Course Syllabus

✎ Edit

Readings. The readings are assigned to complement the lectures. Readings can be consulted before lecture as preparation or after lecture to review.

Book 1: **[Goldberg]** Goldberg: A Primer on Neural Network Models for Natural Language Processing (https://u.cs.biu.ac.il/~yogo/nnlp.pdf) .

Book 2: **[Eisenstein]** Eisenstein: Natural Language Processing (https://canvas.ucsd.edu/courses/55037/files/12085505?wrap=1) (https://canvas.ucsd.edu/courses/55037/files/12085505/download?download_frd=1)

Book 3: **[ZLLS]** Aston Zhang, Zack C. Lipton, Mu Li, Alex J. Smola, Dive into Deep Learning (https://d2l.ai/index.html)

Book 4: **[J&M]** Jurafsky and Martin: Speech and Language Processing (3rd ed. draft) (https://web.stanford.edu/~jurafsky/slp3/)

Lecture # (No content. Not visible to students)	Title	Readings
Lecture 1: 2 April	Introduction (https://canvas.ucsd.edu/courses/55037/files/12081909?wrap=1) (https://canvas.ucsd.edu/courses/55037/files/12081909/download?download_frd=1)	Chris Manning: Computati Deep Learning (https://a
Lecture 2: 4 April	Text Classification with Linear Models (https://canvas.ucsd.edu/courses/55037/files/12105459?wrap=1) (https://canvas.ucsd.edu/courses/55037/files/12105459/download?download_frd=1)	• **[Eisenstein]** 2.1; 2.5.0 - • A Few Useful Things t (https://homes.cs.washir
Lecture 3: 9 April	Text Classification with FeedForward Neural Networks (https://canvas.ucsd.edu/courses/55037/files/12126326?wrap=1) (https://canvas.ucsd.edu/courses/55037/files/12126326/download?download_frd=1)	• **[Eisenstein]** 3.0-3.3 • Natural Language Pro (https://aclanthology.org
Lecture 4: 11 April	Word Embeddings (https://canvas.ucsd.edu/courses/55037/files/12136161?wrap=1) (https://canvas.ucsd.edu/courses/55037/files/12136161/download?download_frd=1)	Mikolov et al. 2013 word2 Pennington et al. 2014 Gl
Lecture 5: 16 April	The Language Modeling Problem (https://canvas.ucsd.edu/courses/55037/files/12177561?wrap=1) (https://canvas.ucsd.edu/courses/55037/files/12177561/download?download_frd=1)	Bengio et al. 2003 a Neura (https://www.jmlr.org/papers
Lecture 6: 18 April	Attention & Transformers (https://canvas.ucsd.edu/courses/55037/files/12217794?wrap=1) (https://canvas.ucsd.edu/courses/55037/files/12217794/download?download_frd=1)	Vaswani et al. 2017 (ht Let's build GPT: from scra (https://www.youtube.com/w Alammar Illustrated Trans Beltagy et al. 2020 Longfo
Lecture 7: 23 April (https://canvas.ucsd.edu/courses/55037/files/12251720?wrap=1)	Pretraining part 1: Encoders (BERT/ELECTRA/DeBERTa) (https://canvas.ucsd.edu/courses/55037/files/12251720?wrap=1)	Peters et al. 2018 ELMo Devlin et al 2019 BERT Alammar Illustrated BERT

(https://canvas.ucsd.edu/courses/55037/files/12251720/download?download_frd=1)	(https://canvas.ucsd.edu/courses/55037/files/12251720/download?download_frd=1)	Liu 2019 RoBERTa (http Clark et al 2020 ELECTRA He 2021 DeBERTa (http
Lecture 8: 25 April	Pretraining part 2: Decoders (GPT/T5/Llama); Decoding Methods; Scaling Laws (https://canvas.ucsd.edu/courses/55037/files/12273404?wrap=1) (https://canvas.ucsd.edu/courses/55037/files/12273404/download?download_frd=1)	Raffel et al 2019 T5 (htt Lewis 2019 BART (https Radford 2019 GPT2 (htt models/language_models_a Brown 2020 GPT3 (http Chowdhery 2021 PaLM Holtzman 2019 Nucleus S
Lecture 9: 30 April	Retrieval & RAG (external memory) (https://canvas.ucsd.edu/courses/55037/files/12316924?wrap=1) (https://canvas.ucsd.edu/courses/55037/files/12316924/download?download_frd=1)	Thoppilan et al. 2021 LaM Karpukhin et al.2020 DPR Borgeaud et al. 2022 RET
Lecture 10 2 May	Knowledge Representation in Transformer LLMs (internal memory) (https://canvas.ucsd.edu/courses/55037/files/12337328?wrap=1) (https://canvas.ucsd.edu/courses/55037/files/12337328/download?download_frd=1)	Geva et al. 2021 FeedForv main.446.pdf) Meng et al. 2022 ROME (https://proceedings.neurips Paper-Conference.pdf)
Lecture 11 May 7	Application #1: Question Answering (https://canvas.ucsd.edu/courses/55037/files/12377794?wrap=1) (https://canvas.ucsd.edu/courses/55037/files/12377794/download?download_frd=1)	Chen et al 2017 DrQA (Seo et al 2019 BiDAF (h
Lecture 12 May 9	Application #2: Code Generation (https://canvas.ucsd.edu/courses/55037/files/12397914?wrap=1) (https://canvas.ucsd.edu/courses/55037/files/12397914/download?download_frd=1)	Chen et al 2021 Codex Li et al 2022 AlphaCode Ahn et al 2022 SayCan
Lecture 13: May 14	Interpretability (https://canvas.ucsd.edu/courses/55037/files/12438186?wrap=1) (https://canvas.ucsd.edu/courses/55037/files/12438186/download?download_frd=1)	Conneau et al, 2018 (h Belinkov 2020 (https://a
Lecture 14: May 16	LLMs, and Society (https://canvas.ucsd.edu/courses/55037/files/12456002?wrap=1) (https://canvas.ucsd.edu/courses/55037/files/12456002/download?download_frd=1)	Jurgens et al 2019 (http Bender et al 2021, stocha
Lecture 15: May 21	Adaptive Pre-training and PEFT (Parameter Composition) (https://canvas.ucsd.edu/courses/55037/files/12493871?wrap=1) (https://canvas.ucsd.edu/courses/55037/files/12493871/download?download_frd=1)	Frankle & Cabin 2019, Lot Ansell et al 2022, Sparse-
Lecture 16: May 23	PEFT (Input Composition & Function Composition) (https://canvas.ucsd.edu/courses/55037/files/12522313?wrap=1) (https://canvas.ucsd.edu/courses/55037/files/12522313/download?download_frd=1)	Hu et al 2021, LoRA (ht Houlsby et al 2019, Adapt (https://proceedings.mlr.pres He et al 2022 (https://op
Lecture 17: May 28	Tokenization in LLMs (https://canvas.ucsd.edu/courses/55037/files/12555631?wrap=1)	