Shreya Shankar

shreyashankar@berkeley.edu
www.sh-reya.com
sh_reya
shreyashankar

Education

Aug 2021–present **University of California, Berkeley**, Berkeley, CA.
Ph.D. in Electrical Engineering and Computer Sciences (Databases)
Advised by Aditya G. Parameswaran

Sep 2015–Dec 2020 **Stanford University**, Stanford, CA.
M.S. in Computer Science (Artificial Intelligence), done part-time
B.S. in Computer Science (Systems)
Advised by Pat Hanrahan

Experience

Industry

May 2022–Nov 2022 **Research Engineer**, *Meta*, Menlo Park, CA.
Researched and engineered an automatic data validation system for ML pipeline monitoring. Published a paper at CIKM 2023.

March 2021–Aug 2021 **Entrepreneur in Residence**, *Amplify Partners*, Menlo Park, CA.
Built open-source tools for machine learning software development (MLOps). Press release here.

June 2019–Jan 2021 **Machine Learning Engineer**, *Viaduct*, Palo Alto, CA.
Built systems and machine learning methods for large-scale time series data as the first ML engineer.
Worked with Airflow, Spark, SQL, Python, TensorFlow 2.0, XGBoost, Spark MLlib, and more.

Sep 2017–April 2019 **Research Intern**, *Google Brain*, Mountain View, CA.
Researched machine learning security and adversarial examples in collaboration with Stanford AI Lab.
Worked with TensorFlow 1.0, Python, and Borg.
Advised by Alex Kurakin and Ian Goodfellow.

June 2017–Sep 2017 **Software Engineering Intern**, *Facebook*, New York, NY.
Worked on Facebook's civic engagement team to connect users to their government representatives.
Worked with Hack (PHP), ReactJS, SQL, and Python.

Teaching

August 2022-Dec 2022 **Graduate Student Instructor**, *UC Berkeley*, Berkeley, CA.
Served as a TA for DATA 101, a course on data engineering fundamentals. Held weekly office hours, taught weekly sections, and wrote some homeworks.

April 2020-June 2020 **Teaching Assistant**, *Stanford University*, Stanford, CA.
Served as a TA part-time for a remote version of CS110 (Principles of Computer Systems). Taught weekly sections and held weekly office hours via Zoom.

June 2018–Dec 2018 **Head Teaching Assistant**, *Stanford University*, Stanford, CA.
Served as head TA for CS106B (Programming Abstractions) and CS101 (Introduction to Computing Principles). Held weekly office hours. Helped write exams and homework grading criteria. Coordinated a staff of undergraduate section leaders.

Jan 2016–April 2018 **Undergraduate Section Leader**, *Stanford University*, Stanford, CA.
Taught weekly sections for CS106A (Programming Methodologies) and CS106B (Programming Abstractions). Held weekly office hours. Graded assignments and exams.

Preprints and Publications

[1] **Shreya Shankar**, Rolando Garcia, Joseph M. Hellerstein, and Aditya G. Parameswaran. "we have no idea how models will behave in production until production": How engineers operationalize machine learning. *Proc. ACM Hum.-Comput. Interact.*, 8(CSCW1), apr 2024. DOI: 10.1145/3653697.

[2] Madelon Hulsebos, Wenjing Lin, **Shreya Shankar**, and Aditya Parameswaran. It took longer than i was expecting: Why is dataset search still so hard? In *Proceedings of the 2024 Workshop on Human-In-the-Loop Data Analytics*, HILDA 24, page 1–4, New York, NY, USA, 2024. Association for Computing Machinery. DOI: 10.1145/3665939.3665959.

[3] Yiming Lin, Madelon Hulsebos, Ruiying Ma, **Shreya Shankar**, Sepanta Zeighami, Aditya G Parameswaran, and Eugene Wu. Towards accurate and efficient document analytics with large language models. *arXiv preprint arXiv:2405.04674*, 2024.

[4] **Shreya Shankar**, Haotian Li, Parth Asawa, Madelon Hulsebos, Yiming Lin, JD Zamfirescu-Pereira, Harrison Chase, Will Fu-Hinthorn, Aditya G Parameswaran, and Eugene Wu. Spade: Synthesizing assertions for large language model pipelines. *arXiv preprint arXiv:2401.03038*, 2024.

[5] **Shreya Shankar** and Aditya G. Parameswaran. Building reactive large language model pipelines with motion. In *Companion of the 2024 International Conference on Management of Data*, SIGMOD/PODS '24, page 520–523, New York, NY, USA, 2024. Association for Computing Machinery. DOI: 10.1145/3626246.3654734.

[6] **Shreya Shankar**, J.D. Zamfirescu-Pereira, Björn Hartmann, Aditya G Parameswaran, and Ian Arawjo. Who validates the validators? aligning llm-assisted evaluation of llm outputs with human preferences. *arXiv preprint arXiv:2404.12272*, 2024.

[7] Aditya G Parameswaran, **Shreya Shankar**, Parth Asawa, Naman Jain, and Yujie Wang. Revisiting prompt engineering via declarative crowdsourcing. *arXiv preprint arXiv:2308.03854*, 2023.

[8] **Shreya Shankar**, Labib Fawaz, Karl Gyllstrom, and Aditya Parameswaran. Automatic and precise data validation for machine learning. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, CIKM '23, page 2198–2207, New York, NY, USA, 2023. Association for Computing Machinery. DOI: 10.1145/3583780.3614786.

[9] Vijay Veerabadran, Josh Goldman, **Shreya Shankar**, Brian Cheung, Nicolas Papernot, Alexey Kurakin, Ian Goodfellow, Jonathon Shlens, Jascha Sohl-Dickstein, Michael C Mozer, et al. Subtle adversarial image manipulations influence both human and machine perception. *Nature Communications*, 14(1):4933, 2023.

[10] **Shreya Shankar**, Rolando Garcia, Joseph M Hellerstein, and Aditya G Parameswaran. Operationalizing machine learning: An interview study. *arXiv preprint arXiv:2209.09125*, 2022.

[11] **Shreya Shankar**, Bernease Herman, and Aditya G. Parameswaran Parameswaran. Rethinking streaming machine learning evaluation. In *ICLR 2022 workshop: ML Evaluation Standards*, 2022.

[12] **Shreya Shankar**, Stephen Macke, Andrew Chasins, Andrew Head, and Aditya Parameswaran. Bolt-on, compact, and rapid program slicing for notebooks. *Proceedings of the VLDB Endowment*, 15(13):4038–4047, 2022.

[13] **Shreya Shankar** and Aditya Parameswaran. Towards observability for production machine learning pipelines. *Proceedings of the VLDB Endowment*, 15(13):4015–4022, 2022.

[14] Sumanth Dathathri, Krishnamurthy Dvijotham, Alexey Kurakin, Aditi Raghunathan, Jonathan Uesato, Rudy R Bunel, **Shreya Shankar**, Jacob Steinhardt, Ian Goodfellow, Percy S Liang, and Pushmeet Kohli. Enabling certification of verification-agnostic networks via memory-efficient semidefinite programming. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 5318–5331. Curran Associates, Inc., 2020.

[15] Gamaleldin F. Elsayed, **Shreya Shankar**, Brian Cheung, Nicolas Papernot, Alexey Kurakin, Ian Goodfellow, and Jascha Sohl-Dickstein. Adversarial examples influence human visual perception. *Journal of Vision*, 19(10):190c–190c, Sep 2019. DOI: 10.1167/19.10.190c.

[16] Gamaleldin F. Elsayed, **Shreya Shankar**, Brian Cheung, Nicolas Papernot, Alexey Kurakin, Ian Goodfellow, and Jascha Sohl-Dickstein. Adversarial examples that fool both computer vision and time-limited humans. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NeurIPS'18, page 3914–3924. Curran Associates, Inc., 2018.

[17] **Shreya Shankar**, Yoni Halpern, Eric Breck, James Atwood, Jimbo Wilson, and D. Sculley. No classification without representation: Assessing geodiversity issues in open data sets for the developing world. In *NIPS 2017 workshop: Machine Learning for the Developing World*, 2017.

Honors, Awards, and Funding

- 2024: Berkeley EECS Evergreen Mentorship Award Recipient
- 2023: Received Funding from Modal Labs (Industry Sponsorship)
- 2023: Heidelberg Laureate Forum Young Scholar
- 2022: NDSEG Fellowship Recipient
- 2022: Hertz Foundation Fellowship Finalist
- 2022: P.D. Soros Fellowship Finalist
- 2021: UC Berkeley EECS Excellence Award Recipient
- 2020: Interact Fellowship Recipient
- 2016: Anita Borg Grace Hopper Scholarship Recipient
- 2016: Palantir Women in Technology Scholarship Recipient

Selected Invited Talks

August 2024 **HCI Seminar**, *HKUST*.
Giving a talk on the SPADE and EvalGen projects (data quality for LLM pipelines).

July 2024 **Databricks**, *Databricks*.
Gave a talk on the SPADE and EvalGen projects (data quality for LLM pipelines). Invited by Alkis Polyzotis.

April 2024 **MLSys Seminar**, *Stanford University*.
Gave a talk on the SPADE and EvalGen projects (data quality for LLM pipelines).

April 2024 **ML in Practice Seminar**, *Carnegie Mellon University*.
Gave a talk on the SPADE and EvalGen projects (data quality for LLM pipelines). Invited by Prof. Virginia Smith.

April 2024 **EPIC Lab Retreat**, *UC Berkeley*.
Gave a talk on the SPADE and EvalGen projects (data quality for LLM pipelines).

January 2024 **Amazon Research**, *Amazon*.
Gave a talk on the SPADE project (assertions for LLM pipelines).

November 2023 **INDE Lab**, *University of Amsterdam*.
Discussed my research on tooling for ML engineers. Invited by Prof. Sebastian Schelter.

March 2023 **Core Data Tech Talk**, *Google*.
Discussed my research on ML observability.

December 2022 **Normconf**, *Normconf*.
Gave a normie talk on data management for ML.

October 2022 **Gradient Dissent Podcast**, *Weights & Biases*.
Discussed our interview study of ML engineers.

February 2022 **CS329S**, Stanford University, Stanford, CA.
Gave a guest lecture on detecting distribution shift in data streams.

November 2021 **Toronto Machine Learning Virtual Summit**, *Toronto ML Society*, Toronto, Canada.
Gave a talk on observability for ML systems and tutorial on building a ML pipeline with testing and monitoring.

October 2021 **Data Observability Summit**, *Facebook*, Menlo Park, CA.
Gave a talk on observability for ML systems.

February 2021 **MLSys Seminar**, *Stanford University*, Stanford, CA.
Gave a talk on debugging ML in production. Code and slides on my Github.

February 2021 **DSC102**, *University of California, San Diego*, San Diego, CA.
Gave a talk on debugging ML in production. Code and slides on my Github.

Software

Motion — This research project provides a framework for building continually-updating ML applications.

GATE — This research project enables automatic and precise data validation for machine learning.

mltrace — This project enables coarse-grained lineage and tracing in complex data pipelines. *400+ stars.*

Toy ML Pipeline — This is a toy example of a standalone ML pipeline written entirely in Python. No external tools are incorporated into the master branch. I built it mainly to experiment with my ideas for ML tooling. *150+ stars.*

Create ML App — This project makes it easier to spin up a machine learning project locally in Python and handle various package dependencies using a Makefile. It abstracts away pip installs and virtual environment commands from the user. *500+ stars.*

GPT3 Sandbox — This project enables users to create cool web demos using OpenAI's GPT-3 API with just a few lines of Python. Co-authored with Bora Uyumazturk. *2.5k+ stars.*

Service

- On the student admissions committee for the UC Berkeley EECS PhD Program.

- Co-organizer of DEEM, a workshop at SIGMOD on end-to-end data management for machine learning.
- Board member of SHE++, a 501(c)(3) nonprofit that improves diversity in tech.
- Former co-director of SHE++, a 501(c)(3) nonprofit that improves diversity in tech.
- Former financial officer of Stanford WiCS (Women in Computer Science).

Mentoring

Current

- Parth Asawa (2 semesters + summer, CRA Undergraduate Award Honorable Mention)
- Rachel Lin (1 semester)

Past

- Boyuan Deng (2 semesters)
- Aditi Mahajan (2 semesters)
- Peter Maldonado (Summer)
- Yujie Wang (3 semesters + summer)

Reviewing

- UIST 2024
- ICLR 2022
- NeurIPS 2021
- ICML 2019 Workshop in Adversarial Machine Learning in Real-World Computer Vision Systems
- ICML 2019 Workshop in Security and Privacy of Machine Learning
- NeurIPS 2018 Workshop on Security in Machine Learning

Interests

Triathlons	Competed for Stanford's Triathlon team. Completed 2021 Ironman 70.3 Santa Cruz.
Writing	Member of a weekly writer's group in San Francisco. Technical writing available at personal website.
Intentional communities	Alum of Phoenix House and Haight Street Commons, a network of co-ops in the Bay Area.