

Siddharth Karamcheti

Academic Objective: PhD in Computer Science

sidd.karamcheti@gmail.com • [linkedin.com/in/siddharthkaramcheti](https://www.linkedin.com/in/siddharthkaramcheti) • siddkaramcheti.com • github.com/siddk

Education

Brown University, Providence, RI

Aug 2015 - Present

Sc.B. Computer Science, A.B. Literary Arts. Technical GPA: 4.0 / 4.0

Transferred from UC Berkeley to pursue an integrated CS + Literature program

CS Advisors: Professors Eugene Charniak and Stefanie Tellex.

Courses: Deep Learning, ML, NLP, Algorithms, Optimization, Reinforcement Learning, Robotics, Probability, Compilers, Formal Methods, Programming Languages, Distributed Systems.

University of California, Berkeley, Berkeley, CA

Aug 2014 - Aug 2015

B.S. Electrical Engineering and Computer Science (Regent's Scholar). Technical GPA: 3.92 / 4.0

Courses: Computer Architecture, Data Structures/Algorithms, Discrete Math, Introduction to Computer Science

Research Experience

CTO Research Group, Bloomberg LP - Advised by Gideon Mann

May 2017 - Present

- Internship as part of the Chief Technology Office - Data Science Dept. at Bloomberg LP in New York.

- *Current Research:* Reinforcement Learning for Smart Program Analysis + Automatic Bug Detection (Program "Fuzzing").

- *Brief Summary:* There are many existing methods for automatic bug detection, utilizing randomness or other naive methods to test programs, resulting in suboptimal performance. How do we develop intelligent agents capable of leveraging these tools in different contexts, in order to optimize the rate of discovering new vulnerabilities?

H2R (Human to Robots Lab), Brown CS - Advised by Stefanie Tellex

Jan 2016 - Present

- *Current Research:* Weakly Supervised Language Grounding via Reinforcement Learning, Language and Partial Observability

- *Previous:* Grounding Language to Actions and Goals (*RoboNLP, 2017*), Interpreting Human-Robot Instructions of Varied Complexities with Abstract Markov Decision Processes (*RSS, 2017*).

- *Brief Summary:* Language provides a clean and intuitive way to communicate with robots. However, interpretation is difficult, as inherent in understanding are problems of identifying specificity and intent. How do we approach these problems, and develop models that can both accurately interpret human-robot instructions, and use the obtained information to efficiently execute tasks?

BLLIP (Natural Language Processing), Brown CS - Advised by Eugene Charniak

Jan 2016 - Present

- *Current Research:* Semantic Parsing, End-to-End Methods for Interpretable Question-Answering

- *Previous:* Interpretable Question Answering via Grounding to External World Updates (*In Review at NAACL 2018*).

- *Brief Summary:* Existing state-of-the-art methods for question-answering on short stories require significant amounts of data, and often result in uninterpretable learned representations that live as numeric vectors in a deep neural network. How do we develop models that more sample efficient, with learned representations that are easily understood by users and other systems?

Accepted Publications

A Tale of Two DRAGGNs: A Hybrid Approach for Interpreting Action and Goal-Oriented Instructions

Siddharth Karamcheti, Edward Williams, Dilip Arumugam, Mina Rhee, Nakul Gopalan, Lawson Wong, and Stefanie Tellex.

Workshop in Language Grounding for Robotics (RoboNLP) at Association for Computational Linguistics, 2017.

Winner of the RoboNLP Best Paper Award. arXiv: <https://arxiv.org/abs/1707.08668>

Accurately and Efficiently Interpreting Human-Robot Instructions of Varying Granularities

Dilip Arumugam*, Siddharth Karamcheti*, Nakul Gopalan, Lawson Wong, and Stefanie Tellex.

Conference for Robotics: Science and Systems (RSS) 2017. arXiv: <https://arxiv.org/abs/1704.06616>

Invited for Submission to Autonomous Robots (AuRo) Journal (RSS Special Issue)

Submissions in Review

World Interaction Networks: Grounding Natural Language to World Updates with Minimal Supervision

Siddharth Karamcheti, Eugene Charniak, and Stefanie Tellex.

Modeling Latent Attention within Neural Networks

Chris Grimm, Dilip Arumugam, Siddharth Karamcheti, David Abel, Lawson Wong, and Michael Littman.

Skills

Programming Languages: *Fluent:* Python, Java, Go, MATLAB | *Comfortable:* C, MIPS/x86 Assembly, Scheme, Lua

Technologies: Tensorflow, Keras, Numpy, CoreNLP, SpaCy, Spark, MapReduce, Redshift, Chef, Flask

Skills: Natural Language Processing, Deep Learning, Reinforcement Learning, Robotics

Teaching Experience

Head Teaching Assistant Oct 2017 - Present

CS 1380: Distributed Systems - Brown, Spring 2018

- Chose to TA over NLP/ML courses for unique opportunity to develop new course material and lectures, with new professor.
- Revised course assignments (written in Golang), adding additional assignments covering block-chains & cryptocurrencies.
- Course taught by Professor Theo Benson: <http://cs.brown.edu/courses/cs138/>

Head Teaching Assistant Jun 2016 - Present

CS 2950K/1470: Deep Learning - Brown, Fall 2016, Fall 2017

- Designed course assignments, pulling from topics in NLP, Vision, and Reinforcement Learning.
- Course taught by Professor Eugene Charniak: <http://cs.brown.edu/courses/cs1470/>

Head Teaching Assistant Jan 2017 - Jun 2017

CS 1460: Computational Linguistics - Brown, Spring 2017

- Revised course assignments, covered language modeling, translation, parsing, topic modeling, and deep learning for NLP.
- Course taught by Professor Eugene Charniak: <http://cs.brown.edu/courses/csci1460/>

Group Tutor/Teaching Assistant Jan 2015 - May 2015

CS 61A: A Structure and Interpretation of Computer Programs - Berkeley, CA

- Taught small recitation sections meant to supplement lecture, holding additional sections when necessary.
- Wrote additional sections and revised parts of the course textbook, *Composing Programs*.

Work Experience

CTO Research Intern May 2017 - Present

Bloomberg LP - New York City, NY

- Research in program “fuzz” testing and analysis supervised by Dr. Gideon Mann.
- See above in Research Experience for more information.

Software Engineering Intern May 2016 - Aug 2016

Wealthfront - Redwood City, CA

- Worked primarily on the backend, building out significant parts of production systems.
- Designed and implemented a brand new service, for use in day-to-day funding operations.

Natural Language Processing/Machine Learning Intern May 2015 - Nov 2015

WriteLab - Berkeley, CA

- Built tools for clausal analysis and modifier usage by examining sentence balance and adjective branching.
- Created a system to track topics, entities, and topic dependencies across an essay or larger text.

Software Development/Research Intern Jun 2013 – Nov 2013

AutoGrid Systems - Redwood City, CA

- Used Hadoop Map-Reduce to load and analyze Electric Smart Meter data in real-time.
- Ran data through a predictive machine learning algorithm to predict demand spikes.

Miscellaneous Projects

Personal Projects

- *ML Research Replications*: Focus on recent papers in Question Answering and Program Induction;
 - + “Tracking the World State with Recurrent Entity Networks” Henaff et. al — (github.com/siddk/entity-network)
 - + “Neural Programmer-Interpreters” Reed, de Freitas — (github.com/siddk/npj)
- *Deep Learning Tutorials with Tensorflow*: Implemented several different state of the art Deep Learning models in Tensorflow.
 - + Variational Autoencoders, Encoder-Decoder MT, Deep-Q Learning, Memory Networks, Image Classifiers.
 - + Github Link: github.com/siddk/deep-nlp

Course Projects

- *CS 1380: Distributed Systems* - Puddlestore
 - + Fully designed and implemented Puddlestore, a fully-functional distributed file system, in Golang.
 - + Leveraged Apache Zookeeper for configuration, the Raft protocol for maintaining consensus across nodes, and Tapestry as the Distributed Object Location and Retrieval Service.
- *CS 1260: Compilers* - PyDecaf
 - + Designed a compiler for the DECAF Programming Language, with full Parser, Abstract Syntax Tree, Semantics, and Code Generation.