learning

A few examples

We will post some sample project reports from previous semesters after getting student permission

- Detection tasks
 - Sentiment detection
 - Sarcasm and humor detection
 - Emoticon detection / learning
- Structured linguistic prediction
 - Targeted sentiment analysis (i liked ___ but hated ___)
 - Relation, event extraction (who did what to whom)
 - Narrative chain extraction
 - Parsing (syntax, semantics, discourse...)

A few examples

We will post some sample project reports from previous semesters after getting student permission

- Detection tasks
 - Sentiment detection
 - Sarcasm and humor detection
 - Emoticon detection / learning
- Structured linguistic prediction
 - Targeted sentiment analysis (i liked ___ but hated ___)
 - Relation, event extraction (who did what to whom)
 - Narrative chain extraction
 - Parsing (syntax, semantics, discourse...)
- Text generation tasks
 - Machine translation
 - Document summarization
 - Story generation
 - Text normalization / “style transfer” (e.g. translate online/Twitter text to standardized English)

A few examples

We will post some sample project reports from previous semesters after getting student permission

- Detection tasks
 - Sentiment detection
 - Sarcasm and humor detection
 - Emoticon detection / learning
- Structured linguistic prediction
 - Targeted sentiment analysis (i liked ___ but hated ___)
 - Relation, event extraction (who did what to whom)
 - Narrative chain extraction
 - Parsing (syntax, semantics, discourse...)
- Text generation tasks
 - Machine translation
 - Document summarization
 - Story generation
 - Text normalization / “style transfer” (e.g. translate online/Twitter text to standardized English)
- End to end systems

A few examples

We will post some sample project reports from previous semesters after getting student permission

- Detection tasks
 - Sentiment detection
 - Sarcasm and humor detection
 - Emoticon detection / learning
- Structured linguistic prediction
 - Targeted sentiment analysis (i liked ___ but hated ___)
 - Relation, event extraction (who did what to whom)
 - Narrative chain extraction
 - Parsing (syntax, semantics, discourse...)
- Text generation tasks
 - Machine translation
 - Document summarization
 - Story generation
 - Text normalization / “style transfer” (e.g. translate online/Twitter text to standardized English)
- End to end systems
 - Question answering

A few examples

We will post some sample project reports from previous semesters after getting student permission

- Detection tasks
 - Sentiment detection
 - Sarcasm and humor detection
 - Emoticon detection / learning
- Structured linguistic prediction
 - Targeted sentiment analysis (i liked ___ but hated ___)
 - Relation, event extraction (who did what to whom)
 - Narrative chain extraction
 - Parsing (syntax, semantics, discourse...)
- Text generation tasks
 - Machine translation
 - Document summarization
 - Story generation
 - Text normalization / “style transfer” (e.g. translate online/Twitter text to standardized English)
- End to end systems
 - Question answering
 - Conversational dialogue systems (hard to eval?)

A few examples

We will post some sample project reports from previous semesters after getting student permission

- Detection tasks
 - Sentiment detection
 - Sarcasm and humor detection
 - Emoticon detection / learning
- Structured linguistic prediction
 - Targeted sentiment analysis (i liked ___ but hated ___)
 - Relation, event extraction (who did what to whom)
 - Narrative chain extraction
 - Parsing (syntax, semantics, discourse...)
- Text generation tasks
 - Machine translation
 - Document summarization
 - Story generation
 - Text normalization / “style transfer” (e.g. translate online/Twitter text to standardized English)
- End to end systems
 - Question answering
 - Conversational dialogue systems (hard to eval?)
- Predict external things from text

A few examples

We will post some sample project reports from previous semesters after getting student permission

- Detection tasks
 - Sentiment detection
 - Sarcasm and humor detection
 - Emoticon detection / learning
- Structured linguistic prediction
 - Targeted sentiment analysis (i liked ___ but hated ___)
 - Relation, event extraction (who did what to whom)
 - Narrative chain extraction
 - Parsing (syntax, semantics, discourse...)
- Text generation tasks
 - Machine translation
 - Document summarization
 - Story generation
 - Text normalization / “style transfer” (e.g. translate online/Twitter text to standardized English)
- End to end systems
 - Question answering
 - Conversational dialogue systems (hard to eval?)
- Predict external things from text
 - Movie revenues based on movie reviews ... or online buzz? [http://www.cs.cmu.edu/~ark/movie\\$-data/](http://www.cs.cmu.edu/~ark/movie$-data/)

A few examples

We will post some sample project reports from previous semesters after getting student permission

- Detection tasks
 - Sentiment detection
 - Sarcasm and humor detection
 - Emoticon detection / learning
- Structured linguistic prediction
 - Targeted sentiment analysis (i liked ___ but hated ___)
 - Relation, event extraction (who did what to whom)
 - Narrative chain extraction
 - Parsing (syntax, semantics, discourse...)
- Text generation tasks
 - Machine translation
 - Document summarization
 - Story generation
 - Text normalization / “style transfer” (e.g. translate online/Twitter text to standardized English)
- End to end systems
 - Question answering
 - Conversational dialogue systems (hard to eval?)
- Predict external things from text
 - Movie revenues based on movie reviews ... or online buzz? [http://www.cs.cmu.edu/~ark/movie\\$-data/](http://www.cs.cmu.edu/~ark/movie$-data/)
- Visualization and exploration (harder to evaluate)

A few examples

We will post some sample project reports from previous semesters after getting student permission

- Detection tasks
 - Sentiment detection
 - Sarcasm and humor detection
 - Emoticon detection / learning
- Structured linguistic prediction
 - Targeted sentiment analysis (i liked ___ but hated ___)
 - Relation, event extraction (who did what to whom)
 - Narrative chain extraction
 - Parsing (syntax, semantics, discourse...)
- Text generation tasks
 - Machine translation
 - Document summarization
 - Story generation
 - Text normalization / “style transfer” (e.g. translate online/Twitter text to standardized English)
- End to end systems
 - Question answering
 - Conversational dialogue systems (hard to eval?)
- Predict external things from text
 - Movie revenues based on movie reviews ... or online buzz? [http://www.cs.cmu.edu/~ark/movie\\$-data/](http://www.cs.cmu.edu/~ark/movie$-data/)
- Visualization and exploration (harder to evaluate)
- Temporal analysis of events, show on timeline

A few examples

We will post some sample project reports from previous semesters after getting student permission

- Detection tasks
 - Sentiment detection
 - Sarcasm and humor detection
 - Emoticon detection / learning
- Structured linguistic prediction
 - Targeted sentiment analysis (i liked ___ but hated ___)
 - Relation, event extraction (who did what to whom)
 - Narrative chain extraction
 - Parsing (syntax, semantics, discourse...)
- Text generation tasks
 - Machine translation
 - Document summarization
 - Story generation
 - Text normalization / “style transfer” (e.g. translate online/Twitter text to standardized English)
- End to end systems
 - Question answering
 - Conversational dialogue systems (hard to eval?)
- Predict external things from text
 - Movie revenues based on movie reviews ... or online buzz? [http://www.cs.cmu.edu/~ark/movie\\$-data/](http://www.cs.cmu.edu/~ark/movie$-data/)
- Visualization and exploration (harder to evaluate)
 - Temporal analysis of events, show on timeline
 - Topic models: cluster and explore documents

A few examples

We will post some sample project reports from previous semesters after getting student permission

- Detection tasks
 - Sentiment detection
 - Sarcasm and humor detection
 - Emoticon detection / learning
- Structured linguistic prediction
 - Targeted sentiment analysis (i liked ___ but hated ___)
 - Relation, event extraction (who did what to whom)
 - Narrative chain extraction
 - Parsing (syntax, semantics, discourse...)
- Text generation tasks
 - Machine translation
 - Document summarization
 - Story generation
 - Text normalization / “style transfer” (e.g. translate online/Twitter text to standardized English)
- End to end systems
 - Question answering
 - Conversational dialogue systems (hard to eval?)
- Predict external things from text
 - Movie revenues based on movie reviews ... or online buzz? [http://www.cs.cmu.edu/~ark/movie\\$-data/](http://www.cs.cmu.edu/~ark/movie$-data/)
- Visualization and exploration (harder to evaluate)
 - Temporal analysis of events, show on timeline
 - Topic models: cluster and explore documents
- Figure out a task with a cool dataset

A few examples

We will post some sample project reports from previous semesters after getting student permission

- Detection tasks
 - Sentiment detection
 - Sarcasm and humor detection
 - Emoticon detection / learning
- Structured linguistic prediction
 - Targeted sentiment analysis (i liked ___ but hated ___)
 - Relation, event extraction (who did what to whom)
 - Narrative chain extraction
 - Parsing (syntax, semantics, discourse...)
- Text generation tasks
 - Machine translation
 - Document summarization
 - Story generation
 - Text normalization / “style transfer” (e.g. translate online/Twitter text to standardized English)
- End to end systems
 - Question answering
 - Conversational dialogue systems (hard to eval?)
- Predict external things from text
 - Movie revenues based on movie reviews ... or online buzz? [http://www.cs.cmu.edu/~ark/movie\\$-data/](http://www.cs.cmu.edu/~ark/movie$-data/)
- Visualization and exploration (harder to evaluate)
 - Temporal analysis of events, show on timeline
 - Topic models: cluster and explore documents
- Figure out a task with a cool dataset
 - e.g. Urban Dictionary

Sources of data

- All projects must use (or make, and use) a textual dataset. Many possibilities.
 - For some projects, creating the dataset may be a large portion of the work; for others, just download and more work on the system/modeling side
- SemEval and CoNLL Shared Tasks:
dozens of datasets/tasks with labeled NLP annotations
 - Sentiment, NER, Coreference, Textual Similarity, Syntactic Parsing, Discourse Parsing, and many other things...
 - e.g. SemEval 2015 ... CoNLL Shared Task 2015 ...
 - <https://en.wikipedia.org/wiki/SemEval> (many per year)
 - <http://ifarm.nl/signll/conll/> (one per year)
- General text data (not necessarily task specific)
 - Books (e.g. Project Gutenberg)
 - Reviews (e.g. Yelp Academic Dataset https://www.yelp.com/academic_dataset)
 - Web
 - Tweets

Tools

- Tagging, parsing, NER, coref, ...
 - Stanford CoreNLP <http://nlp.stanford.edu/software/corenlp.shtml>
 - spaCy (English-only, no coref) <http://spacy.io/>
 - Twitter-specific tools (ARK, GATE)
- Many other tools and resources
 - tools* ... word segmentation ... morph analyzers ...
 - resources* ... pronunciation dictionaries ... wordnet, word embeddings, word clusters ...
- Long list of NLP resources
<https://medium.com/@joshdotai/a-curated-list-of-speech-and-natural-language-processing-resources-4d89f94c032a>
- Deep learning? Try out AllenNLP, PyTorch, Tensorflow (<https://allennlp.org>, <https://pytorch.org/>, <https://www.tensorflow.org/>)