

Robin Jia

Email: robinjia@usc.edu Address: 1031 Downey Way, Los Angeles, CA 90089
Website: <https://robinjia.github.io/> Office: GCS 405E Phone: (213) 740-5906

Education	Stanford University September 2014 – August 2020 Ph.D. in Computer Science Advisor: Percy Liang Thesis: Building Robust Natural Language Processing Systems
	Stanford University September 2010 – June 2014 Bachelor of Science with Honors in Computer Science Minor in Biology GPA: 4.103 / 4.0
Employment	Assistant Professor , Computer Science Department, University of Southern California Los Angeles, CA August 2021 – Present
	Visiting Researcher , Facebook AI Research Menlo Park, CA August 2020 – August 2021 Hosts: Luke Zettlemoyer and Douwe Kiela
	Research Intern , Microsoft Research Redmond, WA June 2018 – September 2018 Host: Hoifung Poon
	Research Intern , Google Research Mountain View, CA June 2016 – September 2016 Host: Larry Heck
Awards	Google ML and Systems Junior Faculty Award 2025
	Google Research Scholar Award 2023 – 2024
	Cisco Research Award 2023 – 2024
	Open Philanthropy research grant 2021 – 2024
	Two-time SoCalNLP Symposium Best Paper Award 2022, 2023
	Best Short Paper ACL 2018
	Outstanding Paper (Best paper runner-up) EMNLP 2017
National Science Foundation Graduate Research Fellow 2014 – 2019	
Publications	Promote, Suppress, Iterate: How Language Models Answer One-to-Many Factual Queries Tianyi Lorena Yan and Robin Jia EMNLP 2025
	Rethinking Backdoor Detection Evaluation for Language Models Jun Yan, Wenjie Jacky Mo, Xiang Ren, and Robin Jia EMNLP 2025
	Why Do Some Inputs Break Low-Bit LLM Quantization? Ting-Yun Chang, Muru Zhang, Jesse Thomason, and Robin Jia EMNLP 2025
	TokenSmith: Streamlining Data Editing, Search, and Inspection for Large-Scale Language Model Training and Interpretability Mohammad Aflah Khan*, Ameya Godbole*, Johnny Tian-Zheng Wei, EMNLP Demo 2025 Ryan Wang, James Flemings, Krishna P. Gummadi, Willie Neiswanger, and Robin Jia
	Teaching Models to Understand (but not Generate) High-risk Data Ryan Wang, Matthew Finlayson, Luca Soldaini, Swabha Swayamdipta, COLM 2025 and Robin Jia

- LLM Unlearning Without an Expert Curated Dataset**
Xiaoyuan Zhu, Muru Zhang, Ollie Liu, Robin Jia, and Willie Neiswanger COLM 2025
- Verify with Caution: The Pitfalls of Relying on Imperfect Factuality Metrics**
Ameya Godbole and Robin Jia ACL Findings 2025
- Mechanistic Interpretability of Emotion Inference in Large Language Models**
Ala N. Tak*, Amin Banayeeanzade*, Anahita Bolourani, Mina Kian, Robin Jia, and Jonathan Gratch ACL Findings 2025
- Robust Data Watermarking in Language Models by Injecting Fictitious Knowledge**
Xinyue Cui, Johnny Tian-Zheng Wei, Swabha Swayamdipta, and Robin Jia ACL Findings 2025
- Interrogating LLM Design under a Fair Learning Doctrine**
Johnny Tian-Zheng Wei*, Maggie Wang*, Ameya Godbole, Jonathan H. Choi, and Robin Jia FAccT 2025
- Language Models can Infer Action Semantics for Classical Planners from Environment Feedback**
Wang Zhu, Ishika Singh, Robin Jia, and Jesse Thomason NAACL 2025
- TLDR: Token-Level Detective Reward Model for Large Vision Language Models**
Deqing Fu, Tong Xiao, Rui Wang, Wang Zhu, Pengchuan Zhang, Guan Pang, Robin Jia, and Lawrence Chen ICLR 2025
- Pre-trained Large Language Models Use Fourier Features to Compute Addition**
Tianyi Zhou, Deqing Fu, Vatsal Sharan, and Robin Jia NeurIPS 2024
- Transformers Learn Higher-Order Optimization Methods for In-Context Learning: A Study with Linear Models**
Deqing Fu, Tian-Qi Chen, Robin Jia, and Vatsal Sharan NeurIPS 2024
SoCalNLP Symposium 2023 Best Paper Award.
- When Parts are Greater Than Sums: Individual LLM Components Can Outperform Full Models**
Ting-Yun Chang, Jesse Thomason, and Robin Jia EMNLP 2024
- Operationalizing content moderation "accuracy" in the Digital Services Act**
Johnny Tian-Zheng Wei, Frederike Zufall, and Robin Jia AIES 2024
- IsoBench: Benchmarking Multimodal Foundation Models on Isomorphic Representations**
Deqing Fu*, Ghazal Khalighinejad*, Ollie Liu*, Bhuwan Dhingra, Dani Yogatama, Robin Jia, and Willie Neiswanger COLM 2024
- Proving membership in LLM pretraining data via data watermarks**
Johnny Tian-Zheng Wei*, Ryan Yixiang Wang*, and Robin Jia ACL Findings 2024
- Do Localization Methods Actually Localize Memorized Data in LLMs?**
Ting-Yun Chang, Jesse Thomason, and Robin Jia NAACL 2024
- Efficient End-to-End Visual Document Understanding with Rationale Distillation**
Wang Zhu, Alekh Agarwal, Mandar Joshi, Robin Jia, Jesse Thomason, and Kristina Toutanova NAACL 2024
- Chain-of-Questions Training with Latent Answers for Robust Multistep Question Answering**

- Wang Zhu, Jesse Thomason, and Robin Jia EMNLP 2023
- SCENE: Self-Labelled Counterfactuals for Extrapolating to Negative Examples**
Deqing Fu, Ameya Godbole, and Robin Jia EMNLP 2023
- Estimating Large Language Model Capabilities without Labeled Test Data**
Harvey Yiyun Fu, Qinyuan Ye, Albert Xu, Xiang Ren, and Robin Jia EMNLP Findings 2023
- How Predictable Are Large Language Model Capabilities?
A Case Study on BIG-bench**
Qinyuan Ye, Harvey Yiyun Fu, Xiang Ren, and Robin Jia EMNLP Findings 2023
- Data Curation Alone Can Stabilize In-context Learning**
Ting-Yun Chang and Robin Jia ACL 2023
- Contrastive Novelty-Augmented Learning: Anticipating Outliers with
Large Language Models**
Albert Xu, Xiang Ren, and Robin Jia ACL 2023
SoCalNLP Symposium 2022 Best Paper Award.
- Are Sample-Efficient NLP Models More Robust?**
Nelson F. Liu, Ananya Kumar, Percy Liang, and Robin Jia ACL 2023
- Do Question Answering Modeling Improvements Hold Across Benchmarks?**
Nelson F. Liu, Tony Lee, Robin Jia, and Percy Liang ACL 2023
- Does VLN Pretraining Work with Nonsensical or Irrelevant Instructions?**
Wang Zhu, Ishika Singh, Yuan Huang, Robin Jia, and Jesse Thomason O-DRUM 2023
- Benchmarking Long-tail Generalization with Likelihood Splits**
Ameya Godbole and Robin Jia EACL Findings 2023
- Generalization Differences between End-to-End and Neuro-Symbolic
Vision-Language Reasoning Systems**
Wang Zhu, Jesse Thomason, and Robin Jia EMNLP Findings 2022
- Knowledge base question answering by case-based reasoning over subgraphs**
Rajarshi Das, Ameya Godbole, Ankita Naik, Elliot Tower, Manzil Zaheer, ICML 2022
Hannaneh Hajishirzi, Robin Jia, and Andrew McCallum
- On the Robustness of Reading Comprehension Models to Entity Renaming**
Jun Yan, Yang Xiao, Sagnik Mukherjee, Bill Yuchen Lin, Robin Jia, NAACL 2022
and Xiang Ren
- Models in the Loop: Aiding Crowdworkers with Generative Annotation Assistants**
Max Bartolo, Tristan Thrush, Sebastian Riedel, Pontus Stenetorp, Robin Jia, NAACL 2022
and Douwe Kiela
- Question Answering Infused Pre-training of
General-Purpose Contextualized Representations**
Robin Jia, Mike Lewis, and Luke Zettlemoyer ACL Findings 2022
- Analyzing Dynamic Adversarial Training Data in the Limit**
Eric Wallace, Adina Williams, Robin Jia, and Douwe Kiela ACL Findings 2022
- On Continual Model Refinement in Out-of-Distribution Data Streams**
Bill Yuchen Lin, Sida Wang, Xi Victoria Lin, Robin Jia, Lin Xiao, Xiang Ren, ACL 2022

and Scott Yih

Dynaboard: An Evaluation-As-A-Service Platform for Holistic Next-Generation Benchmarking

Zhiyi Ma*, Kawin Ethayarajh*, Tristan Thrush*, Somya Jain, Ledell Wu, Robin Jia, Christopher Potts, Adina Williams, and Douwe Kiela NeurIPS 2021

Masked Language Modeling and the Distributional Hypothesis: Order Word Matters Pre-training for Little

Koustuv Sinha, Robin Jia, Dieuwke Hupkes, Joelle Pineau, Adina Williams, and Douwe Kiela EMNLP 2021

Improving Question Answering Model Robustness with Synthetic Adversarial Data Generation

Max Bartolo, Tristan Thrush, Robin Jia, Sebastian Riedel, Pontus Stenetorp, and Douwe Kiela EMNLP 2021

To What Extent do Human Explanations of Model Behavior Align with Actual Model Behavior?

Grusha Prasad, Yixin Nie, Mohit Bansal, Robin Jia, Douwe Kiela, and Adina Williams BlackBoxNLP 2021

The statistical advantage of automatic NLG metrics at the system level

Johnny Tian-Zheng Wei and Robin Jia ACL 2021

Evaluation Examples Are Not Equally Informative: How Should That Change NLP Leaderboards?

Pedro Rodriguez, Joe Barrow, Alexander Hoyle, John P. Lalor, Robin Jia, and Jordan Boyd-Graber ACL 2021

Do Explanations Help Users Detect Errors in Open-Domain QA? An Evaluation of Spoken vs. Visual Explanations

Ana Valeria Gonzalez, Gagan Bansal, Angela Fan, Yashar Mehdad, Robin Jia, and Srinivasan Iyer ACL Findings 2021

Swords: A Benchmark for Lexical Substitution with Improved Data Coverage and Quality

Mina Lee*, Chris Donahue*, Robin Jia, Alexander Iyabor, and Percy Liang NAACL 2021

Dynabench: Rethinking Benchmarking in NLP

Douwe Kiela, Max Bartolo, Yixin Nie, Divyansh Kaushik, Atticus Geiger, Zhengxuan Wu, Bertie Vidgen, Grusha Prasad, Amanpreet Singh, Pratik Ringshia, Zhiyi Ma, Tristan Thrush, Sebastian Riedel, Zeerak Waseem, Pontus Stenetorp, Robin Jia, Mohit Bansal, Christopher Potts, and Adina Williams NAACL 2021

On the Importance of Adaptive Data Collection for Extremely Imbalanced Pairwise Tasks

Stephen Mussmann*, Robin Jia*, and Percy Liang EMNLP Findings 2020

With Little Power Comes Great Responsibility

Dallas Card, Peter Henderson, Urvashi Khandelwal, Robin Jia, Kyle Mahowald, and Dan Jurafsky EMNLP 2020

Selective Question Answering under Domain Shift

Amita Kamath, Robin Jia, and Percy Liang ACL 2020

Robust Encodings: A Framework for Combating Adversarial Typos

Erik Jones, Robin Jia*, Aditi Raghunathan*, and Percy Liang ACL 2020

Certified Robustness to Adversarial Word Substitutions
Robin Jia, Aditi Raghunathan, Kerem Göksel, Percy Liang EMNLP 2019

MRQA 2019 Shared Task: Evaluating Generalization in Reading Comprehension
Adam Fisch, Alon Talmor, Robin Jia, Minjoon Seo, Eunsol Choi, and Danqi Chen MRQA 2019

Document-Level N-ary Relation Extraction with Multiscale Representation Learning
Robin Jia, Cliff Wong, and Hoifung Poon NAACL 2019

Know What You Don't Know: Unanswerable Questions for SQuAD
Pranav Rajpurkar*, Robin Jia*, and Percy Liang ACL 2018
Best Short Paper Award.

Delete, Retrieve, Generate: A Simple Approach to Sentiment and Style Transfer
Juncen Li, Robin Jia, He He, and Percy Liang NAACL 2018

Adversarial Examples for Evaluating Reading Comprehension Systems
Robin Jia and Percy Liang EMNLP 2017
Outstanding Paper Award.

Learning Concepts through Conversations in Spoken Dialogue Systems
Robin Jia, Larry Heck, Dilek Hakkani-Tür, and Georgi Nikolov ICASSP 2017

Data Recombination for Neural Semantic Parsing
Robin Jia and Percy Liang ACL 2016

"Reverse Genomics" Predicts Function of Human Conserved Noncoding Elements
Amir Marcovitz, Robin Jia, and Gill Bejerano MBE 2016

Mx1 and Mx2 Key Antiviral Proteins are Surprisingly Lost in Toothed Whales
Benjamin A. Braun, Amir Marcovitz, J. Gray Camp, Robin Jia, and Gill Bejerano PNAS 2015

Note: * denotes equal contribution.

Grants

NVIDIA Hardware Grant, NVIDIA, 32,000 NVIDIA GPU hours (\approx \$200k value)
PI: Robin Jia. Co-PI's: Vatsal Sharan and Mahdi Soltanolkotabi. October 2025 – March 2026
Title: *Robust Mathematical Reasoning with Fourier Language Models*

Google ML and Systems Junior Faculty Award, Google, \$100k
PI: Robin Jia Awarded June 2025

NAIRR Pilot Allocation, NSF, 200,000 NVIDIA DGX Cloud GPU hours (\approx \$1.3M value)
PI: Robin Jia January 2025 – August 2025
Title: *Hubble: a suite to advance the study of LLM memorization*

IARPA BENGAL Award, IARPA, \$568k total (my share: \$485k)
PI: Robin Jia. Co-PI: Jesse Thomason. February 2026 – July 2027
Title: *Enabling Efficient Unlearning in Pretrained Large Language Models through Information Localization*

USC-Capital One CREDIF Award, Capital One, \$84k total (my share: \$42k)
Co-PI's: Sai Praneeth Karimireddy and Robin Jia January 2025 – December 2025
Title: *Privacy Preserving Synthetic Data Generation with Differentially Private Reward Models*

USC-Amazon Center Award, Amazon, \$75k + \$75k AWS credits

PI: Robin Jia

August 2024 – August 2025

Title: *Mechanistically-Aware Decoding for Ensuring Factuality During Retrieval Augmented Generation*

NSF Collaborative Research Medium Award, NSF, \$1.2M total (my share: \$207.5k)

Co-PI: Robin Jia. PI: Jordan Boyd-Graber.

Jul 2024 – Jun 2028

Other Co-PI's: Swabha Swayamdipta, John Lalor, Alvin Grissom.

Title: *Collaborative Research: RI: Medium: Hard Data to the Model: Personalized, Diverse Preferences for Language Models*

Google Research Scholar Award, Google, \$60k

PI: Robin Jia

May 2023 – May 2024

Title: *Stabilizing In-Context Learning by Understanding the Value of Demonstrations*

Cisco Research Award, Cisco, \$70k

PI: Robin Jia

Apr 2023 – Apr 2024

Title: *Estimating Large Language Model Capabilities without Labeled Data*

Open Philanthropy Research Grant, Open Philanthropy, \$320k

PI: Robin Jia

Aug 2021 – Aug 2024

Title: *Adversarial Robustness Research*

Students Supervised

Ph.D. students

Johnny Tian-Zheng Wei	Jan 2021 – present	
Ameya Godbole	Aug 2021 – present	
Wang (Bill) Zhu	Oct 2021 – present	Joint with Jesse Thomason
Ting-Yun (Charlotte) Chang	Jan 2022 – present	Joint with Jesse Thomason
Deqing Fu	Mar 2023 – present	Joint with Vatsal Sharan
		Best Research Assistant, USC CS (2025)
Tianyi Zhou	Aug 2024 – present	Joint with Vatsal Sharan
Yuqing Yang	Aug 2024 – present	
Muru Zhang	Aug 2024 – present	Joint with Swabha Swayamdipta

Masters and undergraduate students

Zitong (Cynthia) Huang (USC UG)	August 2025 – Present	
Nitya Kashyap (USC UG)	May 2025 – Nov 2025	
Miaosen Chai (USC UG)	May 2025 – Present	
Qitong (Tony) Yi (USC UG)	May 2025 – Present	
Debasish Modak (USC MS)	Apr 2025 – Aug 2025	
Xiaoyuan Zhu (USC UG)	Jan 2025 – Present	
Maggie Wang (Princeton UG)	Jun 2024 – Jan 2025	
Lorena Yan (USC UG)	Apr 2024 – Aug 2025	Now: Columbia Ph.D. student
Daniel Firebanks-Quevedo (USC MS)	Jan 2024 – Dec 2024	
Gustavo Adolpho Lucas De Carvalho (USC MS)	Jan 2024 – Present	
Tianyu Wang (USC UG)	Aug 2023 – Dec 2023	
Rahel Selemom (Brown UG)	Jun 2023 – Aug 2023	
Qilin Ye (USC UG)	Jun 2023 – Present	Now: Duke M.S. student
Ryan Wang (USC UG)	Apr 2023 – Aug 2025	Now: UC Berkeley Ph.D. student Viterbi Outstanding Research Award NSF GRFP Recipient Provost's Research Fellowship
Tianqi Chen (USC MS)	Mar 2023 – Aug 2025	
Anthony Sauer (USC UG)	Jan 2023 – Mar 2024	
Yuan Huang (USC MS)	Jun 2022 – Oct 2023	
Harvey Fu (USC UG)	May 2022 – Aug 2024	Now: UChicago Ph.D. student Provost's Research Fellowship
Adam Reynolds (USC MS)	Aug 2021 – Dec 2021	
Amita Kamath (Stanford MS)	Sep 2018 – Jun 2020	Now: UCLA Ph.D. student
Erik Jones (Stanford UG)	Jun 2019 – Dec 2019	Now: UC Berkeley Ph.D. student
Kerem Göksel (Stanford MS)	Jan 2019 – Jun 2019	Now: Semantic Machines

Faculty research project mentor for CAIS++ students Leslie Moreno, Aryan Gulati, and Aditya Kumar.

Ph.D. qualifying exam committee member: Hexiang Hu, Yury Zemlyanskiy, Zalan Fabian, Negar Mokhberian, Michiel de Jong, Pei Zhou, Qinyuan Ye, Jun Yan, Ming-Chang Chiu, Xisen Jin, Fei Wang, Yun Cheng Wang, Soumya Sanyal, Jake Bremerman, Johnny Wei, Bingyi Zhang, Ali Omrani, Lee Kezar, Wang Zhu, Brihi Joshi, Ting-Yun Chang, Justin Cho, Dong-Ho Lee, Ameya Godbole, James Huang, Yifei Huang, Huihan Li, Alireza Ziabari, Sahana Ramnath, Sayan Ghosh, Bahareh Harandizadeh, Yuliang Cai (32 total).

Ph.D. thesis proposal committee member: Hexiang Hu, Yury Zemlyanskiy, Yuchen Lin, Aaron Chan, Wenxuan Zhou, Michiel de Jong, Woojeong Jin, Ming-Chang Chiu, Jun Yan, Fei Wang, Xisen Jin, Ali Omrani, Qinyuan Ye, Alexander Spangher, Soumya Sanyal, Hae Jin Song, Dong-Ho Lee, Changzhi Xie, Tejas Srinivasan, James Huang, Johnny Wei (21 total).

Ph.D. thesis defense committee member: Yury Zemlyanskiy, Hanqing Zeng, Aaron Chan, Wenxuan Zhou, Yun-Cheng Wang, Ali Omrani, Jun Yan, Hassan Hamad, Dong-Ho Lee, Soumya Sanyal (10 total).

M.S. thesis committee member: Jeong Hyun An, Abid Hassan, Xinyue Cui (3 total).

Teaching

Instructor, CSCI 467: Introduction to Machine Learning
University of Southern California 4 times between Spring 2023 – Spring 2025

Instructor, CSCI 699: The Science of Large Language Models
University of Southern California Fall 2024

Instructor, CSCI 699: Generalization and Robustness in Natural Language Processing
University of Southern California Spring 2022

Teaching Fellow, CS 221 (Artificial Intelligence)
Stanford University Summer 2019
Instructor for 100-student class on artificial intelligence.

Teaching Assistant, CS 124 (Introduction to Natural Language Processing)
Stanford University Winter 2018

Head Teaching Assistant, CS 221 (Artificial Intelligence)
Stanford University Fall 2015
Head TA of 550-student class, managed a team of 18 TA's.

Section Leader, CS 106A (Introduction to Programming)
Stanford University Winter 2012
Taught a section of twelve students, graded assignments and exams.

Tutor, Stanford University Mathematical Organization
Stanford University Winter 2011 – Spring 2012
Tutored students in linear algebra and multivariable calculus. Served as coordinator of the tutoring program from September 2011 to June 2012.

Professional Service

Co-organizer of the First Workshop on Large Language Model Memorization (L2M2) at ACL 2025.

Volunteers Co-chair for NAACL 2025.

Faculty Committee Organizer for AAAI 2025 Symposium on Child-AI Interaction in the Era of Foundation Models.

General Chair for 2023 SoCal NLP Symposium.

Steering Committee member for Workshop on Instruction Tuning and Instruction Following at NeurIPS 2023.

Co-instructor of Tutorial on Uncertainty Estimation for Natural Language Processing at COLING 2022.

Co-organizer of the First Workshop on Dynamic Adversarial Data Collection (DADC) at NAACL 2022.

Co-organizer of the Third Workshop on Machine Reading and Question Answering (MRQA) at EMNLP 2021.

Co-instructor of Tutorial on Robustness and Adversarial Examples in Natural Language Processing at EMNLP 2021.

Co-organizer of the Second Workshop on Machine Reading and Question Answering (MRQA) at EMNLP 2019.

Co-organizer of the First Workshop on Machine Reading and Question Answering (MRQA) at ACL 2018.

Area chair for ACL (2021, 2023, 2024, 2025), EMNLP (2021, 2022, 2023, 2024, 2025), NAACL (2021, 2025), COLM (2025) and AKBC (2021).

Reviewer for NeurIPS (2024, 2025), ACL Rolling Review (2021, 2022, 2023, 2024, 2025), ACL

(2018, 2019, 2020), EMNLP (2018, 2019, 2020), NAACL (2019), TACL (2022, 2023, 2024), EACL (2022), COLM (2024), ICLR (2025), AACL (2020), ICML (2019), CoNLL (2018), AKBC (2019, 2022), RobustSeq Workshop (2022), ML Safety Workshop (2022), DistShift Workshop (2021, 2022, 2023), BlackboxNLP Workshop (2021, 2022, 2023, 2024), Repl4NLP Workshop (2021, 2023), GenBench Workshop (2023), ACL Student Research Workshop (2021), RobustML Workshop (2021), EMNLP DeepLo Workshop (2019), and NAACL GenDeep Workshop (2018). Outstanding Reviewer for EMNLP 2020.

NSF panel reviewer (2x).

Invited Talks	Princeton Language and Intelligence Talk	Nov 2025
	UPenn Computational Linguistics Invited Talk	Nov 2025
	Amazon BCN Invited Talk	Sep 2025
	JHU CLSP Seminar	Sep 2025
	UMD CLIP Colloquium Talk	Sep 2025
	Northeastern NLP Seminar	Sep 2025
	Harvard ML Foundations Seminar	Sep 2025
	MIT NLP Seminar	Sep 2025
	UChicago CS Invited Talk	Jun 2025
	Northwestern CS Seminar	Jun 2025
	Cornell AI Seminar	May 2025
	Solventum ML/AI Invited Talk	Jan 2025
	CMU AI Seminar Invited Talk	Dec 2024
	Simons Institute Domain Adaptation Workshop Invited Talk	Nov 2024
	UCLA CS Seminar Invited Talk	Nov 2024
	UC Berkeley EECS Invited Talk	Oct 2024
	UCSD CS Invited Talk	Oct 2024
	DSAA Panelist on GenAI	Oct 2024
	Stanford NLP Seminar	Jun 2024
	ICLR Safe and Trustworthy LLM Workshop Invited Talk	May 2024
	USC Generative AI Symposium Spotlight Talk	Mar 2024
	Knowing Machines Podcast	Oct 2023
	CHAI Workshop Plenary Talk	Jun 2023
	Capital One Research Invited Talk	Jun 2023
	UC Irvine AI/ML Seminar	May 2022
	Amazon Research Invited Talk	Apr 2022
	Princeton NLP Group Seminar	Jul 2021
	NLP Highlights Podcast	Mar 2021
	USC ISI Seminar	Feb 2021
	UC Santa Barbara NLP Seminar	Feb 2021
Other	Frederick Emmons Terman Engineering Scholastic Award, Stanford University	2014
	Finalist, Lunsford Oral Presentation of Research Award, Stanford University	2012
	Finalist, Boothe Prize for Excellence in Writing, Stanford University	2011
	Top 500 Scorer, William Lowell Putnam Mathematical Competition	2011
	Three-time Qualifier, USA Mathematics Olympiad	2008–2010
	Top Twenty Finalist, US National Chemistry Olympiad	2009