

CSCI 662 Fall 2021 course page

Jonathan May

August 3, 2021

Website	https://www.isi.edu/~jonmay/cs662_fa21_web/
Lectures	Mondays and Wednesdays 3:30–5:20 pm, KAP 166
Instructor & office hours	Jonathan May, Mondays and Wednesdays 2:30–3:30 pm or by appointment, location TBD
TAs & Office hours	TBD
Textbook	Required: Natural Language Processing - Eisenstein ¹ Required: 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 (PRELIMINARY AND SUBJECT TO CHANGE)	10 %: In-class participation 10 %: Posted questions before each in-class selected paper presentation 10 %: In-class selected paper presentation 30 %: Three Homeworks (10% each) 40 %: Project, comprising proposal (10%), final conference-quality paper (15%), and 15-minute in-class presentation (15%) (may be done in small groups). Final report is due December 13, 2021, 4:00 PM PST
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

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/>

- Semantics: lexical, formal, inference tasks
- Information Extraction: Named Entities, Relations, Events
- Generation: Machine Translation, Summarization, Dialogue, Creative Generation

SCHEDULE SUBJECT TO CHANGE; READINGS CURRENTLY OUT OF ORDER

date	material	reading	presentation	Other
8/23	intro, applications	Eisenstein 1 (not mandatory)		project assignment out (due 9/1)
8/25	end of intro, probability basics	Eisenstein Appendix A, Goldwater probability tutorial ⁴		
8/30	ethics, naive bayes, text processing	Eisenstein 2.1, Nathan Schneider's unix notes ⁵ , Unix for poets ⁶ , sculpting text ⁷		
9/1	Perceptron, Logistic Regression	Eisenstein 2.2, 2.4, Charniak 1.		project proposal due
9/6	LABOR DAY NO CLASS			
9/8	Nonlinear classifiers, backpropagation, gradient descent	Eisenstein 3		HW1 out (due 9/29)
9/13	POS tags, HMMs	Eisenstein 7		
9/15	YOM KIPPUR NO CLASS			
9/20	constituencies, cky, treebank	Eisenstein 9.2, 10		
9/22	restructuring, dependencies, shift-reduce	Eisenstein 11		
9/27	arc-eager, evaluation, human annotation			
9/29	semantics: word sense, propbank, amr, distributional	Eisenstein 13, 14.		HW1 due
10/4	language models: ngram, feed-forward, recurrent	Eisenstein 6, 18.1, 18.2		
10/6	Machine Translation history, evaluation			
10/6	Statistical, Neural Machine Translation, summarization, generation	Eisenstein 18.3, 19.1, 19.2		HW2 out (due 10/25)
10/11	Transformers	Attention is all you need ⁸ , Illustrated Transformer ⁹		
10/13	Large Contextualized Language Models (ELMo, BERT, GPT-N, etc.)	Illustrated BERT, ELMo, and co. ¹⁰		
10/18	Guest Lecture TBD			
10/20	Guest Lecture TBD			

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://arxiv.org/abs/1706.03762>

⁹<http://jalammar.github.io/illustrated-transformer/>

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

10/25 10/27	Information Extraction: Entity/Relation, CRF Information Extraction: Events, Zero-shot	Eisenstein 17.1, 17.2 Eisenstein 17.3		HW2 due HW3 out (due 11/17)
11/1 11/3	Blade Runner NLP/Bertology Dialogue	Eisenstein 19.3		
11/8 11/10	Text Games and Reinforcement Learning Power and Ethics			
11/15 11/17	How to write a paper TBD	Neubig slides on Piazza		HW3 due
11/22 11/24 11/29 12/1	TBD THANKSGIVING BREAK; NO CLASS Presentations Presentations			