

Harini Suresh

hsuresh@mit.edu
https://harinisuresh.com

Education

Massachusetts Institute of Technology (MIT) - Cambridge, MA

Bachelor of Science in Computer Science, 2016

Master of Engineering in Computer Science, 2017

PhD in Computer Science, current

Thesis committee: John Guttag, Arvind Satyanaran, Catherine D’Ignazio

Research labs: Clinical and Applied Machine Learning Group (CSAIL), Visualization Group (CSAIL),

Data + Feminim Lab (DUSP)

GPA: 4.8/5.0

Selected Publications

Harini Suresh, Divya Shanmugam, Anna Bryan, Tiffany Chen, Alexander D’Amour, John V. Guttag, Arvind Satyanarayan. [Kaleidoscope: Semantically-grounded, context-specific ML model evaluation](#). CHI Conference on Human Factors in Computing Systems (CHI ’23).

Harini Suresh, Rajiv Movva, Amelia Lee Dogan, Rahul Bhargava, Isadora Cruxên, Ángeles Martinez Cuba, Giulia Taurino, Wonyoung So, and Catherine D’Ignazio. [Towards Intersectional Feminist and Participatory ML: A Case Study in Supporting Femicide Counterdata Collection](#). ACM Conference on Fairness, Accountability, and Transparency (FAccT ’22). ★ **Best paper award** ★

Harini Suresh, Kathleen M. Lewis, John V. Guttag, Arvind Satyanarayan. [Intuitively Assessing ML Model Reliability through Example-Based Explanations and Editing Model Inputs](#). ACM Conference on Intelligent User Interfaces (IUI ’22).

Harini Suresh, John Guttag. [A Framework for Understanding Sources of Harm throughout the Machine Learning Life Cycle](#). ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO ’21).

Harini Suresh, Steven R. Gomez, Kevin K. Nam, Arvind Satyanarayan. [Beyond Expertise and Roles: A Framework to Characterize the Stakeholders of Interpretable Machine Learning and their Needs](#). CHI Conference on Human Factors in Computing Systems (CHI ’21).

Susanne Gaube*, **Harini Suresh*** (co-first author), Martina Raue, Alexander Merritt, Seth J. Berkowitz, Eva Lermer, Joseph F. Coughlin, John V. Guttag, Errol Colak, Marzyeh Ghassemi. [Do as AI say: susceptibility in deployment of clinical decision-aids](#). npj Digital Medicine, 2021.

Harini Suresh, Natalie Lao, Ilaria Liccardi. [Measuring the Interference of Machine Learning in Human Decision-Making](#). ACM Conference on Web Science (WebSci '20). ACM Conference Proceedings.

Harini Suresh*, Jen Gong*, John Guttag. [Learning Tasks for Multitask Learning: Heterogeneous Patient Populations in the ICU](#). Conference on Knowledge Discovery and Data Mining (KDD 2018).

Harini Suresh, Nathan Hunt, Alistair Johnson, Leo Anthony Celi, Peter Szolovits, Marzyeh Ghassemi. [Clinical Event Prediction and Understanding using Neural Networks](#). Machine Learning for Healthcare Conference (MLHC '17).

See my [Google Scholar profile](#) for a full list of publications.

Teaching, Leadership and Service

Advanced Natural Language Processing, Teaching Assistant

<https://www.mit.edu/~jda/teaching/6.864/sp21> | Spring 2021

Graduate teaching assistant for Prof. Jacob Andreas. Developed and taught a new course module on dataset design and its ethical implications with a hands-on problem set. This material received positive feedback from ~250 students, has continued to be taught in subsequent years, made widely available on MIT OpenCourseWare, and adapted for different classes.

ML Tidbits Inc., Co-Director

<https://mltidbits.github.io> | Fall 2018 - Present

Co-director of ML Tidbits, an educational non-profit intended to empower the public to understand and discuss machine learning concepts and their societal effects. Wrote, illustrated and published engaging short [videos](#) on YouTube with quantitative [pedagogical benefits](#).

AI Ethics Discussion Group, Lead Organizer

<https://mitaiethics.github.io> | October 2018 - May 2020

Co-founded and led an MIT-wide initiative to connect community members on topics around AI Ethics and facilitate interdisciplinary conversations. Organized bi-weekly meetings, curated readings, and facilitated discussions.

Intro to Deep Learning, Lead Organizer

<http://introtodeeplearning.com/2017> | January 2017

Created and led an extensive introduction to the field of deep learning, covering applications to machine translation, image recognition, game playing, image generation and more. Counted as a for-credit course at MIT. Included hands-on labs in TensorFlow and peer brainstorming sessions. 250+ students attended.

MIT Graduate Student Advisory Group, Member

<https://engineering.mit.edu/about/leadership/graduate-student-advisory-group> | 2019 - 2020

Met regularly with the Dean of Engineering to discuss and address student needs, including student-advisor relationships and DEI issues.

Program Committees and Reviewing

ACM Transactions on Interactive Intelligent Systems	2022
ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization	2021
Harvard CRCS Workshop on AI for Social Good	2020
Conference on Neural Information Processing Systems (NeurIPS)	2019
Fair ML for Health Workshop @ NeurIPS	2019
Machine Learning for Health Workshop @ NeurIPS	2019

Professional Experience

Research Intern, Google Research, People + AI Research (PAIR) Team

Fall 2020 | Remote

- Designed and implemented methods for visualizing long narrative text with sentence embeddings
- Produced a backend pipeline for applying methods to new text and an [open-sourced interactive visualization](#)

Research Intern, Google Research, Brain Fairness Team

Summer 2018 | Cambridge, MA

- Implemented a multipart production system to train embedding models and quantify sources of downstream bias with large text datasets
- Contributed to a resulting JMLR [publication](#) on underspecification in ML models

Flybridge Venture Capital, ML Consultant

2018 - 2019 | Boston, MA

- Attended startup pitches and served as an expert for ML-related issues
- Conducted additional interviews with portfolio companies and prospects, wrote up summaries of technical strengths and weaknesses

Data Science Intern, Jawbone

Summer 2015 | San Francisco, CA

- Built models of how different behaviors affected users' heart rates using longitudinal timeseries data from Jawbone UP bands
- Shipped personalized insights around auto travel detection, workout patterns, and heart rate trends, which received double the average App Content Score within first 48 hours of release

Data Science Intern, Zephyr Health

Summer 2014 | San Francisco, CA

- Using raw Medicare data, developed a multipurpose analytic application in R for pharmaceutical companies to view and compare prescriptions from specific disease areas, extract information about relevant drugs and healthcare providers, and design an efficient go-to-market strategy

External Engagement

Invited Talks, Lectures, and Panels

- 2022 **National Institute of Standards and Technology (NIST)** AI Risk Management Framework Workshop. *Map: Values in Context* (panel).
- Ethical ML in Human Deployments** at MIT. *Sources of Harm through the Machine Learning Lifecycle* (invited lecture).
- 2021 **Institute for Artificial Intelligence and Fundamental Interactions (IAIFI)**. *A Framework for Understanding Sources of Harm throughout the Machine Learning Life Cycle* (talk).
- Information+ Conference**. *Communicating Uncertainty in Machine Learning Systems* (talk and panel).
- MIT AI Ethics and Policy**. *Data generation through the lens of Data Feminism* (talk).
- Mechanism Design for Social Good** working group. *Data generation through the lens of Data Feminism* (talk).
- 2020 **Priscilla King Gray Public Service Center @ MIT**. *Femicide & Machine Learning* (invited lecture).
- Experiential Ethics @ MIT**. *Sources of harm in ML* (talk and panel).
- UCL-Toronto Ethical Innovation for AI workshop**. *Understanding and Preventing Unintended Consequences of ML* (talk and panel).
- MIT Political Science**. *Understanding and Mitigating Unintended Consequences of ML* (invited lecture).
- 2019 **MIT Better World symposium** in Atlanta, GA. *Trust Issues in Machine Learning* (talk and panel).
- Fair ML in Health at Data & Society Research Institute** in New York, NY. *Deploying decision-aids: real-world considerations* (talk).
- Computational Cultures: Uncommon Knowledge** at MIT Department of Philosophy. *Ethics across disciplines* (talk and panel).
- Diversity and Inclusion Symposium** by True Blue Inclusion and the Rockefeller Foundation in New York, NY. *Tackling Harm and Improving Accountability in the Automation of Talent Management* (talk and panel).
- 2018 **Systems that Learn @ CSAIL Annual Meeting**. *Bias in Machine Learning and Implications for Healthcare* (talk).
- MIT Science Policy Initiative**. *Combating Bias in Artificial Intelligence* (talk and discussion).

Conference Presentations

- 2022 **FAccT** *Towards Intersectional Feminist and Participatory ML: A Case Study in Supporting Femicide Counterdata Collection* (talk).
- IUI** *Intuitively Assessing ML Model Reliability through Example-Based Explanations and Editing Model Inputs* (talk).
- 2021 **EAAMO** *A Framework for Understanding Sources of Harm throughout the Machine Learning Life Cycle* (talk).
- CHI** *Beyond Expertise and Roles: A Framework to Characterize the Stakeholders of Interpretable Machine Learning and their Needs* (talk).
- 2020 **WebSci** *Measuring the Interference of Machine Learning in Human Decision-Making* (talk).
- MD4SG** *Femicide & Machine Learning: Detecting Gender-based Violence to Strengthen Civil Sector Activism* (talk).
- 2018 **KDD** *Learning Tasks for Multitask Learning: Heterogeneous Patient Populations in the ICU* (talk).
- 2017 **MLHC** *Clinical Event Prediction and Understanding using Neural Networks* (talk).
- WiML workshop @ NeurIPS** *Disparities in the Performance of Natural Language Processing Tools* (poster).
- 2016 **ML4H workshop @ NeurIPS** *The Use of Autoencoders for Discovering Patient Phenotypes* (poster).

Mentoring

Niki Karanikola (Bennington College '23)	<i>Fall 2022 - current</i>
Helen Lu (MIT '22)	<i>Fall 2021 - Spring 2022</i>
Rajiv Movva (MIT '22)	<i>Fall 2021 - Spring 2022</i>
Annie Bryan (MIT '22)	<i>Fall 2021 - Summer 2022</i>
Tiffany Chen (MIT '22)	<i>Fall 2021 - Spring 2022</i>
Angela Zhang (MIT '23)	<i>Fall 2021</i>

Honors and Awards

Best Student Paper Award , ACM Conference on Fairness, Accountability and Transparency (FAccT).	2022
Most Inspiring Paper Award , Mechanism Design for Social Good Workshop (MD4SG).	2020
\$10,000 grant from MIT Sandbox for building an ML Literacy Non-profit Initiative.	2019
MIT Viterbi Fellowship.	2017-2018
Louis Kampff Writing Prize for best essay in Women and Gender Studies.	2016
\$1000 grant from MIT Sandbox for a platform supporting women in tech.	2016