

Swabha Swayamdipta

GABILAN ASSISTANT PROFESSOR OF COMPUTER SCIENCE, UNIVERSITY OF SOUTHERN CALIFORNIA

✉ swabhas@usc.edu | 🏠 https://swabhs.com | 🌐 swabhs | 🐦 @swabhz

Appointments

University of Southern California

GABILAN ASSISTANT PROFESSOR OF COMPUTER SCIENCE

- Co-Associate Director of the Center for AI and Society

Los Angeles, CA, USA

Aug 2022 - PRESENT

Allen Institute for AI

POSTDOCTORAL INVESTIGATOR

- **Advisor:** Yejin Choi
- **Project Team:** MOSAIC

Seattle, WA, USA

Aug 2019 - June 2022

Oracle Server Technologies

MEMBER TECHNICAL STAFF

- **Project:** Unified messaging for Fusion Middleware

Bangalore, India

2010 - 2011

Education

Carnegie Mellon University

PHD IN LANGUAGE AND INFORMATION TECHNOLOGIES

- **Advisors:** Noah A. Smith, Chris Dyer
- **University of Washington**, Seattle, WA — *Visiting PhD student Fall 2015 - Summer 2019*
- **Thesis:** Syntactic Inductive Biases in NLP

Pittsburgh, PA, USA

2013 - 2019

Columbia University

MASTERS IN COMPUTER SCIENCE

- **Advisors:** Owen Rambow, Michael Collins
- **Specialization:** Natural Language Processing

City of New York, NY, USA

2011 - 2012

National Institute of Technology

BACHELORS OF TECHNOLOGY IN COMPUTER SCIENCE AND ENGINEERING

Calicut, Kerala, India

2006 - 2010

Awards and Grants

2024	Outstanding Paper Award , OATH-Frames: Characterizing Online Attitudes Towards Homelessness with LLM Assistants	EMNLP
2024	Simons Institute Research Fellowship , Special Program on LLMs and Transformers, Spring 2025	UC Berkeley
2024	NSF CISE Core Medium , RII: Hard Data to the Model	NSF
2024	Zumberge Fellowship , DEI in Research	USC OORI
2023	Intel Rising Stars Award , LLM Inference for Scaling Data Creation	Intel Labs
2022	AI2 Young Investigators Award , Synthesizing Data Towards Commonsense and Causal Reasoning	Allen Institute for AI
2022	Gabilan Assistant Professor Fellowship ,	USC WISE
2022	Outstanding Paper Award , Understanding Dataset Difficulty with V-Usable Information	ICML
2021	Outstanding Paper Award , MAUVE: Measuring the Gap Between Neural Text and Human Text using Divergence Frontiers	NeurIPS
2020	Best Paper: Honorable Mention , Don't Stop Pretraining: Adapt Language Models to Domains and Tasks	ACL
2019	Rising Star in EECS , UIUC	
2014	Best Student Presentation , Diversity in Dependency Parsing	CMU LTI Student Research Symposium
2009	Sun Campus Ambassador Scholarship , Bangalore, India	Sun Microsystems
2003	National Talent Search Scholarship , Rourkela, India	NCERT, India

Publications

CONFERENCE PAPERS

- Compare without Despair: Reliable Preference Evaluation with Generation Separability *EMNLP 2024*
• S. Ghosh, T. Srinivasan, **S. Swayamdipta**,
- Out-of-Distribution Detection through Soft Clustering with Non-Negative Kernel Regression *EMNLP 2024*
• A. Gulati, X. Dong, C. Hurtado, S. Shekkizhar, **S. Swayamdipta**, A. Ortega
- OATH-Frames: Characterizing Online Attitudes towards Homelessness via LLM Assistants *EMNLP 2024*
• J. Ranjit, B. Joshi, R. Dorn, L. Petry, O. Koumoundouros, J. Bottarini, P. Liu, E. Rice, and **S. Swayamdipta**
• **Outstanding Paper Award @ EMNLP 2024 & Best Poster @ ShowCAIS 2024**
- Logits of API-Protected LLMs Leak Proprietary Information *COLM 2024*
• M. Finlayson, X. Ren, **S. Swayamdipta**,
- Crowd-Calibrator: Can Annotator Disagreement Inform Calibration in Subjective Tasks? *COLM 2024*
• U. Khurana, E. Nalisnick, A. Fokkens, **S. Swayamdipta**,
- Annotating FrameNet via Structure-Conditioned Language Generation *ACL 2024*
• X. Cui, and **S. Swayamdipta**
- Closing the Curious Case of Neural Text Degeneration *ICLR 2024*
• M. Finlayson, J. Hewitt, A. Koller, **S. Swayamdipta**, A. Sabharwal
- Does Video Summarization Require Videos? Quantifying the Effectiveness of Language in Video Summarization *ICASSP 2024*
• Y. Nam, A. Lehavi, D. Yang, D. Bose, **S. Swayamdipta**, and S. Narayanan
- NeuroComparatives: Neuro-Symbolic Distillation of Comparative Knowledge *NAACL-Findings 2024*
• P. Howard, J. Wang, V. Lal, G. Singer, Y. Choi, and **S. Swayamdipta**
- Generative Explanations for Program Synthesizers *ICSE + VL/HCC 2024*
• A. Nazari, S. Chattopadhyay, **S. Swayamdipta**, M. Raghothaman
- We're Afraid Language Models Aren't Modeling Ambiguity *EMNLP 2023*
• A. Liu, Z. Wu, J. Michael, A. Suhr, P. West, A. Koller, **S. Swayamdipta**, N. A. Smith, and Y. Choi
- MAUVE Scores for Generative Models: Theory and Practice *JMLR 2023*
• K. Pillutla, L. Liu, J. Thickstun, S. Welleck, **S. Swayamdipta**, R. Zellers, S. Oh, Y. Choi, and Z. Harchaoui
- I2D2: Inductive Knowledge Distillation with NeuroLogic and Self-Imitation *ACL 2023*
• C. Bhagavatula, J. D. Hwang, D. Downey, R. Le Bras, X. Lu, K. Sakaguchi, **S. Swayamdipta**, P. West, and Y. Choi
- REV: Information-Theoretic Evaluation of Free-Text Rationales *ACL 2023*
• H. Chen, F. Brahman, X. Ren, Y. Ji, Y. Choi, and **S. Swayamdipta**
- COBRA Frames: Contextual Reasoning about Effects and Harms of Offensive Statements *ACL-Findings 2023*
• X. Zhou, H. Zhu, A. Yerukola, T. Davidson, J. D. Hwang, **S. Swayamdipta**, and M. Sap
- Investigating the Benefits of Free-Form Rationales *EMNLP-Findings 2022*
J. Sun, **S. Swayamdipta**, J. May, and X. Ma
- NeuroCounterfactuals: Beyond Minimal-Edit Counterfactuals for Richer Data Augmentation *EMNLP-Findings 2022*
P. Howard, G. Singer, V. Lal, Y. Choi, and **S. Swayamdipta**
- WaNLI: Worker and AI Collaboration for Natural Language Inference Dataset Creation *EMNLP-Findings 2022*
A. Liu, **S. Swayamdipta**, N. A. Smith, and Y. Choi
- Reframing Human-AI Collaboration for Generating Free-Text Explanations *NAACL 2022*
S. Wiegreffe, J. Hessel, **S. Swayamdipta**, M. Riedl, and Y. Choi
- Annotators with Attitudes: How Annotator Beliefs And Identities Bias Toxic Language Detection *NAACL 2022*
M. Sap, **S. Swayamdipta**, L. Vianna, X. Zhou, Y. Choi, and N. A. Smith
- Understanding Dataset Difficulty with V-Usable Information *ICML 2022*
• K. Ethayarajh, Y. Choi, and **S. Swayamdipta**
• **Outstanding Paper Award**

MAUVE: Measuring the Gap Between Neural Text and Human Text using Divergence Frontiers *NeurIPS 2021*

- K. Pillutla, **S. Swayamdipta**, R. Zellers, J. Thickstun, S. Wellecks, Y. Choi, and Z. Harchaoui
- **Outstanding Paper Award**

Contrastive Explanations for Model Interpretability *EMNLP 2021*

A. Jacovi, **S. Swayamdipta**, S. Ravfogel, Y. Elazar, Y. Choi, and Y. Goldberg

On-the-Fly Controlled Text Generation with Experts and Anti-Experts *ACL 2021*

A. Liu, M. Sap, X. Lu, **S. Swayamdipta**, C. Bhagavatula, N. A. Smith, and Y. Choi

Challenges in Automated Debiasing for Toxic Language Detection *EACL 2021*

- X. Zhou, M. Sap, **S. Swayamdipta**, N. A. Smith, and Y. Choi

Dataset Cartography: Mapping and Diagnosing Datasets with Training Dynamics *EMNLP 2020*

- **S. Swayamdipta**, R. Schwartz, N. Lourie, Y. Wang, H. Hajishirzi, N. A. Smith, Y. Choi

Generative Data Augmentation for Commonsense Reasoning *EMNLP-Findings 2020*

- Y. Yang, C. Malaviya, J. Fernandez, **S. Swayamdipta**, R. LeBras, J. Wang, C. Bhagavatula, Y. Choi, and D. Downey

Adversarial Filters of Dataset Biases *ICML 2020*

- R. LeBras, **S. Swayamdipta**, C. Bhagavatula, R. Zellers, M. E. Peters, A. Sabharwal, and Y. Choi

The Right Tool for the Job: Matching Model and Instance Complexities *ACL 2020*

- R. Schwartz, G. Stanovsky, **S. Swayamdipta**, J. Dodge, and N. A. Smith

Don't Stop Pretraining: Adapt Language Models to Domains and Tasks *ACL 2020*

- S. Gururangan, A. Marasović, **S. Swayamdipta**, K. Lo, I. Beltagy, D. Downey, and N. A. Smith
- **Best Paper Honorable Mention**

Syntactic Scaffolds for Semantic Structures *EMNLP 2018*

- **S. Swayamdipta**, S. Thomson, K. Lee, L. Zettlemoyer, C. Dyer, N. A. Smith.

Learning Joint Semantic Parsers from Disjoint Data *NAACL 2018*

- H. Peng, S. Thomson, **S. Swayamdipta**, N. A. Smith

Annotation Artifacts in Natural Language Inference Data *NAACL 2018*

- S. Gururangan*, **S. Swayamdipta***, O. Levy, R. Schwartz, S. Bowman, and N. A. Smith
- * equal contribution

Polyglot Semantic Role Labeling *ACL 2018*

- P. Mulcaire, **S. Swayamdipta**, and N. A. Smith

Multi-Mention Learning for Reading Comprehension with Neural Cascades *ICLR 2018*

- **S. Swayamdipta**, A. Parikh, T. Kwiatkowski

Greedy, Joint Syntactic and Semantic Parsing with Stack LSTMs *CoNLL 2016*

- **S. Swayamdipta**, M. Ballesteros, C. Dyer, N. A. Smith

A Dependency Parser for Tweets *EMNLP 2014*

- L. Kong, N. Schneider, **S. Swayamdipta**, A. Bhatia, C. Dyer, N. A. Smith

The Pursuit of Power and its Manifestation in Written Dialog *ICSC 2012*

- **S. Swayamdipta**, O. Rambow

WORKSHOP PAPERS

Sister Help: Data Augmentation for Frame-Semantic Role Labeling *LAW-DMR @ EMNLP 2021*

- A. Pancholy, **S. Swayamdipta**, M. R. L. Petruck

Multi-Task Learning for Incremental Parsing using Stack LSTMs *WiML @ NeurIPS 2016*

- **S. Swayamdipta**, M. Ballesteros, C. Dyer, N. A. Smith

CMU: Arc-Factored, Discriminative Semantic Dependency Parsing *SemEval 2014*

- S. Thomson, D. Bamman, J. Dodge, **S. Swayamdipta**, N. Schneider, C. Dyer, N. A. Smith

The CMU Machine Translation Systems *WMT 2014*

- A. Matthews, C. Dyer, A. Lavie, G. Hanneman, W. Ammar, A. Bhatia, **S. Swayamdipta**, E. Schlinger, Y. Tsvetkov

Invited Talks

Nov 07, 2024	NYU NLP/Text-as-Data Talk Series , Ensuring Safety and Accountability in LLMs, Pre- and Post-Training	<i>New York University</i>
Sep 23, 2024	Amazon AWS Bedrock AI Seminar , Ensuring Safety and Accountability in LLMs, Pre- and Post-Training	<i>Amazon AWS, Virtual</i>
May 3, 2024	ISI AI Seminar , Understanding LLMs through their Generative Behavior, Successes and Shortcomings	<i>USC ISI</i>
Mar 26, 2024	NSF-Open Source Generative AI Workshop , Towards (Closed-Source) LLM Accountability via Logit Signatures	<i>Cornell Tech</i>
Mar 13, 2024	Utah Data Science Seminar , Understanding LLMs through their Generative Behavior, Successes and Shortcomings	<i>University of Utah</i>
Feb 13, 2024	Intel Rising Stars Talk , LLM Inference for Scaling Data Creation	<i>Intel Labs</i>
Feb 8, 2024	Cambridge LTL Seminar , Understanding LLMs via their Generative Successes and Shortcomings	<i>University of Cambridge</i>
Dec 15, 2023	ATTRIB Workshop Keynote , Understanding LLMs via their Generative Successes and Shortcomings	<i>NeurIPS 2023</i>
Nov 16, 2023	UCLA Communications Symposium , Understanding Online Discourse through Social Context and Structured Pragmatics	<i>UCLA</i>
Oct 3, 2023	DataComp Workshop , Understanding Data with \mathcal{V} -Information	<i>ICCV 2023</i>
Sep 22, 2023	LACoE CS Speaker Series , The Role of Language Models in Natural Language Processing	<i>Los Angeles County of Education</i>
Aug 24, 2023	CMU LTI Student Research Symposium , Understanding Datasets and Explanations through \mathcal{V} -Information	<i>Alumni Keynote: CMU, Pittsburgh</i>
Jul 13, 2023	Rep4NLP Workshop , Contextualizing Representations in Varied Annotator Perspectives	<i>ACL Toronto</i>
Jun 14, 2023	Google Responsible Machine Learning , Contextualizing Data in Varied Annotator Perspectives	<i>Google Brain</i>
Mar 14, 2023	Cohere for AI Reading Group , Understanding Dataset Difficulty with V-Usable Information	<i>Cohere for AI</i>
Feb 28, 2023	Spotify Research Seminar , Designing Controls and Filters for Dataset Generation	<i>Spotify Research Labs</i>
Feb 03, 2023	Amazon Data-Centric AI Seminar , What's in a Dataset? Interpreting datasets to enable better data creation	<i>Amazon</i>
Feb 01, 2023	USC CAIS++ Seminar , Contextualizing Bias in Hate Speech Detection	<i>USC</i>
Nov 18, 2022	SoCal NLP Symposium , Generating Datasets for Robust Generalization	<i>UC Santa Barbara</i>
Nov 09, 2022	USC Center for AI in Society (CAIS) Seminar , Contextualizing Bias in Hate Speech Detection through Annotator Perspectives	<i>USC</i>
May 23, 2022	ACL Spotlight Talks for Young Rising Stars , The Devil's in the Data: Mapping and Generating Datasets for Robust Generalization	<i>Dublin, Ireland</i>
May 13, 2022	UC Irvine CS Seminar , Mapping and Generating Datasets for Robust Generalization	<i>UC Irvine</i>
Mar 16, 2022	UC Santa Cruz IFDS Ethics Seminar , Rethinking Dataset Construction: : The Role of Generative Modeling and Annotator Perspectives	<i>UC Santa Cruz</i>
Oct 20, 2021	Oracle ML Talks , What's in your Data? Mapping Datasets and Exploring Data Usability	<i>Oracle</i>
Feb 24, 2020	Georgetown NERT Seminar , Addressing Biases for Robust, Generalizable AI	<i>Georgetown University</i>
Feb 12, 2020	Georgia Tech NLP Seminar , Addressing Biases for Robust, Generalizable AI	<i>Georgia Tech</i>
Nov 02, 2020	Microsoft E+D Product Leads , Responsible AI: Addressing Biases in Datasets and Models	<i>Microsoft</i>
Oct 30, 2019	Rising Stars in EECs , Biases and Learning Challenges in Natural Language Processing	<i>UIUC</i>
Oct 18, 2019	UW Linguistics Colloquium , Sprucing up a Dataset: Adversarially Filtering Dataset Artifacts	<i>UW Linguistics</i>
Apr 08, 2019	AI2 Seminar , Learning Challenges in Natural Language Processing	<i>Allen Institute for AI</i>
Apr 04, 2019	Stanford NLP Seminar , Learning Challenges in Natural Language Processing	<i>Stanford University</i>
Mar 13, 2019	JHU CLSP Seminar , Learning Challenges in Natural Language Processing	<i>Johns Hopkins University</i>
Mar 11, 2019	NYU CILVR Seminar , Learning Challenges in Natural Language Processing	<i>New York University</i>
Jan 22, 2019	MSR Seminar , Learning Challenges in Natural Language Processing	<i>Microsoft Research, New York</i>
Sep 15, 2018	MSR AI Breakthroughs Workshop , Representation Learning with Linguistic Structure	<i>Microsoft Research</i>
Apr 27, 2018	NorthWest NLP , Syntactic Scaffolds for Semantic Structures	<i>Microsoft Research</i>

Mentorship

PostDocs

Fall 2024-Present **Gregory Yauney**, PostDoc

PHD STUDENTS

Fall 2024-Present **Xinyue Cui**, PhD

Fall 2024-Present **Muru Zhang**, PhD

Co-Advisor: Robin Jia

Fall 2024-Present **Atharva Kulkarni**, PhD

Spring 2024-Present **Brihi Joshi**, PhD, passed Quads

Co-Advisor: Xiang Ren

Fall 2023-Present **Matthew Finlayson**, PhD

Co-Advisor: Xiang Ren

Fall 2022-Present **Sayan Ghosh**, PhD

Fall 2022-Present **Jaspreet Ranjit**, PhD, passed Quads

MASTERS STUDENTS

Spring 2024-Present **Shahzaib Saqib Warraich**, Masters

Spring 2024-Present **Shauryasikt Jena**, Masters

Spring 2024-Present **Xingjian Dong**, Masters

UNDERGRADS

Spring 2024-Present **Ryan Wang**, BS

Co-Advisor: Robin Jia

Spring 2024-Present **Joseph Liu**, BS

Spring 2024-Present **Keyu He**, BS

Fall 2023-Present **Risha Surana**, BS

Summer 2023-Present **Aryan Gulati**, BS

Co-Advisor: Antonio Ortega

Summer 2023-Present **Catherine He**, BS

ALUMNI

Spring 2023-Spring 2024 **Xinyue Cui**, Masters

Now PhD@USC

Fall 2023-Spring 2024 **Jacky Mo**, BS

Summer 2023-Fall 2023 **Jayne Bottarini**, BS

Summer 2023-Fall 2023 **Kritin Dhoka**, BS

Spring 2023-Spring 2024 **Yoonsoo Nam**, Masters

Now at RenderWolf

Spring 2023-Summer 2023 **Adam Lehavi**, BS

Summer 2023 **Ruyuan Zuo**, Masters

Now at Google

Spring 2022 **Hanjie Chen / Junlin Wang / Abdallah Bashir**, PhD, UVA / Masters, UCI / Masters, Saarland Universit

AI2 Intern

Fall 2021 **Jillian Fisher / Liwei Jiang**, PhD, UW

Summer 2021 **Kawin Ethayarajh**, PhD, Stanford University

AI2 Intern

Summer 2021 **Sarah Wiegrefe**, PhD, George Institute of Technology

AI2 Intern

Summer 2021 **Ximing Liu**, BS, UW CSE

AI2 Intern

Fall 2020-2023 **Alisa Liu**, PhD, UW CSE

Fall 2020 **Alon Jacovi**, PhD, Bar Ilan University

AI2 Intern

Summer 2020-1 **Jenny Liang**, BS, UW CSE

AI2 Intern

Summer 2020 **Jize Cao**, BS, UW CSE

AI2 Intern

Spring 2020-2023 **Xuhui Zhou**, CLMS, UW Linguistics

Fall 2019 - Spring 2020 **Yiben Yang**, PhD, Northwestern University

Fall 2019 - Spring 2020 **Chaitanya Malaviya**, PYI, Allen Institute for AI

2018 - 2019 **Ron Fan / Karishma Mandyam**, MS, UW CSE / BS, UW CSE

Invited Panels

Oct 18, 2024 **Alfred P. Sloan Foundation Science Seminar**,
 Jul 13, 2023 **Limitations of Large Language Models**, Rep4NLP Workshop
 Jun 29, 2023 **Generative AI and Education**, Roundtable
 Apr 5, 2023 **Almaginings**, Polymathic Pizza
 Jul 14, 2022 **Adversarial Data Augmentations**, DADC Workshop
 Jun 17, 2022 **Role of LLMs**, Responsible AI Symposium

USC School of Cinematic Arts
ACL Toronto
ASEE 2023
USC Sidney Harmon Academy
for Polymathic Studies
NAACL Seattle
AILA

Professional Service

EXTERNAL SERVICE

Senior Area Chair

- EMNLP 2024: Machine Learning in NLP
- ACL 2024: Sentence-level Semantics

Area Chair

- ICLR 2025
- COLM 2024
- ICLR 2024
- EMNLP 2023: Machine Learning for NLP
- ACL 2023: Interpretability
- EMNLP 2022: Language Models
- EMNLP 2021: Machine Learning for NLP
- NAACL 2021: Sentence-Level Semantics
- EACL 2021: Sentence-Level Semantics
- ACL 2020: Semantics (Long)

Reviewer: Conferences

- ACL (ARR) 2015-2022
- NAACL 2015-2022
- EMNLP 2015-2022
- NeurIPS 2018-2022
- ICML 2015, 2019, 2020, 2023
- EACL 2017
- AAAI 2017-2020
- CoNLL 2017-2018, 2020

Reviewer: Workshops

- NAACL Student Research Workshop 2016
- Workshop for Women in Machine Learning (WiML) at NeurIPS 2016
- Machine Reading for Question Answering (MRQA) Workshop 2017

Reviewer: Journals

- Transactions of ACL 2020 - PRESENT
- Computational Linguistics 2019 - 2022
- Journal of AI Research 2019

Co-organizer

- EACL 2024: Workshop on Uncertainty in NLP 2024
- NAACL 2022: Workshop on Deep Learning in Low Resource NLP (DeepLo) 2022
- EMNLP 2019: Workshop on Deep Learning in Low Resource NLP (DeepLo) 2019
- West Coast NLP Workshop 2018
- Women's Research Day 2017
- UW-NLP Retreat 2015-2016

Mentorship at Conferences

- EMNLP 2020, ACL 2020, EMNLP 2018

INTERNAL SERVICE

Univeristy of Southern California

• Diversity, Equity and Inclusion Committee	2023-present
• Apple and IBM Fellowship Committee	2024
• Distinguished Lectures Committee	2023-2024
• Annual Faculty Review Committee	2023-2024
• Apple and IBM Fellowship Committee	2023
• PhD Fellowship Admissions Committee	2022-2023
• Viterbi Undergraduate Admission Scholarship	2023-2024

Los Angeles, CA, USA

Allen Institute for AI

• Diversity, Equity and Inclusion Committee	2019-2022
• AllenAI Outstanding Engineer Scholarship Committee	2019

Seattle WA, USA

Paul G. Allen School of CSE, University of Washington

• Faculty Recruiting Liaison	2019
• Graduate Student Advisory Council	2018
• Outreach: Highline Public High School	2017
• CS: Gradswomen	2016-2019

Seattle WA, USA

Teaching

CLASSES

Fall 2024	CSCI 544: Applied NLP , Students: 260	USC
Spring 2024	CSCI 499: Language Models in NLP , Students: 28	USC
Fall 2023	CSCI 499: Language Models in NLP , Students: 30	USC
Fall 2022	CSCI 699: Data-Centric NLP , Students: 29	USC

TUTORIALS

2022	OAMLS: Online Asian Machine Learning School , Attendees: 350+	ACML
2019	Transfer Learning in Natural Language Processing , Attendees: 350+	NAACL
2018	Frame Semantics across Languages: Towards a Multilingual FrameNet , Attendees: 30+	CoLING

TEACHING ASSISTANTSHIP

Spring 2015	CMU 11-411/611: Natural Language Processing , Instructors: Chris Dyer, Alan W Black, Shomir Wilson	Head TA
-------------	---	----------------

GUEST LECTURES

Summer 2024	Introduction to (Large) Language Models , Host: Darin Gray	<i>USC Viterbi K-12 Discover Engineering</i>
Spring 2024	Language Generation and Generative Evaluation , Host: Xuezhe Ma	<i>USC CSCI 544: Advanced NLP</i>
Summer 2023	Introduction to NLP , Host: Darin Gray	<i>USC Viterbi K-12 Discover Engineering</i>
Spring 2023	Contextualizing Bias , Host: Maja Matarić	<i>USC CSCI 697: Seminar in CS Research</i>
Spring 2022	What's in a dataset? , Host: Cyrus Shahabi	<i>USC CSCI 697: Seminar in CS Research</i>
Spring 2022	Data-Centric NLP , Host: Dongyeop Kang	<i>UMinn. CSCI 8980: Intro / NLP Research</i>
Spring 2022	Annotating Ambiguous NLI , Host: Dirk Hovy	<i>MilaNLP: Coding Aperitivo</i>
Spring 2021	Transfer Learning , Host: Zaid Harachoui	<i>UW DATA 598: Statistical Deep Learning</i>
Fall 2020	Biases and Interpretability in NLP , Host: Eunsol Choi	<i>UT-Austin CS 395T: Topics in NLP</i>
Winter 2019	Phrase-Structure Parsing , Host: Noah A. Smith	<i>UW CSE 447/547M: NLP</i>
Spring 2018	Minimum Bayes Risk Decoding , Host: Waleed Ammar	<i>UW CSE 599 D1: Advanced Topics in NLP</i>
Spring 2018	Dependency Parsing , Host: Noah A. Smith	<i>UW CSEP 517: NLP</i>
Fall 2017	Unsupervised Learning , Host: Noah A. Smith	<i>UW CSE 446: Machine Learning</i>
Fall 2014	Dependency Parsing with Chu-Liu-Edmonds , Hosts: Chris Dyer, Alon Lavie, Bob Frederking	<i>CMU 11-271: Algorithms for NLP</i>