

CSCI 662 Fall 2019 course page

Jonathan May

August 28, 2019

Website	https://www.isi.edu/~jonmay/cs662_fa19_web/
Lectures	LVL (Leavy Library) 16, Mondays and Wednesdays 3:30-5:20 pm
Instructor & office hours	Jonathan May, LVL (exact location on piazza), Mondays and Wednesdays 2-3pm or by appointment
TA & Office hours	Samar Haider, TBD
Textbook	Required: Natural Language Processing - Eisenstein ¹ Selected Papers: From NLP literature, see (evolving) schedule <i>Optional:</i> Introduction to Deep Learning - Charniak ² <i>Optional:</i> Speech and Language Processing 3rd edition -Jurafsky, Martin ³
Grading	10 %: In-class participation 10 %: In-class quizzes 40 %: Four Homeworks 40 %: Project, including extended abstract, final conference-quality paper, and 20-minute in-class presentation (may be done in small groups)
Contact us	On Piazza or in class/office hours. Please do not email (unless notified otherwise).

Topics (subject to change per instructor/class whim) (will not be presented in this order):

- Linguistic Stack (graphemes/phones - words - syntax - semantics - pragmatics - discourse)
- Tools:
 - Corpora, Corpus statistics, Data cleaning and munging
 - Annotation and crowdwork
 - Evaluation
 - Models/approaches: rule-based, automata/grammars, perceptron, logistic regression, neural network models
 - Effective written and oral communication
- Components/Tasks/Subtasks:
 - Language Models
 - Syntax: POS tags, constituency tree, dependency tree, parsing
 - Semantics: lexical, formal, inference tasks
 - Information Extraction: Named Entities, Relations, Events
 - Generation: Machine Translation, Summarization, Dialogue, Creative Generation

¹<https://mitpress.mit.edu/books/introduction-natural-language-processing> or free version <https://github.com/jacobeisenstein/gt-nlp-class/blob/master/notes/eisenstein-nlp-notes.pdf>

²<https://mitpress.mit.edu/books/introduction-deep-learning> (first three chapters at <https://cs.brown.edu/courses/csci1460/assets/files/deep-learning.pdf>)

³<https://web.stanford.edu/~jurafsky/slp3/>

instructor	date	material	reading	Other
JM	8/26	intro, applications	Eisenstein 1 (not mandatory)	
JM	8/28	probability basics	Eisenstein Appendix A, Goldwater probability tutorial ⁴	
N/A	9/2	LABOR DAY NO CLASS		
JM	9/4	corpora, text processing, Linear Classifiers (Naive Bayes, Logistic Regression, Perceptron), part of speech tagging	Eisenstein 2, Nathan Schneider's sunix notes ⁵ , Unix for poets ⁶ , sculpting text ⁷	HW1 out (4 weeks)
TG	9/9	Nonlinear classifiers, feed forward neural networks, backpropagation, gradient descent	Eisenstein 3, Charniak 1	
TG	9/11	pytorch and google cloud basics		
JM	9/16	evaluation, annotation, mechanical turk		
JM	9/18	POS tags, HMMs, search	Eisenstein 7	
JM	9/23	parsing and syntax 1: treebanks, evaluation, cky, grammar induction, pcfgs	Eisenstein 9,2, 10	HW2 out (4 weeks)
JM	9/25	parsing and syntax 2: dependencies, shift-reduce, chiu-liu-edmonds	Eisenstein 11, A Fast and Accurate Dependency Parser using Neural Networks ⁸	
N/A	9/30	NO CLASS		
JM	10/2	language models, RNNs	Eisenstein 7	HW1 due
JM	10/7	semantics 1 (word sense disambiguation, prop-bank? AMR)		HW3 out (4 weeks)
XY	10/9	Text Games and Reinforcement Learning		
JM	10/14	distributional lexical semantics	Eisenstein 14	
JM	10/16	Machine Translation history, evaluation, statistical	Eisenstein 18.1, 18.2	project proposal due
JM	10/21	Neural Machine Translation, summarization, generation	Eisenstein 18.3, 19.1, 19.2	HW 2 Due
JM	10/23	Information Extraction: Entity/Relation, CRF	Eisenstein 17.1, 17.2	
JM	10/28	Information Extraction: Events, Zero-shot	Eisenstein 17.3	
JM	10/30	Transformers	Attention is all you need ⁹ , Illustrated Transformer ¹⁰	
JM	11/4	Large Contextualized Language Models (ELMo, BERT, GPT(2), etc.)	Illustrated BERT, ELMo, and co. ¹¹	HW3 due, HW4 out (4 weeks)
JM	11/6	Power and Ethics	Energy and Policy Considerations for Deep Learning in NLP ¹²	
JM	11/11	How to write a paper		
JM	11/13	Creative Generation, structure-to-text, text-to-text	Eisenstein 19.1, 19.2	
JM	11/18	Dialogue	Eisenstein 19.3	
JM	11/20	TBD		project paper due (if submitting to ACL)
JM	11/25	TBD		
N/A	11/27	THANKSGIVING NO CLASS		
Class	12/2	Presentations		HW4 due
Class	12/4	Presentations		

⁴<http://homepages.inf.ed.ac.uk/sgwater/teaching/general/probability.pdf>

⁵<https://github.com/nschneid/unix-text-commands>

⁶<https://www.cs.upc.edu/~padro/Unixforpoets.pdf>

⁷<http://matt.might.net/articles/sculpting-text/>

⁸<https://www.aclweb.org/anthology/D14-1082/>

⁹<https://arxiv.org/abs/1706.03762>

¹⁰<http://jalammar.github.io/illustrated-transformer/>

¹¹<http://jalammar.github.io/illustrated-bert/>

¹²<https://aclweb.org/anthology/papers/P/P19/P19-1355/>