# Siddharth Karamcheti

## Education

2019 –	<ul> <li>Ph.D. in Computer Science at Stanford University, Stanford, CA.</li> <li>Co-advised by Dorsa Sadigh and Percy Liang.</li> <li>Supported by the Open Philanthropy AI Fellowship (5+ years of support).</li> <li>Thesis (Planned): Language-Driven Learning for Interactive Robotics.</li> </ul>
2015 – 2018 GPA 3.95	<ul> <li>Sc.B in Computer Science, A.B. in Literary Arts at Brown University, Providence, RI.</li> <li>Graduated Magna Cum Laude with Honors in Computer Science.</li> <li>Co-advised by Eugene Charniak and Stefanie Tellex.</li> <li>Thesis: Grounding Natural Language to Goals for Abstraction, Generalization, and Interpretability.</li> </ul>
2014 – 2015 GPA 3.92	<ul> <li>Electrical Engineering &amp; Computer Science, UC Berkeley, Berkeley, CA.</li> <li>Awarded the <i>Regents and Chancellor's Scholarship</i> (top 2% of incoming class).</li> <li>Transferred to Brown after first year to pursue interdisciplinary education in CS &amp; Creative Writing.</li> </ul>
	Research Experience
2023 -	<b>Toyota Research Institute</b> , <i>Los Altos, CA</i> . <i>Research Intern – Robotics</i> ; Large Behavior Models Team, led by Russ Tedrake. <b>Focus Areas:</b> Robot learning, data collection and curation at scale, vision-language-action models.
2021 - 2022	Hugging Face, Remote – Palo Alto, CA.Research Intern – Multimodal Learning; Core Science Team, led by Victor Sanh and Douwe Kiela.Focus Areas: Large-scale multimodal (text, vision, video, audio) learning, large language models.
2018 - 2019	<b>Facebook AI Research</b> , <i>New York</i> , <i>NY</i> . <i>AI Resident</i> ; Advised by Rob Fergus, Douwe Kiela, Jason Weston, and Arthur Szlam. <b>Focus Areas:</b> Deep reinforcement learning and natural language processing.
2017 – 2018	Bloomberg LP – CTO Research Group, New York, NY. Research Intern (Summer '17, '18, Part-Time '17-'18); Advised by Gideon Mann and David Rosenberg. Focus Areas: Machine learning for automated bug detection ("fuzzing"), active learning.
2016 - 2018	Human to Robots (H2R) Lab – Brown University, Providence, RI. Undergraduate Research Assistant; Advised by Stefanie Tellex. Focus Areas: Human-robot interaction: instruction following for mobile manipulation.
2016 – 2018	<ul> <li>Brown Lab for Linguistic Information Processing (BLLIP), Providence, RI.</li> <li>Undergraduate Research Assistant; Advised by Eugene Charniak.</li> <li>Focus Areas: Semantic parsing, question-answering, deep learning for natural language processing.</li> </ul>
	Honors and Awards
2024	Outstanding Paper Award Finalist at CoRL 2024 for "OpenVLA: An Open-Source Vision- Language-Action Model." Awarded to the top 6 of 671 submissions (1%).
2024	Selected for RSS Pioneers – a program for early-career roboticists (15% acceptance rate).
2023	Best Paper Award Finalist at RSS 2023 for "Language-Driven Representation Learning for Robotics." Awarded to the top 5 of 333 submissions (1.5%).
2022	Best Paper Award Finalist at the 2nd Workshop on Scaling Robot Learning at RSS 2022 for "What Makes Representation Learning from Videos Hard for Control?"

- 2022 Highlighted Reviewer at ICLR 2022.
- 2021 Outstanding Reviewer Award at EMNLP 2021.

- 2021 Outstanding Paper Award at ACL-IJCNLP 2021 for "Mind Your Outliers! Investigating the Negative Impact of Outliers on Active Learning for Visual Question Answering." Awarded to the top 5 of 3000+ submissions (0.2%).
- 2021 Outstanding Reviewer Award at ACL-IJCNLP 2021.
- 2019 Open Philanthropy Project AI Fellowship, offering full funding for 5+ years of the PhD, and access to broader network of fellows and affiliated research scientists.
- 2018 Selected for the Facebook AI Residency, 1/11 out of 2000+ applications.
- 2018 University Distinguished Thesis Nominee, Brown Computer Science. Sole CS-department nominee for University-wide undergraduate thesis prize.
- 2018 Honorable Mention Computing Research Association (CRA) Outstanding Undergraduate Researcher Prize. One of 45 student recognized in nationwide competition.
- 2018 Senior Prize, Brown University Computer Science Department. Recognized for outstanding record in research, teaching, and service.
- 2017 Best Paper Award at RoboNLP Workshop @ ACL 2017 for "A Tale of Two DRAGGNs: Interpreting Action and Goal-Oriented Instructions."
- 2014 Regents and Chancellor's Scholarship UC Berkeley. One of 200 incoming undergraduates recognized with the scholarship, as determined by the Academic Senate.

## Research Mentorship

### Stanford Students

- 2021-2023 Raj Palleti Undergraduate (continuing at Stanford for a coterminal Masters)
- 2022-2023 Michael Piseno Masters (continuing Masters at Stanford)
- 2021-2022 Madeline Liao Undergraduate (now a Machine Learning Engineer at Gridmatic)
- 2019-2021 Suvir Mirchandani Undergraduate (now a PhD student at Stanford)

### Visiting Summer Students

- 2024 Isabel Santoyo-Garcia Visiting as part of CS LINXS (now a senior at CSU Sacramento)
- 2022 Elise Corwin Visiting as part of Stanford CSLI (now a senior at Occidental College)
- 2021 Albert Wu Visiting as part of Caltech SURF (now a PhD student at UIUC)

### Teaching Experience

- Spring 2023 **Course Assistant for Stanford CS 224U: Natural Language Understanding**. Taught by Chris Potts. Gave two guest lectures integrated into the official online course offering. Held weekly sessions, supervised final projects.
- Autumn 2022 Head Course Assistant for Stanford CS 221: Introduction to Artificial Intelligence. Co-taught by Percy Liang and Dorsa Sadigh. Managed a staff of 14 TAs to design assignments, hold weekly sections, write exams, supervise final projects for an enrollment of 400+ students.
  - 2017-2018 **Head Teaching Assistant for Brown CS 1380: Distributed Systems**. Taught by Theophilius Benson. Managed a staff of 5 TAs to design assignments, grade, and hold additional tutorial sections and office hours. Built new projects on block-chain & cryptocurrency.
  - 2016-2017, Head Teaching Assistant for Brown CS 2950K/1470: Deep Learning.
  - 2017-2018 Taught by Eugene Charniak. Developed the first offering of the class (2016-2017), leading a group of 2 other TAs in designing and implementing all projects in Tensorflow 0.4. The following year (2017-2018), led a team of 12+ TAs as enrollment grew to 150+ students.
  - 2016-2017 **Head Teaching Assistant for Brown CS 1460: Computational Linguistics**. Taught by Eugene Charniak. Rewrote assignments in language modeling, machine translation, parsing, topic modeling, and deep learning for representation learning, managing a team of 7 TAs.

## Peer-Reviewed Academic Publications

CoRL 2024	<b>Vocal Sandbox: Continual Learning for Situated Human-Robot Collaboration</b> . Jennifer Grannen <sup>*</sup> , <b>Siddharth Karamcheti</b> <sup>*</sup> , Suvir Mirchandani, Percy Liang, Dorsa Sadigh <i>Conference on Robot Learning (CoRL), 2024</i> Oral Presentation
CoRL 2024	<b>OpenVLA: An Open-Source Vision-Language-Action Model</b> . Moo Jin Kim <sup>*</sup> , Karl Pertsch <sup>*</sup> , <b>Siddharth Karamcheti</b> <sup>*</sup> , Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan Foster, Grace Lam, Pannag Sanketi, Quan Vuong, Thomas Kollar, Benjamin Burchfiel, Russ Tedrake, Dorsa Sadigh, Sergey Levine, Percy Liang, Chelsea Finn <i>Conference on Robot Learning (CoRL), 2024</i> <b>Outstanding Paper Award Finalist</b>
RSS 2024	<b>DROID: A Large-Scale In-the-Wild Robot Manipulation Dataset</b> . DROID Collaboration (Research Lead – Data Curation & Annotation, Lab Lead) <i>Robotics: Science and Systems (RSS), 2024</i>
ICML 2024	Prismatic VLMs: Investigating the Design Space of Vision-Language Models. Siddharth Karamcheti, Suraj Nair, Ashwin Balakrishna, Percy Liang, Thomas Kollar*, Dorsa Sadigh* International Conference on Machine Learning (ICML), 2024
ICRA 2024	<b>Towards Grounded Social Reasoning</b> . Minae Kwon, Hengyuan Hu, Vivek Myers, <b>Siddharth Karamcheti</b> , Anca Dragan, Dorsa Sadigh <i>IEEE International Conference on Robotics and Automation (ICRA)</i> 2024
ICRA 2024	Open X-Embodiment: Robotic Learning Datasets and RT-X Models. Open X-Embodiment Collaboration IEEE International Conference on Robotics and Automation (ICRA), 2024 Best Paper Award Winner
NeurIPS 2023	OBELICS: An Open Web-Scale Dataset of Interleaved Image-Text Documents.
Datasets & Benchmarks	Hugo Laurençon, Lucile Saulnier, Léo Tronchon, Stas Bekman, Amanpreet Singh, Anton Lozhkov, Thomas Wang, <b>Siddharth Karamcheti</b> , Sasha Rush, Douwe Kiela, Matthieu Cord, Victor Sanh Advances in Neural Information Processing Systems (NeurIPS): Datasets and Benchmarks Track, 2023
RSS 2023	Language-Driven Representation for Robotics. Siddharth Karamcheti, Suraj Nair, Annie Chen, Thomas Kollar, Chelsea Finn, Dorsa Sadigh, Percy Liang <i>Robotics: Science and Systems (RSS), 2023</i>   Invited to IJRR Special Issue Best Paper Award Finalist
HRI 2023	"No, to the Right" – Online Language Corrections for Robotic Manipulation via
	Shared Autonomy.
	Yuchen Cui*, <b>Siddharth Karamcheti</b> *, Raj Palleti, Nidhya Shivakumar, Percy Liang, Dorsa Sadigh ACM/IEEE International Conference on Human Robot Interaction (HRI), 2023
CoRL 2022	Eliciting Compatible Demonstrations for Multi-Human Imitation Learning. Kanishk Gandhi, Siddharth Karamcheti, Madeline Liao, Dorsa Sadigh Conference on Robot Learning (CoRL), 2022
NeurIPS 2021	ELLA: Exploration through Learned Language Abstraction. Suvir Mirchandani, Siddharth Karamcheti, Dorsa Sadigh Advances in Neural Information Processing Systems (NeurIPS), 2021
CoRL 2021	LILA: Language-Informed Latent Actions. Siddharth Karamcheti <sup>*</sup> , Megha Srivastava <sup>*</sup> , Percy Liang, Dorsa Sadigh Conference on Robot Learning (CoRL), 2021
ACL 2021	Mind Your Outliers! Investigating the Negative Impact of Outliers on Active
	Learning for Visual Question Answering.
	<b>Siddharth Karamcheti</b> , Ranjay Krishna, Li Fei-Fei, Christopher D. Manning Annual Meeting of the Association of Computational Linguistics (ACL-IJCNLP), 2021 Outstanding Paper Award Winner   Main Stage Presentation

ICML 2021	Targeted Data Acquisition for Evolving Negotiation Agents.
	International Conference on Machine Learning (ICML), 2021
L4DC 2021	Learning Visually Guided Latent Actions for Assistive Teleoperation.
	Siddharth Karamcheti, Albert J. Zhai, Dylan P. Losey, Dorsa Sadigh Learning for Dynamics and Control (L4DC), 2021
IntEx-SemPar	Learning Adaptive Language Interfaces through Decomposition.
2020	Siddharth Karamcheti, Dorsa Sadigh, Percy Liang Workshop for Interactive and Executable Semantic Parsing (IntEx-SemPar) @ EMNLP, 2020 Oral Presentation
AAAI 2020	Generating Interactive Worlds with Text.
	Angela Fan*, Jack Urbanek*, Pratik Ringshia, Emily Dinan, Emma Qian, <b>Siddharth Karamcheti</b> , Shrimai Prabhumoye, Douwe Kiela, Tim Rocktäshel, Arthur Szlam, Jason Weston Association for the Advancement of Artificial Intelligence (AAAI), 2020
EMNLP 2019	Finding Generalizable Evidence by Learning to Convince Q&A Models.
	Ethan Perez, <b>Siddharth Karamcheti</b> , Rob Fergus, Jason Weston, Douwe Kiela, Kyunghyun Cho <i>Empirical Methods in Natural Language Processing (EMNLP), 2019</i>
	Learning to Speak and Act in a Fantasy Text Adventure Game.
	Jack Urbanek, Angela Fan, <b>Siddharth Karamcheti</b> , Saachi Jain, Samuel Humeau, Emily Dinan, Tim Rocktäshel, Douwe Kiela, Arthur Szlam, Jason Weston <i>Empirical Methods in Natural Language Processing (EMNLP), 2019</i>
ML4SE 2019	Improving Grey-Box Fuzzing by Modeling Program Control Flow.
	Siddharth Karamcheti, Gideon Mann, David Rosenberg Workshop on Machine Learning for Software Engineering (ML4SE), 2019
AuRO 2019	Grounding Natural Language Instructions to Semantic Goal Representations for Ab-
	Dilip Arumugam <sup>*</sup> , <b>Siddharth Karamcheti</b> <sup>*</sup> , Nakul Gopalan, Edward C. Williams, Mina Rhee, Lawson L.S. Wong, Stefanie Tellex
	Journal – Autonomous Robots (AuRO), 2019
AISEC 2018	Adaptive Grey-Box Fuzz Testing with Thompson Sampling.         Siddharth Karamcheti, Gideon Mann, David Rosenberg         ACM Workshop on Artificial Intelligence and Security (AISEC), 2018
DoboNII D	Oral Presentation A Tale of Two DPACCNes Interpreting Action and Coal Oriented Instructions
2017	Siddharth Karamcheti, Edward C. Williams, Dilip Arumugam, Mina Rhee, Nakul Gopalan,
	Lawson L.S. Wong, Stefanie Tellex
	Workshop in Language Grounding for Robotics (RoboNLP) @ ACL, 2017 Best Paper Award Winner   Oral Presentation
RSS 2017	Accurately and Efficiently Interpreting Instructions of Varying Granularities. Dilip Arumugam <sup>*</sup> , Siddharth Karamcheti <sup>*</sup> , Nakul Gopalan, Lawson L.S. Wong, Stefanie Tellex <i>Robotics: Science and Systems (RSS), 2017</i>   Invited to Autonomous Robots Special Issue
	Non-Archival Publications & Preprints
Preprint	<b>ProVox: Personalization &amp; Proactive Planning for Human-Robot Collaboration</b> . Jennifer Grannen, <b>Siddharth Karamcheti</b> , Blake Wulfe, Dorsa Sadigh Under Review – ACM/IEEE International Conference on Human Robot Interaction (HRI)
Preprint	Voltron-X: Capability-Aware Representation Learning for Robotics. Siddharth Karamcheti, Suraj Nair, Annie Chen, Thomas Kollar, Chelsea Finn, Dorsa Sadigh, Percy Liang Under Review – International Journal of Robotics Research (IJRR)

Workshop SRL	What Makes Representation Learning from Videos Hard for Control?.
2022	Tony Z. Zhao, <b>Siddharth Karamcheti</b> , Thomas Kollar, Chelsea Finn, Percy Liang 2nd Workshop on Scaling Robot Learning (SRL) @ RSS, 2022 Best Paper, Award Finalist
Workshop	Shared Autonomy for Robotic Manipulation with Language Corrections
MU-Sup 2022	Siddharth Karamahati, Bai Pallati, Yuchan Cui, Parcy Liang, Darca Sadigh
14L-5up 2022	Workshop on Learning with Natural Language Supervision (NL-Sup) @ ACL, 2022
Blog Post	The Annotated S4 – Structured State Space Models for Sequence Modeling.
	Sasha Rush and <b>Siddharth Karamcheti</b> – <i>ICLR Blog Track, 2022</i> A literate implementation and associated blog post for the Structured State Spaces for Sequence Modeling (S4) architecture, with analysis of long-range generalization capabilities. Blog    Library: https://github.com/srush/annotated-s4    ICLR Blog Reviews
Blog Post	Mistral: A Journey towards Reproducible Language Model Training.
	Siddharth Karamcheti <sup>*</sup> and Laurel Orr <sup>*</sup> Codebase & associated writeup detailing our journey in scaling large-scale language model pretraining as part of the development team for the Center for Research on Foundation Models. Blog    Code & Checkpoints: https://github.com/stanford-crfm/mistral
Manuscript	On the Opportunities & Risks of Foundation Models.
	<ul> <li>Center for Research on Foundation Models, led by Percy Liang. Authored the following sections:</li> <li><i>Robotics</i>: Siddharth Karamcheti (lead author), Annie Chen, Suvir Mirchandani, Suraj Nair, Krishnan Srinivasan, Kyle Hsu, Jeannette Bohg, Dorsa Sadigh, Chelsea Finn</li> <li><i>Interaction</i>: Joon Sung Park, Chris Donahue, Mina Lee, Siddharth Karamcheti, Dorsa Sadigh</li> </ul>
	Michael Bernstein
Undergraduate	Grounding Language to Goals for Abstraction, Generalization and Interpretability.
Thesis	Siddharth Karamcheti – Submitted for fulfillment of Sc. B. Honors in Computer Science
	Nominated for University-Level Honors (CS Department Nominee)
	Invited Talks
2024	Language-Driven Learning for Interactive Robotics.
	MIT – Computational Sensorimotor Learning Seminar
	NYU – GRAIL Lab Seminar (PI: Lerrel Pinto)
	Bay Area Robotics Symposium – Lightning Talk
	Apple – AI/ML Seminar
	RSS Pioneers – Spotlight Talk
	UC Berkeley – Semiautonomous Lab Seminar (PI: Shankar Sastry)
2024	Robotics in the Era of Foundation Models.
	Northwestern University – Shirley Ryan AbilityLab Seminar
2000	Joint taik with Quan vuong (Google Deepmind Robotics)
2023	Grounding Language Models for Perception and Action.
	Georgia Tech – Guest Lecture for CS 7643: Deep Learning (Fall 2023, taught by Danfei Xu)
2023	Vocal Sandbox – Continual Learning for Collaborative Robotics.
	Princeton University – Safe Autonomy Lab (PI: Jaime Fisac) Stanford University – Natural Language Processing Group
2023	Language-Driven Representation Learning for Robotics.
	Office of Naval Research – ONR Science of Autonomy Program
	Toyota Research Institute – Internal Research Seminar
	1X Technologies (Halodi) – Internal Talk
2022-2023	Fantastic Language Models and How to Build Them.
	Stanford University – Guest Lecture for CS 224U: Natural Language Understanding (Spring 2023) Georgia Tech – Guest Lecture for CS 7643: Deep Learning (Fall 2022, taught by Danfei Xu)

- 2023 Language and Interaction for Robotics. *Stanford University* – Guest Lecture for CS 237B: Principles of Robot Autonomy II (Spring 2023)
- 2022 **Bridging Language and Robotics for Interaction, Learning, and Teaching**. *NeurIPS Workshop* – Language and Reinforcement Learning (LaReL) Joint talk with Dorsa Sadigh (Stanford)
- 2022 Shared Autonomy and Natural Language for Robotic Manipulation. Johns Hopkins University – Applied Physics Lab Seminar Stanford University – Natural Language Processing Group
- 2021 **Robot Learning for Shared Autonomy**. *NeurIPS Workshop* – 4th Robot Learning Workshop: Self-Supervised and Lifelong Learning Joint talk with Dorsa Sadigh (Stanford)
- 2020 **Mistral: Transparent Foundations for Training Large Language Models**. *Stanford University* – Center for Research on Foundation Models Symposium Joint talk with Laurel Orr (Stanford)

## **Professional Activities**

### Workshop & Symposia Organization

- 2024 Co-organizer for the "Generative Modeling Meets HRI Workshop (GenAI-HRI)" at RSS with Felix Wang (MIT), Vignesh Prasad (TU Darmstadt), Georgia Chalvatzaki (TU Darmstadt), Pete Florence (Google Deepmind), Nadia Figueroa (UPenn), Fabio Ramos (NVIDIA), Julie Shah (MIT), Claudia D'Arpino (NVIDIA).
- 2023 Co-organizer for the **"2nd Workshop on Learning from Diverse, Offline Data (L-DOD)"** at *ICRA* with Ted Xiao (Google Brain), Fei Xia (Google Brain), Jackie Kay (Deepmind), Dhruv Shah (UC Berkeley), Suraj Nair (Stanford), Victoria Dean (CMU), Jeff Clune (UBC), Dorsa Sadigh (Stanford), and Ed Johns (ICL).
- 2022 Co-organizer for the **"Workshop on Foundation Models for Decision Making (FMDM)"** at *NeurIPS*, with Sherry Yang (UC Berkeley), Yilun Du (MIT), Jack Parker-Holder (Deepmind), Igor Mordatch (Google Brain), Shane Gu (Google Brain), and Ofir Nachum (Google Brain).
- 2022 Lead organizer for the **"Workshop on Learning from Diverse, Offline Data (L-DOD)"** at *RSS*, with Suraj Nair (Stanford), Dhruv Shah (UC Berkeley), Victoria Dean (CMU), and Professors Percy Liang (Stanford), Chelsea Finn (Stanford), and Dorsa Sadigh (Stanford).
- 2021 Co-organizer for the **"Bay Area Robotics Symposium (BARS)"** with Dorsa Sadigh (Stanford), Mark Mueller (UC Berkeley), and fellow student Erdem Biyik (Stanford).

#### Journal & Conference Reviewer

- Robotics Autonomous Robots (AuRo): 2023-
  - Robotics & Automation Letters (RA-L): 2020-
  - Conference on Robot Learning (CoRL): 2021-
  - Robotics: Science and Systems (RSS): 2024-
  - International Conference on Robotics & Automation (ICRA): 2020-
  - International Conference on Intelligent Robots & Systems (IROS): 2021-
  - ACM/IEEE Conference on Human-Robot Interaction (HRI): 2022-
  - ML Conference on Language Modeling (COLM): 2024-
    - Conference on Neural Information Processing Systems (NeurIPS): 2021-
    - International Conference on Learning Representations (ICLR): 2021-
    - International Conference on Machine Learning (ICML): 2021-

- NLP ACL Rolling Review (ARR): October 2021-2023 (Monthly)
  - Association for Computational Linguistics (ACL): 2020-
  - Empirical Methods in Natural Language Processing (EMNLP): 2020-
  - North American Association for Computation Linguistics (NAACL): 2020-2023

### Workshop Reviewer

- Robotics Workshop on Pretraining for Robotics (PTR) @ ICRA: 2023
  - Workshop on Foundation Models for Decision Making (FMDM) @ NeurIPS: 2022, 2023
  - Workshop on Learning from Diverse, Offline Data (L-DOD) @ RSS: 2022, 2023
  - Robot Learning: Self-Supervised & Lifelong Learning @ NeurIPS: 2021
  - Social Intelligence in Humans and Robotics @ ICRA: 2021
  - ML $\,$   $\,\circ\,$  Mathematical & Empirical Understanding of Foundation Models @ ICLR: 2023, 2024
    - $\circ~$  Workshop on Deep Reinforcement Learning @ NeurIPS: 2021, 2022
    - Interactive Machine Learning Workshop (IML) @ AAAI: 2022
    - Controllable Generation (CtrlGen) @ NeurIPS: 2021
  - NLP o Learning with Natural Language Supervision (NL-Sup) @ ACL: 2022
    - Interactive Learning for Natural Language Processing (InterNLP) @ ACL: 2021
    - Visually Grounded Interaction and Language (ViGiL) @ NAACL: 2021
    - Novel Ideas in Learning to Learn from Interaction (NILLI) @ EMNLP: 2021
    - Insights from Negative Results in NLP (Insights) @ EMNLP: 2021
    - Spatial Language and Grounding for Robotics (Splu-RoboNLP) @ ACL: 2021

### University Service

- 2020 2022 Editor for the Stanford Artificial Intelligence Blog; responsible for editing student-authored posts and managing the blog website.
- 2019 2021 Organized Grounding & Interaction Reading Group, bringing together students and faculty from NLP, CV, RL, Robotics, and beyond to discuss and present recent work.
- 2020 2021 Reviewed Ph.D. applications for Stanford's Computer Science department as part of the Applications Committee.