

# Harini Suresh

hsuresh@mit.edu  
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  
GPA: 4.8/5.0

## Research

### Data Driven Inference Group, MIT

*Fall 2017 - Present | Cambridge, MA*

Pursuing a PhD in EECS; studying how people trust automated systems, and how to build models that work with end users effectively. Advised by John Guttag.

### Clinical Decision Making Group, MIT

*Summer 2016 - Summer 2017 | Cambridge, MA*

Integrated multimodal clinical data and used deep learning models to improve the state-of-the-art for intervention prediction in intensive care units and provide interpretability for model decisions (presented at Machine Learning for Healthcare, 2017). Advised by Peter Szolovits.

## Experience

### ML Tidbits Inc.

*Fall 2018 - Present | Cambridge, MA*

- 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
- Writing, filming, illustrating and publishing engaging short videos on YouTube

### Research SWE Intern at Google, Brain Fairness Team

*Summer 2018 | Cambridge, MA*

- Quantified “dataset bias” in large text datasets
- Identified features of a dataset that led to bias in resulting word embeddings and downstream tasks
- Implemented a multipart production system to train several embedding models with different datasets

## Teaching and Leadership

### AI Ethics Reading Group, Lead Organizer

*mitaiethics.github.io | October 2018 - Present*

Co-founded and currently lead an MIT-wide initiative to connect people interested in topics around AI Ethics and facilitate interdisciplinary conversations. Organize bi-weekly meetings with curated readings and discussions.

### Intro to Deep Learning, Lead Organizer

[introtodeeplearning.com/2017](http://introtodeeplearning.com/2017) | January 2017

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.

## Publications

### IN SUBMISSION

**Harini Suresh**, John Guttag. [A Framework For Understanding Sources of “Bias” in Machine Learning.](#)

Natalie Lao\*, **Harini Suresh\***, Ilaria Llicardi. [Measuring the Interference of Machine Learning in Human Decision-Making.](#) (\* = equal contribution)

### CONFERENCE PROCEEDINGS

**Harini Suresh\***, Jen Gong\*, John Guttag. [Learning Tasks for Multitask Learning: Heterogenous Patient Populations in the ICU.](#) *Conference on Knowledge Discovery and Data Mining (KDD) 2018*, London, UK. ACM Conference Proceedings. (\* = equal contribution)

**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) 2017*, Boston MA. JMLR Workshop and Conference Track.

Willie Boag, **Harini Suresh**, Leo Celi, Peter Szolovits and Marzyeh Ghassemi. [Racial Disparities and Mistrust in End-of-Life Care.](#) *Machine Learning for Healthcare Conference 2018*, Stanford CA. JMLR Workshop and Conference Track.

### WORKSHOPS AND POSTERS

Willie Boag, **Harini Suresh**, Leo Celi, Peter Szolovits and Marzyeh Ghassemi. [Modelling Mistrust in End-of-Life Care.](#) *Fairness, Accountability, and Transparency in Machine Learning Workshop*, ICML 2018, Stockholm.

**Harini Suresh\***, Divya Shanmugam\*, John Guttag. [Disparities in the Performance of Natural Language Processing Tools.](#) *Women in Machine Learning (WiML) Workshop*, NIPS 2017, Palm Springs CA. (\* = equal contribution)

**Harini Suresh**, Peter Szolovits, Marzyeh Ghassemi. [The Use of Autoencoders for Discovering Patient Phenotypes.](#) *Machine Learning for Healthcare Workshop*, NIPS 2016, Barcelona.

## Invited Talks and Panels

**MIT Better World symposium** in Atlanta, GA (October 2019). *Trust Issues in Machine Learning* (talk and panel).

**Fair ML in Health** at Data & Society Research Institute in New York, NY (October 2019). *Deploying decision-aids: real-world considerations* (talk).

**Computational Cultures: Uncommon Knowledge** at MIT Department of Philosophy (May 2019). *Ethics across disciplines* (lightning talk and panel).

**Diversity and Inclusion Symposium** by True Blue Inclusion in New York, NY (May 2019). *Tackling Harm and Improving Accountability in the Automation of Talent Management* (talk and panel).

**Systems that Learn @ CSAIL Annual Meeting** (August 2018). *Bias in Machine Learning and Applications to Healthcare* (talk).

**Machine Learning for Healthcare Conference** in Boston, MA (August 2017). *Clinical Event Prediction and Understanding using Neural Networks* (lightning talk).

## Projects

### Microsoft's Bias & Discrimination Hackathons

May 2018: Designed a dashboard for more transparent automated decisions. 3rd place team.

([algbias.weebly.com](http://algbias.weebly.com))

June 2017: Built interface to understand the tradeoffs between different frameworks of fairness, with the goal of making them easier to understand and adopt for key users. Winning team of the hackathon.

([challengethebias.github.io](http://challengethebias.github.io))

Feb 2017: Made a tool to help understand bias in word embeddings by visualizing genres of words across the "gender subspace" of a word embedding ([wordbias.org](http://wordbias.org))

### Commonwealth Stats - [massnonprofitnet.herokuapp.com](http://massnonprofitnet.herokuapp.com)

An application for non-profit organizations to access and analyze important demographic data from the US Census Bureau in an easy and visual way, for use in grants, reports, or research. Made for the Massachusetts Nonprofit Network.