# EECS RISING STARS 2018

**VIVIENNE SZE**
**Associate Professor of EECS, MIT**
Rising Stars, 2012

OCTOBER 28-30, 2018

"**Welcome to MIT!** EECS is delighted to host such an outstanding group of early-career electrical engineers and computer scientists. We've designed an intensive workshop that will allow you to **share your work, hear from leading professors and researchers,** and **develop a strong peer network** for career support and collaboration long beyond your time on campus."

**ASU OZDAGLAR**
**Department Head | MIT Department of Electrical Engineering and Computer Science**
**MIT School of Engineering Distinguished Professor of Engineering**

"The **2018 Rising Stars workshop** is once again bringing together some of the world's most talented women working in the fields of electrical engineering and computer science. The program will provide you with **great opportunities to present your research, meet top faculty, and network with your peers.** Ultimately, we believe you'll form connections that will persist throughout your careers."

**ANANTHA CHANDRAKASAN**
**Dean | MIT School of Engineering**
**Vannevar Bush Professor of Electrical Engineering and Computer Science**

**COVER PHOTO OF VIVIENNE SZE:**
Lillie Paquette, MIT School of Engineering

FROM THE 2018

# WORKSHOP CHAIRS

## WELCOME TO THE 2018 RISING STARS IN EECS WORKSHOP AT MIT!

We launched Rising Stars in 2012 to identify and mentor outstanding women electrical engineers and computer scientists interested in exploring careers in academia. We are delighted that the program has grown substantially since its beginning. This year's workshop will bring together more than 70 of the world's brightest women in EECS for two days of scientific interactions and discussions about navigating the early stages of careers in academia.

Our program includes invited presentations on building research networks, teaching and mentoring students, funding research groups, and applying for jobs, among other critical topics. Panel discussions will focus on career trajectories, search committees, the promotion process, research labs, and more.

Workshop participants will also give presentations covering a wide range of specialties representative of the breadth of EECS research. These presentations will range from materials, devices, and circuits to signal processing, communications, computer science theory, artificial intelligence, and systems.

Past Rising Stars attendees have gone on to secure faculty positions at top universities or research positions in leading industry labs. Toward this end, we are pleased to highlight and feature workshop participants in this guide and on the Rising Stars website.

In addition, we hope that Rising Stars will give participants a chance to network with their peers, opening the door for ongoing collaboration and professional support for years to come.

We'd like to thank advisors who have supported current and past Rising Stars participants. We're also grateful for the support of MIT's School of Engineering and the four EECS-affiliated research labs: the Computer Science and Artificial Intelligence Laboratory (CSAIL), the Laboratory for Decision and Information Systems (LIDS), the Microsystems Technology Laboratories (MTL), and the Research Laboratory of Electronics (RLE). Finally, we thank the Rising Stars administrative team for all their assistance.

We look forward to meeting and interacting with you all.

### Anantha Chandrakasan

WORKSHOP ADVISOR

Dean | MIT School of Engineering
Vannevar Bush Professor of Electrical Engineering and Computer Science

### Asu Ozdaglar

WORKSHOP CHAIR

Department Head | MIT Department of Electrical Engineering and Computer Science
MIT School of Engineering Distinguished Professor of Engineering

### Stefanie Mueller

WORKSHOP TECHNICAL CO-CHAIR

X-Consortium Career Development Assistant Professor of Electrical Engineering and Computer Science, MIT

### Farnaz Niroui

WORKSHOP TECHNICAL CO-CHAIR

Assistant Professor of Electrical Engineering and Computer Science, MIT

### Vivienne Sze

WORKSHOP TECHNICAL CO-CHAIR

Associate Professor of Electrical Engineering and Computer Science, MIT

### Virginia Williams

WORKSHOP TECHNICAL CO-CHAIR

Steven G. (1968) and Renee Finn Career Development Associate Professor of Electrical Engineering and Computer Science, MIT

# 2018
## RISING STARS IN EECS

**YOMNA ABDELRAHMAN**
University of Stuttgart

**SANGEETHA ABDU JYOTHI**
University of Illinois at Urbana-Champaign

**AISHWARYA AGRAWAL**
Georgia Institute of Technology

**FATEMEH AKBAR**
University of Michigan

**MARYAM ALIAKBARPOUR**
Massachusetts Institute of Technology

**BEHNAZ ARZANI**
Microsoft Research

**SUGUMAN BANSAL**
Rice University

**ZOYA BYLINSKII**
Adobe Research

**SARAH CANNON**
University of California, Berkeley

**MARGARITA PAZ CASTRO**
University of Toronto

**STEVIE CHANCELLOR**
Georgia Institute of Technology

**SARAH CHASINS**
University of California, Berkeley

**ANGELA DAI**
Stanford University

**BITA DARVISH ROUHANI**
University of California, San Diego

**ORIANNA DEMASI**
University of California, Berkeley

**TAYLA EDEN**
Tel Aviv University

**AISHWARYA GANESAN**
University of Wisconsin–Madison

**CHAYA GANESH**
Aarhus University

**ESHA GHOSH**
Microsoft Research

**ALYSSON REBECCA GOLD**
Stanford University

**HANNAH GOMMERSTADT**
Carnegie Mellon University

**DIANA CRISTINA GONZALEZ GONZALEZ**
Massachusetts Institute of Technology

**NIL ZEYNEP GUREL**
Georgia Institute of Technology

**SEPIDEH HASSAN MOGHADDAM**
University of Southern California

**ELIZABETH FARRELL HELBLING**
Harvard University

**QIUYUAN HUANG**
Microsoft Research

**SANDY HUANG**
University of California, Berkeley

**RUPAMATHI JADDIVADA**
Massachusetts Institute of Technology

**AKANKSHA JAIN**
University of Texas, Austin

**SEYEDEH MAHSA KAMALI**
California Institute of Technology

**ANGJOO KANAZAWA**
University of California, Berkeley

**MANDY BARRETT KORPUSIK**
Massachusetts Institute of Technology

**PAN LI**
University of Washington

**QI LI**
University of Illinois at Urbana-Champaign

**MENGXIA LIU**
University of Toronto

**SEVIL ZEYNEP LULEC**
University of Toronto

**DERYA MALAK**
Massachusetts Institute of Technology

**ZELDA MARIET**
Massachusetts Institute of Technology

**NEGAR MEHR**
University of California, Berkeley

**AASTHA KETAN MEHTA**
Max Planck Institute for Software Systems

**BAHARAN MIRZASOLEIMAN**
Stanford University

**NASIM MOHAMMADI ESTAKHRI**
University of Pennsylvania

**RAJALAKSHMI NANDAKUMAR**
University of Washington

**SHADI NOGHABI**
University of Illinois at Urbana-Champaign

**JELENA NOTAROS**
Massachusetts Institute of Technology

**SHIJIA PAN**
Carnegie Mellon University

**EMMA PIERSON**
Stanford University

**NEELAM PRABHU GAUNKAR**
Iowa State University

**MAITHRA RAGHU**
Cornell University and Google Brain

**NIRANJINI RAJAGOPAL**
Carnegie Mellon University

**ALENA RODIONOVA**
University of Pennsylvania

**RACHEL RUDINGER**
Johns Hopkins University

**EMMA R. SCHMIDGALL**
University of Washington

**ALEXANDRA KATHRYN SCHOFIELD**
Cornell University

**SIVARANJANI SEETHARAMAN**
University of Notre Dame

**YUANYUAN SHI**
University of Washington

**MISHA SRA**
Massachusetts Institute of Technology

**LILI SU**
Massachusetts Institute of Technology

**WENJING SU**
Google

**YIHAN SUN**
Carnegie Mellon University

**YIXIN SUN**
Princeton University

**MINA TAHMASBI ARASHLOO**
Princeton University

**CAROLINE JUNE TRIPPEL**
Princeton University

**ELLEN VITERCIK**
Carnegie Mellon University

**CHING-HUA WANG**
Stanford University

**ZI WANG**
Massachusetts Institute of Technology

**FAN WEI**
Stanford University

**NORA S. WILLETT**
Princeton University

**ANQI WU**
Princeton University

**HUIJUAN XU**
Boston University

**MENGJIA YAN**
University of Illinois at Urbana-Champaign

**HAZAL YUKSEL**
Columbia University

**AMY XIAN ZHANG**
Massachusetts Institute of Technology

**SHANGHANG ZHANG**
Carnegie Mellon University

**YIYING I. ZHU**
Georgia Institute of Technology

**MARINKA ZITNIK**
Stanford University

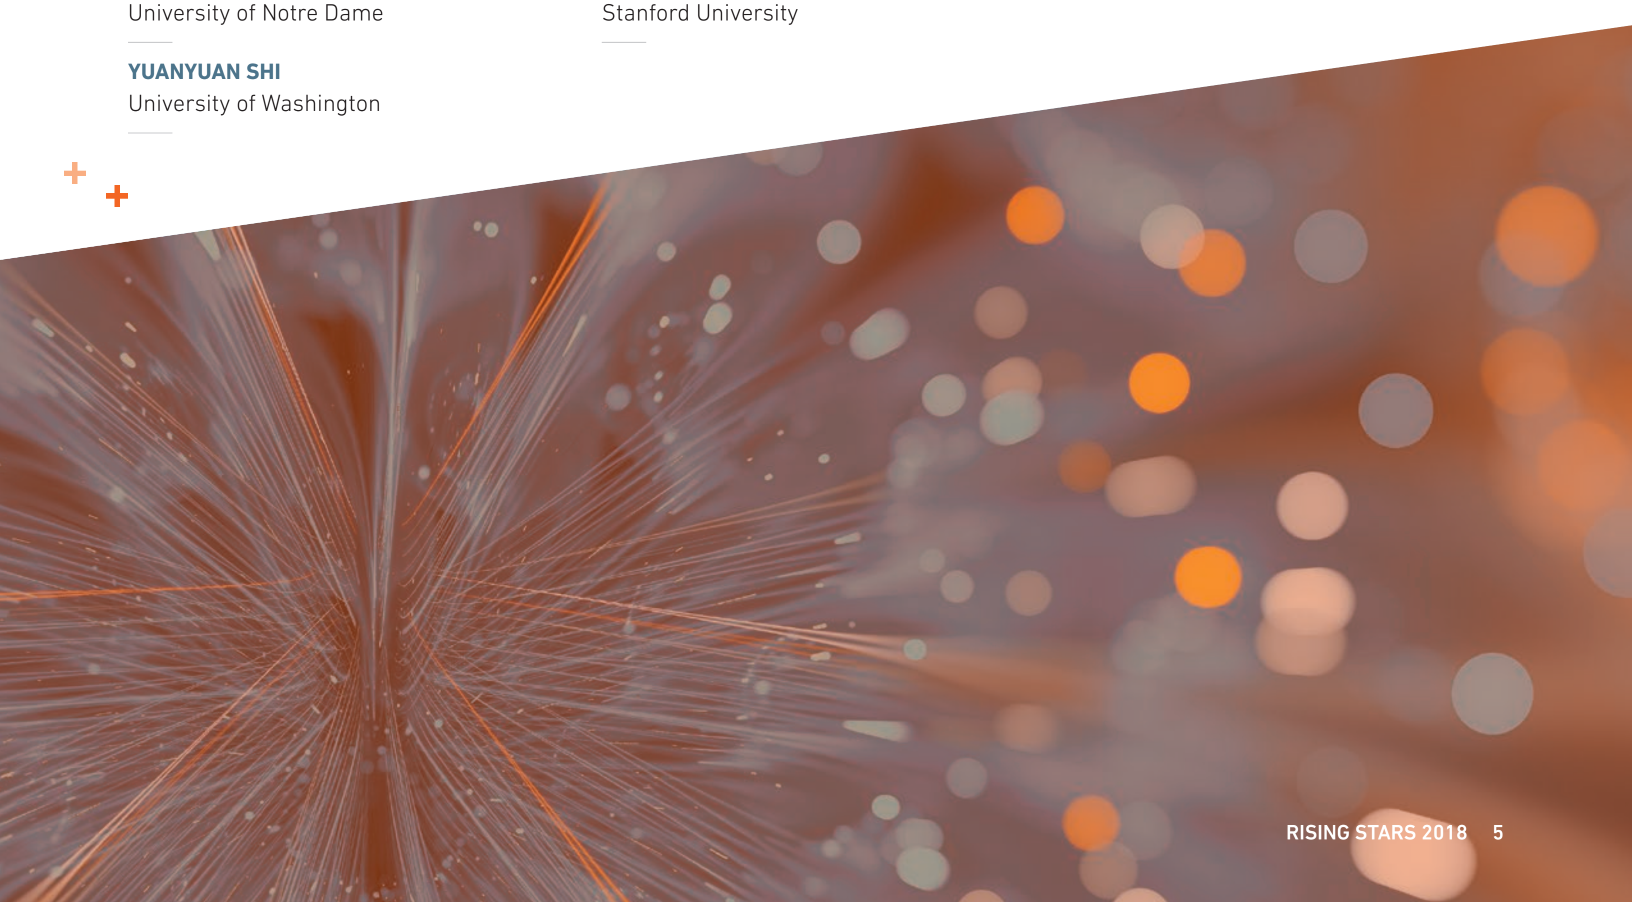


## Yomna Abdelrahman

**PHD CANDIDATE,** University of Stuttgart



## Sangeetha Abdu Jyothi

**PHD STUDENT,** University of Illinois at Urbana-Champaign


### Thermal Imaging in HCI

Thermal cameras have recently drawn the attention of HCI researchers as a new sensory system enabling novel interactive systems. They are robust to illumination changes and make it easy to separate objects from the scene background. Far-infrared radiation, however, has another characteristic that distinguishes thermal cameras from their RGB or depth counterparts as it operates in the non-visual spectrum. On the other hand, the visual spectrum, i.e., human visual perception, is limited to only 1 percent of the electromagnetic spectrum. Research has shown that extending visual perception can be beneficial, but it is unclear if this is useful for a broader range of applications in daily setups by novice users. To investigate the potential of domestication and adoption of thermal imaging, we are conducting a set of explorative studies in different forms, including: user/lab studies, interviews, proof-of-concept prototype technology, and cultural probe. Throughout the conducted research, we collected various thermal data in different formats, such as videos, photos, and raw temperature representation. Our findings reflected the potential of thermal imaging domestication and the adoption of such a technology by novice users. Additionally, we found that the users were excited about using thermal cameras in their everyday lives and found many practical uses for them. Our studies provide insights into how novice users wish to use thermal imaging technology to augment their vision in daily setups as well as identify and classify common thermal imaging use cases. Our work contributes implications for designing thermal imaging devices targeted towards novice users.

#### BIOGRAPHY

Yomna Abdelrahman is a PhD candidate in the hciLab at the University of Stuttgart, supervised by Professor Dr. Albrecht Schmidt. During her PhD studies, she explored thermal imaging to build novel interactive systems. Her work has been published in conference papers and journals such as CHI and IMWUT. Her research was recognized by the community and received a best paper award at MUM2015 and honorable mention at CHI2017. She was awarded the Melbourne School of Engineering Fellows Grant during her SocialNUI Microsoft internship. She did her master's thesis in the hciLab with the work titled "Thermal Imaging for Interactive Systems," which was published in CHI2014. Additionally, during her undergraduate studies at the German University in Cairo, she was part of the team of RoboCup 2010, which won first place in the logistics league.

### Automated Resource Management in Large-Scale Networked Systems

Data centers (DCs) constitute a critical component in today's Internet infrastructure with most applications relying on DCs partly or wholly. DCs are typically multi-tenanted — the servers and the interconnecting network are shared across multiple users. In this environment, the goals of the various stakeholders are diverse. The objective of the provider is to increase revenue by utilizing the resources effectively. The applications on the other hand have a variety of performance requirements and time-varying demands. My research is centered on automated, closed-loop control in large-scale networked systems to simultaneously satisfy the utilization requirements of the provider and varied performance requirements of applications with dynamic loads. Towards this goal, I have worked on (1) estimating the bounds of performance achievable in a given environment and (2) designing and building automated systems to utilize the resources efficiently. We built Morpheus, a system for cloud resource management, which is currently deployed in Microsoft and open-sourced in Hadoop 2.9. Morpheus achieves high resource utilization in enterprise clusters using three main components: a learning module that accurately learns users' requirements, a reservation module that schedules jobs based on the estimated requirements, and a dynamic reprovisioning module that adapts the allocation in real-time. I also extend these ideas to other large-scale systems with distinct challenges: (i) geo-distributed micro data centers deployed by cellular providers supporting applications with stringent latency constraints and (ii) Deep Neural Network (DNN) frameworks with iterative network-heavy workloads.

#### BIOGRAPHY

Sangeetha Abdu Jyothi is a PhD candidate at UIUC, where she is advised by Brighten Godfrey. Her research interests lie in the areas of computer networking and systems with a focus on automated resource management in large-scale networked systems. She received the Facebook Graduate Fellowship (2017–2019) and the Mavis Future Faculty Fellowship (2017–2018). She was a research intern at Microsoft Research during the summer of 2015 and a software engineering intern at Google during the summer of 2014. Prior to UIUC, she received a master's degree from the University of Pennsylvania in 2012. She received a bachelor's degree from the National Institute of Technology, Calicut, India, in 2010 where she was awarded the gold medal for being the top graduating student.

## Aishwarya Agrawal

**PHD CANDIDATE,**
Georgia Institute of Technology

## Fatemeh Akbar

**PHD CANDIDATE, RESEARCH ASSISTANT,** University of Michigan

### Towards Intelligent Vision and Language Systems

My research goal is to develop artificial intelligence (AI) systems that can "see" (i.e., understand the contents of an image: who, what, where, doing what?) and "talk" (i.e., communicate the understanding to humans in free-form natural language). Applications of such vision and language systems include: (1) Aiding visually impaired users in understanding their surroundings (Human: "What is on the shelf above the microwave?" AI: "Canned containers"); (2) Aiding analysts in making decisions based on large quantities of surveillance data (Human: "What kind of car did the man in red shirt leave in?" AI: "Blue Toyota Prius"); (3) Teaching children through interactive demos. (Kid: "What animal is that?" AI: "That is a Dall sheep. You can find those in Alaska"); (4) Interacting with personal AI assistants, such as Alexa or Siri (Human: "Is my laptop in my bedroom upstairs?" AI: "Yes." Human: "Is the charger plugged in?" AI: "Yes"); and (5) Making visual social media content more accessible (AI: "Your friend Bob just uploaded a picture from his Hawaii trip." Human: "Great, is he at the beach?" AI: "No, on a mountain"). As a first step towards making intelligent vision and language systems in my PhD research so far, I have worked on building datasets, models, and evaluation protocols for answering free-form and open-ended natural language questions about images. I will talk about the Visual Question Answering (VQA) task dataset baseline neural models (Antol et al. ICCV15), various VQA challenges we have organized until now and findings from those challenges, our findings from analyzing the behavior of VQA models (Agrawal et al. EMNLP16), and how we propose to build VQA models that are more visually grounded and can better deal with VQA under changing priors (Agrawal et al. CVPR18).

#### BIOGRAPHY

Aishwarya Agrawal is a fifth-year PhD student in the School of Interactive Computing at Georgia Tech, advised by Dhruv Batra. She received a bachelor's degree in electrical engineering with a minor in computer science and engineering from the Indian Institute of Technology (IIT) Gandhinagar in 2014. Her research interests lie at the intersection of computer vision, machine learning, and natural language processing with a focus on developing AI systems that that can "see" and "talk." She received an NVIDIA Graduate Fellowship for 2018-2019.

### Novel Architectures for Low-Complexity Scalable Phased Arrays

Inspired by the unique advantages of phased arrays in communication and radar systems such as increasing the channel capacity signal-to-noise ratio directivity and radar resolution, my research has focused on presenting new architectures for low-complexity scalable phased arrays to facilitate their widespread use in communication and radar systems. In phased arrays, phase shifters are the key components responsible for adjusting the signal phase across the array elements. In general, phase shifters and their control circuitry play a significant role in determining the complexity and size of conventional phased arrays. Novel architectures and design techniques for scalable phased arrays with significantly reduced number of phase shifters, control complexity, and circuit size have been devised and will be presented. Design approach and performance of a limited-scan Ku-band and a wide-scan K-band scalable phased array where the number of phase shifters are reduced by at least a factor of two will be presented. The integrated phased arrays designed based on the proposed architectures have a potential to be utilized in commercial applications such as 5G communications and automotive radars for advanced driver assistance systems (ADAS) and autonomous vehicles. I will also present new circuit topologies for integrated phase shifters operating at K and Ka bands with low power consumption, compact size, and simple control mechanisms designed for reducing the complexity size and power consumption of phased arrays. Future work building upon the proposed low-complexity scalable phased arrays will also be presented.

#### BIOGRAPHY

Fatemeh Akbar is a PhD candidate in the Electrical Engineering and Computer Science Department at the University of Michigan. She received bachelor's and master's degrees (with highest honors) in electrical engineering from Shahid Beheshti University in 2010 and Sharif University of Technology in 2012. The focus of her current research is on RF/mm-wave integrated phased arrays. Her research interests include analog mixed-signal RF microwave and mm-wave integrated circuits and systems for applications in wireless transceivers, radars, sensors, imaging, biomedical devices, and photonics. She was a recipient of the Engineering Graduate Symposium Technical Award of the University of Michigan College of Engineering in 2014, 2016, and 2017. She is a full member of the Sigma Xi, IEEE Solid-State Circuits, IEEE Circuits and Systems, and IEEE Microwave Theory and Techniques societies.

## Maryam Aliakbarpour

**PHD STUDENT,** Massachusetts Institute of Technology

### Differentially Private Identity and Equivalence Testing of Discrete Distributions

We study the fundamental problems of identity and equivalence testing over a discrete population from random samples. Our goal is to develop efficient testers while guaranteeing differential privacy to the individuals of the population. We provide sample-efficient differentially private testers for these problems. Our theoretical results significantly improve over the best known algorithms for identity testing and are the first results for private equivalence testing. The conceptual message of our work is that there exist private hypothesis testers that are nearly as sample-efficient as their non-private counterparts. We perform an experimental evaluation of our algorithms on synthetic data. Our experiments illustrate that our private testers achieve small type I and type II errors with sample size sublinear in the domain size of the underlying distributions. Specifically, the sample complexity of our private identity tester significantly outperforms the sample complexity of recently proposed methods for this problem.

#### BIOGRAPHY

Maryam Aliakbarpour is a PhD student in the Department of Electrical Engineering and Computer Science (EECS) at MIT, conducting research under the supervision of Professor Ronitt Rubinfeld. She received a master's degree in EECS from MIT and a bachelor's degree in computer engineering from Sharif University of Technology. Her research interest is in sublinear algorithm learning theory and property theory and property testing, in particular property testing of distributions. She conducted several research projects on the topic, including differentially private testing of distribution testing and learning junta distributions incrementally improving visualizations of datasets. She received the Neekeyfar Fund Award from the MIT Office of the Dean for Graduate Education in 2013. She was also awarded the silver medal in the Iranian National Olympiad in Informatics 2008.

## Behnaz Arzani

**POSTDOCTORAL RESEARCHER,** Microsoft Research

### Towards Self-Managing and Performant Networks

Today, the vast majority of the applications we use on a day-to-day basis operate by relying on computer networks. For this reason, any performance degradation on these networks can result in significant impact on not just user experience but the revenues of the service providers (e.g., Google, Azure, AWS) and the services themselves. My research focuses on improving user quality of experience (QoE) in data center and wide area networks. To this end, I have looked at many different aspects of networking and systems research, including traffic engineering diagnosis, resilience to failure, and security. In my future research, I plan to work on creating self-managed networks. New machine-learning techniques and new programmable data-planes makes the time ripe for moving toward such networks. In theory, a self-managed network would identify where and when to execute management tasks and identify problems when and where they occur, then automatically resolve them. This includes reconfiguring the network when necessary, predicting performance problems and mitigating them, and finding link failures, security breaches, and even congestion events. As a postdoc, I am exploring problems that set a foundation useful for solving this overarching problem. My intern, Robert MacDavid of Princeton University, and I are looking at improving ticket resolution in Microsoft data centers and the security space, finding compromised VMs without visibility into customer VMs.

#### BIOGRAPHY

Behnaz Arzani is a postdoctoral researcher at Microsoft Research in the Mobility and Networking group. She received a PhD in computer science from the University of Pennsylvania in 2017, where she was advised by Professors Boon Thau Loo and Roch Guerin. She completed a dual master's degree in computer science and electrical engineering at the University of Pennsylvania in the same year. She received a bachelor's degree in electrical engineering at Sharif University of Technology, where she worked with Professor Javad Salehi. She recently received the University of Pennsylvania Rubinoff Award for the best dissertation in computer science.


## Suguman Bansal

**PHD CANDIDATE,**
Rice University



## Zoya Bylinskii

**RESEARCH SCIENTIST,**
Adobe Research


# Scalable Quantitative Verification and Its Applications

On one hand, our increasing dependence on computerized systems emphasizes the importance of systems correctness. On the other hand, their increasing complexity makes correctness harder to achieve. The formal reasoning about complex quantitative properties of these systems, such as cost reward, runtime, and the like, involves automated techniques to ensure system correctness. Unsurprisingly, it has real-world applications in verification of decentralized cryptocurrency protocols, network-track monitoring, optimal-cost motion-planning in robotics, and more. So far, however, the formal reasoning of quantitative properties has been addressed only from a theoretical point of view and not with practical considerations. The goal of my dissertation is to develop new techniques to bridge the gap between theory and practice of quantitative reasoning to harness its untapped potential in real-world applications. At the core of reasoning about quantitative properties of systems lies the fundamental problem of comparison of the quantities associated with system runs. In my research, I have developed an automata-theoretic technique for the comparison of quantities of system runs and have utilized it to provide the first and, to the best of our knowledge, only automatically generated proof, as opposed to a manually generated proof of the non-incentive compatibility of the Bitcoin protocol. I have also applied my automata-theoretic comparison technique in quantitative model-checking to design resource-efficient systems and to establish new theoretical complexity bounds, specifically the PSPACE-completeness of discounted-sum inclusion with integer discount-factors.

### BIOGRAPHY

Suguman Bansal is a PhD candidate in the Department of Computer Science at Rice University. Her research interests span formal methods, specifically the application of quantitative verification and reactive synthesis to real-time decision making in multi-agent systems. She received the Andrew Ladd Fellowship and Gold Medal at the Association for Computing Machinery (ACM) Student Research Competition at POPL 2016. She was also a visiting graduate student in spring 2018 at the "Real-Time Decision Making" program at the Simons Institute for the Theory of Computing at UC Berkeley and spent the summers of 2017 and 2018 interning at Nokia Bell Labs.

# Computational Perception for Multimodal Document Understanding

Multi-modal documents occur in a variety of forms as graphs in technical reports, diagrams in textbooks, and graphic designs in bulletins. Humans can efficiently process the visual and textual information contained within to make decisions on topics including business, health care, and science. Building the computational tools to understand multi-modal documents can therefore have important applications for web search, information retrieval, automatic captioning, and design tools. My research has used machine learning for detecting and parsing the visual and textual elements in multi-modal documents for topic prediction and automatic summarization. Inspired by human perception, I have developed novel data collection methods and built models that predict where people look in graphic designs and information visualizations. These predictions have enabled interactive design applications. My work has made contributions to the fields of human vision, computer vision, and human-computer interaction.

### BIOGRAPHY

Zoya Bylinskii recently started as a research scientist at Adobe Research in Cambridge, Massachusetts. She received a PhD from MIT in September 2018, a master's degree in electrical engineering and computer science from MIT in 2015, and an honors bachelor's degree in computer science and statistics from the University of Toronto in 2012. She has been a 2016 Adobe Research Fellow, a 2014–2016 NSERC Postgraduate Scholar, and a 2013 Julie Payette Research Scholar; she also received a 2011 Google Anita Borg Memorial Scholarship. She works at the interface of human vision, computer vision, and human-computer interaction: building computational models of people's memory and attention and applying the findings to graphic designs and data visualization with automatic summarization and redesign applications (AI for creativity).


## Sarah Cannon

**POSTDOCTORAL RESEARCH FELLOW,** University of California, Berkeley



## Margarita Paz Castro

**PHD STUDENT,** University of Toronto


### Markov Chains, Programmable Matter, and Emergent Behavior

How can collections of simple computational elements accomplish complicated system-level goals? This question motivates the study of programmable matter, swarm robotics, and even biological systems such as ant colonies. In simplified settings, we've shown that simple collections of abstract particles executing our local stochastic algorithms can accomplish remarkable global objectives, including the fundamental behaviors of compression separation and alignment. In several of our algorithms, changing just a single parameter results in a different, but equally desirable, emergent global behavior. Because our distributed algorithms come from Markov chains, we can use tools and insights from Markov chain analysis to predict and rigorously explain their behavior.

#### BIOGRAPHY

Sarah Cannon is a National Science Foundation (NSF) Mathematical Sciences Postdoctoral Research Fellow at UC Berkeley. She recently completed a PhD in algorithms, combinatorics, and optimization at the Georgia Institute of Technology, where she was an NSF Graduate Research Fellow and a Clare Boothe Luce Outstanding Graduate Fellow. She also received a Simons Award for Graduate Students in Theoretical Computer Science. Previously, she earned a master's degree in mathematics and the foundations of computer science from the University of Oxford in 2013 and a bachelor's degree in mathematics from Tufts University in 2012.

### Relaxed Decision Diagrams for Discrete Optimization Problems

Discrete optimization is one of the main research fields in operational research and computer science due to its challenging problems and many applications. Several techniques have been used to solve these problems such as Integer programming, constraints programming, and heuristic search. Relaxed Decision Diagrams (DDs) are one of the newest approaches to tackle discrete optimization problems, which use a graphical structure to represent the set of feasible solutions and compute bounds. My research aims to investigate and extend the use of relaxed DDs as a tool to solve discrete optimization problems across different fields. I want to uncover the advantages of the approach, its relationship to existing techniques, and its uses in conjunction with other approaches such as integer programming. I am currently working with challenging (NP-hard) problems from the operations research and artificial intelligence (AI) communities including vehicle routing, scheduling, and AI planning problems.

#### BIOGRAPHY

Margarita Castro is a fourth-year PhD student in the Mechanical and Industrial Engineering Department at the University of Toronto. Margarita is part of the Toronto Intelligent Decision Engineering Laboratory (TIDEL) and is under the supervision of Professors J. Christopher Beck and Andre Cire. Margarita received a full PhD scholarship from the Chilean government. She received a master's degree in 2014 from the Pontificia Universidad Catolica de Chile, where she graduated with maximum distinction. Her research focuses on discrete optimization problems and the use of relaxed decision diagrams to solve them. She is interested in constraint programming, integer programming, stochastic programming, AI applications, machine learning, scheduling, and vehicle routing problems.

## Stevie Chancellor

**PHD CANDIDATE,**
Georgia Institute of Technology

## Sarah Chasins

**PHD CANDIDATE,** University of California, Berkeley

### Computational Methods to Understand Deviant Health Behaviors in Online Communities

Online groups provide advice and a sense of community for health and wellness, especially for stigmatized mental health disorders. However, individuals also turn to online communities to promote deliberate self-injury, disordered eating habits, and suicidal ideation. These behaviors can have contagion-like effects on those outside of these communities as well as impacts on social networks that struggle with managing such dangerous content. These behaviors are "deviant"—actions that violate the social norms and behaviors of a particular community. Deviant behaviors that promote self-injury violate platform policies and social expectations that individuals do not harm themselves. To precisely identify and manage these pernicious issues is challenging, as is designing appropriate interventions. I study deviant behavior in online communities using computational techniques on large-scale social data. My area of focus is deviant mental-health behaviors, such as pro-eating disorder. I draw from machine learning and computational linguistics to analyze social media datasets to understand patterns of behavior. Taking a human-centered approach to this complex topic, I explore issues of content moderation/management and ethics in managing these behaviors. I integrate these computationally robust techniques with human knowledge and domain expertise alongside ethics to design "human-in-the-loop" systems for answering our toughest questions about deviant behavior online. My dissertation research focuses on pro-eating disorder communities, a clandestine group that glorifies eating disorders. Previously, I analyzed patterns of elaborate yet semantically meaningful lexical variants to evade content bans on these communities on Instagram, such as changing tags from "thighgap" to "thyghgappp." I also developed a new model of generalized mental-illness severity on online data that predicts future severity levels up to eight months in advance.

#### BIOGRAPHY

Stevie Chancellor is a PhD candidate in Human Centered Computing in Interactive Computing at Georgia Tech. Her research interest lies in quantitative social computing and human-computer interaction (HCI). Specifically, she studies computational methods to understand deviant mental-health behavior in online communities and combines techniques from natural language processing, machine learning, and data science. Her work has won several best paper honorable mention awards at CHI and CSCW, premiere venues in human computer interaction. She received a bachelor's degree from the University of Virginia in 2012 and a master's degree from Georgetown University in 2014. Her work is supported by a Snap Inc. Research Fellowship and has appeared in national publications such as Wired and Gizmodo.

### Helena: A Web Automation Language for End Users

Web data is becoming increasingly important for data scientists. Social scientists in particular envision a diverse range of applications and insights driven by web data — for example, tracking rents in real time by collecting online apartment listings or studying how charitable foundations communicate with supporters by collecting social media posts. With the wide variety of web scripting libraries on offer, programmers have access to increasing language support for web scraping tasks. However, these libraries are inaccessible to non-programmers, and empowering non-programmers to complete these same tasks is a long-standing open problem. To solve this problem, we designed the Helena web automation language. Helena brings together the following key innovations, which together empower end users to write robust web scraping programs: (1) The Helena programming environment uses Programming by Demonstration (PBD), which makes scripts easy to write; the tool takes a single-shot learning approach, creating scripts based on recording a single interaction of the user with a set of webpages. Empirically, users can learn the tool and use it to write a robust large-scale scraping script in less than 10 minutes, while programmers tackling the same task with the traditional Selenium language time out after an hour. (2) Helena's adaptive replayer makes scripts robust to webpage redesigns and obfuscation, which enables longitudinal experiments. (3) Helena's novel runtime can parallelize and distribute scraping programs for speedups over 50x, facilitating large-scale scraping. Our approach relied on novel insights into the web scraping domain but also on bringing new techniques to bear. By combining techniques from the programming languages community and the human-computer Interaction (HCI) community, we arrived at a language design that meets real users' needs.

#### BIOGRAPHY

Sarah Chasins is a PhD candidate at UC Berkeley, advised by Ras Bodik. Her research interests lie at the intersection of programming languages and HCI. She works on end-user programming program synthesis and programming language design. Much of her work is shaped by ongoing collaborations with social scientists from fields ranging from sociology to economics to public policy. She believes well-designed languages and programming environments can put complicated programming tasks in range for people who consider themselves non-coders. She has received a National Science Foundation Graduate Research Fellowship and a first-place award in the Association for Computing Machinery (ACM) Student Research Competition.

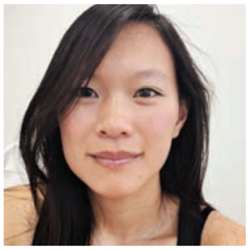


## Angela Dai

**PHD STUDENT,** Stanford University


### Using Generative Deep Learning to Create High-Quality Models from 3D Scans

With recent developments in both commodity range sensors and mixed reality devices, capturing and creating 3D models of the world around us has become increasingly important. Such 3D models will not only facilitate, capture, and display for content creation, but will also provide a basis for fundamental scene understanding from semantic understanding to virtual interactions, which must be formulated in 3D for many applications such as augmented or virtual reality. Leveraging data from commodity range sensors to reconstruct 3D scans of a scene has shown significant promise towards 3D model creation of real-world environments. However, the quality of reconstructed 3D scans has yet to reach that of artist-created 3D models—in particular, 3D scans always suffer from incompleteness due to occlusions in real-world scenes as well as physical limitations of range sensors. Such incomplete 3D models are both unsuitable visually and, moreover, provide only a limited basis for higher-level scene reasoning (for example, virtual interactions will not be accurate in unknown or missing regions). My work introduces a generative formulation for the task of scan completion, leveraging deep learning techniques to create high-quality complete models from 3D scans. We approach this problem as a conditional generative task; conditioned on an input partial scan, we aim to learn "part"-wise similarity between scans to infer the complete model. We begin by focusing on the more constrained problem of completing scans of isolated shapes. We then expand upon this to design a generative approach for completion of general 3D scans, directly addressing the challenge of varying scene sizes in 3D. This not only provides scan completion at scale producing geometrically complete 3D models but also provides a basis for higher-level scene reasoning such as that required for virtual interactions or physical simulations.

#### BIOGRAPHY

Angela Dai is a PhD student in computer science at Stanford University, advised by Pat Hanrahan. Her research focuses on 3D reconstruction and understanding with commodity sensors. She received a master's degree from Stanford University and a bachelor's degree from Princeton University, and she has received a Stanford Graduate Fellowship.


## Bita Darvish Rouhani

**PHD CANDIDATE,** University of California, San Diego


### Succinct and Assured Machine Learning: Training and Execution

Physical viability and safety consideration are key factors in devising machine learning systems that are both sustainable and trustworthy. Learning and analysis of massive data is a trend that is ubiquitous among different computing platforms ranging from smartphones and Internet-of-Things (IoT) devices to personal computers and many-core cloud servers. Concerns over the functionality (accuracy), physical performance, and reliability (safety) are major challenges in building automated learning systems that can be employed in real-world settings. My research work aims to address these three critical aspects of emerging computing scenarios: the functionality, physical performance, and reliability. What makes my work distinctive is that it provides holistic automated solutions that simultaneously capture the best of the computer architecture, machine learning (ML), and security fields. I introduced, realized, and automated a resource-aware deep learning (DL) framework (called Deep3) that enables succinct training and execution of DL models by simultaneously leveraging three levels of parallelism: data, network, and hardware. The resource efficiency of my solution enables, for the first time, dynamic training and execution of DL networks in resource constraints settings applied to both distributed cloud servers with limited communication bandwidth and embedded devices with bounded computational power (for example, smartphones). I further enriched my proposed DL framework by providing algorithmic and custom hardware-accelerated tools to ensure the reliability of model prediction in the face of adversarial attacks. The reliability aspect of my developed solution in turn empowers the use of assured DL models in sensitive scenarios where data privacy and/or robustness against adversarial samples is crucial.

#### BIOGRAPHY

Bita Darvish Rouhani is a last-year PhD student at UC San Diego with an expected degree in electrical and computer engineering. She received a master's degree in computer engineering from Rice University. She is a recipient of a Microsoft PhD Fellowship. Her research interests include reconfigurable computing, distributed optimization, big data analytics, and deep learning.

## Orianna DeMasi

**PHD CANDIDATE,**
University of California, Berkeley

## Talya Eden

**PHD STUDENT,**
Tel Aviv University

### Developing a Dialog System to Augment SMS Helpline Counselor Training

The rise of smartphone usage has changed communication patterns and paralleled an increase in individuals reporting feeling lonely and disconnected, especially among young adults. As a result, crisis helplines are expanding their text services to keep pace with new communication trends during a time of growing need. However, tools for training helpline counselors provide insufficient practice and remain time consuming. I am working in partnership with a crisis helpline to build a dialog system that can be used to train crisis helpline counselors on how to intervene and de-escalate crises. This tool is designed as an interactive, automated tutoring agent that will simulate a person in crisis over text and allow counselors to practice intervening while receiving feedback on the particular approaches they choose to take. I will discuss the unique opportunities and structure within the problem setting and my work on the various associated challenges. Challenges include collecting appropriate data, developing a meaningful baseline system, comparing response selection with generation, and designing approaches to evaluate input quality and guide users towards improvement. Many of these challenges are frequently encountered with language systems, so I will discuss how the domain poses a novel opportunity to mitigate some of them. Such a system could help novice counselors learn, practice, and build confidence with crisis intervention strategies in a low pressure environment that puts no individual in danger. As a result, such a system will hopefully help crisis centers meet the rising demand for their services.

#### BIOGRAPHY

Orianna DeMasi is a PhD candidate in computer science at UC Berkeley and a Data Science Fellow at the Berkeley Institute of Data Science. Her work focuses on applications of machine learning to improve mental health, well-being, and care delivery. Her work experience includes time at Lawrence Berkeley National Lab and Twitter. Previously, she held an ARCS Fellowship and received a bachelor's degree from McGill University. Her work has won awards from the Big Ideas competition, the CITRIS Mobile App Challenge, and the CITRIS Tech for Social Good program. She has worked to make the STEM fields more accessible by mentoring 12 undergraduate research projects, serving as co-president of Women in Computer Science and Electrical Engineering at Berkeley, and working as a teaching assistant for the Fung Fellowship for Wellness and Innovation, which focuses on designing tech solutions to address public-health challenges.

### Sublinear-Time Algorithms

The amount of recorded information in the world grows by the split-second. Each minute, more than 200 million emails move across the Internet; 350,000 tweets are sent worldwide; and 3.6 million web searches are being performed. This increase in information is not limited to the internet. Scientists, too, have more information than ever before. Biologists collect enormous numbers of measurements, astronomers fill banks of hard drives with observations of stars and galaxies, and earth scientists assemble detailed snapshots of the weather. While this abundance of information offers many opportunities, the problem of analyzing the data is becoming more and more difficult. How do we run efficient algorithms on this massive amount of ever-growing and constantly-changing data? What does "efficient" even mean at these very large scales? Classical algorithms are considered to be efficient if they run in time that is polynomial in the input size. However, in these massive data sets, even a single pass on the input could take too long. My research tackles this exact problem. Sublinear-time algorithms are only allowed to read a small fraction of the input at the cost of providing an approximate answer. Typically, these algorithms rely on randomization and guarantee to return a highly accurate answer with high probability. My research focus is on devising sublinear-time algorithms for estimating graph parameters and for property testing of graphs and Boolean functions.

#### BIOGRAPHY

Talya Eden is a PhD student in the Electrical Engineering School of Tel Aviv University, advised by Professor Dana Ron. Her research focuses on sublinear-time algorithms and property testing and, in particular, sublinear-time estimation of graph parameters. She received a bachelor's degree in medical sciences and computer science from Tel Aviv University and a master's in computer science (graduating magna cum laude). Her research is supported by the Weinstein Award for Graduate Studies, the Sephora Berrebi Scholarship for Women in Advanced Computer Science, and the Azrieli Fellows Scholarship. During her studies, she was a teaching assistant in the graduate courses Computational Learning Theory (2018) and Design and Analysis of Algorithms (2015–2017), and in the undergraduate courses Computational Models (2015) and Data Structures and Algorithms (2016–2018).

## Aishwarya Ganesan

**PHD CANDIDATE,**
University of Wisconsin–Madison

## Chaya Ganesh

**POSTDOCTORAL RESEARCHER,**
Aarhus University

### Towards Truly Reliable Distributed Storage

Distributed storage systems have a simple goal: to reliably store and provide access to user data. However, realizing this apparently simple goal is fraught with challenges. Power loss, a system crash, or a faulty storage device can cause such systems to lose data or become unavailable. To protect against failures, these systems have traditionally have paid a cost: performance. By understanding how failures occur, the seemingly conflicting goals of reliability and performance can both be realized. We take a two-pronged approach to improving distributed storage system reliability: (1) We build two reliability-testing frameworks to analyze the effects of faulty disks and crash failures and use them to examine 10 systems (e.g., ZooKeeper, Cassandra, and Redis). Our analysis leads to new fundamental insights on how current reliability measures fall short. For example, most systems don't effectively use redundancy to recover from a local corruption: a single fault can cause disastrous outcomes such as data loss, silent corruption, and unavailability. Many vulnerabilities exposed by our tools have been acknowledged and fixed by developers. (2) Using these insights, we develop protocol-aware recovery (PAR), a new technique improving resiliency of faulty storage devices. A key aspect of PAR is that it isn't specific to a system; rather, it exploits the properties of protocols common to many distributed systems. For instance, we apply PAR to two different systems that implement a replicated state machine by exploiting their common properties. We show that the PAR versions safely recover from storage faults and provide high availability, while the unmodified versions can lose data or become unavailable, and that the PAR versions have little or no performance overhead.

#### BIOGRAPHY

Aishwarya Ganesan is a PhD candidate in computer sciences at the University of Wisconsin-Madison, advised by Professors Andrea Arpaci-Dusseau and Remzi Arpaci-Dusseau. Her research interests are in distributed systems, storage and file systems, and operating systems, with an emphasis on improving distributed systems' reliability without compromising performance. She received the best paper award at FAST 18 and a best paper nomination at FAST 17. Previously, she was a research fellow in the Mobility Networks and Systems group at Microsoft Research India. There, she built systems that leverage smart glasses to enable new applications such as physical analytics and near-vision communication. She received a master's degree from the Indian Institute of Technology Bombay and a bachelor's degree from Coimbatore Institute of Technology.

### Efficient Zero-Knowledge Proof Systems for Composite Statements

Zero-knowledge proofs provide a powerful tool, which allows a prover to convince a verifier that a statement is true without revealing any further information. It is known that every language in NP has a zero knowledge proof system, thus opening up several cryptographic applications. While true in theory, designing proof systems that are efficient enough to be used in practice remains challenging. Known approaches in prior work are each suited for certain representations of statements. But statements that arise in practice are composite statements that have components represented in different ways: Boolean/arithmetic circuit, algebraic representation. For instance, verifying an RSA signature involves checking a hash function H (represented as a boolean circuit) and computing exponentiations (algebraic group operations). Given a message m and a purported signature s, verification involves checking if s^e mod N=H(m) for the RSA public key (e,N). The state of the art fails to take advantage of the best of all worlds and has to forgo the efficiency of one approach to obtain the other's. My research focuses on new zero-knowledge proofs techniques that are suitable for proving such composite statements. In particular, my work is in designing efficient protocols for proving a wide class of statements motivated by applications in different settings: protocols that use symmetric-key operations (and very few public-key operations) that are suited where interaction is acceptable; non-interactive proofs that allow efficient verification for applications on the blockchain where the proof needs to be short and posted on the chain; and protocols that achieve the strong notion of adaptive composable zero-knowledge that is necessary for applications in secure multi-party computation.

#### BIOGRAPHY

Chaya Ganesh is a postdoctoral researcher in the Cryptography group at Aarhus University. She received a PhD from New York University (NYU) in 2017, supervised Professor Yevgeniy Dodis. She has completed internships at Microsoft Research, Bell Labs, and Visa Research. Her research is on secure computation protocols and zero-knowledge proofs. She maintains research interest in theoretical and applied Cryptography. She was awarded the Henry MacCracken Fellowship by NYU during her doctoral studies. She has served on the program committees of CANS 2018 and SEC 2018.


## Esha Ghosh

**POSTDOCTORAL RESEARCHER,**
Microsoft Research



## Alysson Rebecca Gold

**PHD STUDENT,**
Stanford University


# Integrity and Privacy in the Cloud: Efficient Algorithms for Secure and Privacy-Preserving Computing on Outsourced Data

An integral component of the modern computing era is the ability to outsource data and computation to remote cloud service providers (CSPs). The advent of cloud services, however, raises important challenges of in terms of integrity and privacy of data and computation. As soon as users delegate computation to cloud platforms (such as Microsoft Azure or Amazon AWS), concerns related to integrity and privacy of the data arise. Moreover, in the face of an alarming number of data breaches and massive surveillance programs around the globe, the privacy of outsourced data is becoming more important than ever. My research focuses on designing efficient privacy-preserving and verifiable mechanisms to outsource data processing and computation along with prototype implementation and experimental validation. In particular, I have been focusing on the following setting: how can a trusted data owner outsource her (possibly encrypted) data to an untrusted server such that the server will not be able to cheat while answering queries on the stored data. My work falls under the following broader areas: secure two-party computation, encrypted search, and zero-knowledge verifiable computation.

### BIOGRAPHY

Esha Ghosh is a postdoctoral researcher in the Cryptography Group at Microsoft Research Redmond. She received a PhD from Brown University in May 2018. She was the winner of a Microsoft Research Dissertation Grant (2017-2018). During her PhD studies, she also won the Paris C. Kanellakis Fellowship at Brown University and the best student paper award at the International Conference on Applied Cryptography and Network Security. Previously, she received a master's degree from the Indian Institute of Technology Madras.

# A Non-Perturbative Numerical Solver for Nonlinear Beam-Wave Interactions in Electron Beam-Driven Radiation Sources

Modern technology relies on complex, nonlinear interactions between electrons and electromagnetic fields. Modeling these interactions is particularly challenging in vacuum electronic sources, the enabling technology in high-power communications and particle accelerators, as approximations such as neglecting relativity, system boundaries, or space charge are invalid. When applied to the next generation of designs, existing simulation codes take on the order of days to run, with results highly sensitive to the mesh and particle distribution. My research focuses on new techniques to numerically model steady state beam-wave interactions efficiently and robustly, reducing simulation times by over two orders of magnitude from the current state of the art. By applying concepts from Lagrangian mechanics and classical field theory in the computation of the fields and particle trajectories, we obtain an unprecedentedly complete treatment of the space charge and current density, which drive the self-fields of the beam. Unlike conventional approaches using edge elements, the resulting nodal formulation is highly conducive to multi-physics problems where the solution is coupled to thermal or structural effects. Moving beyond vacuum electronic sources, including additional interaction terms in the Lagrangian, would render this approach suitable for the similar modeling challenges of semiconductor chip design. By solving for the electromagnetic four-potential instead of the fields, this approach could also have intriguing applications in metamaterials and quantum devices to model physical phenomena that do not manifest in interactions with the fields, such as the Aharonov-Bohm effect.

### BIOGRAPHY

Alysson Gold is pursuing a PhD in electrical engineering at Stanford University, focusing on the modeling, design, and testing of high-power radio-frequency sources with Professor Tantawi. Her research interests are centered on the interaction of intense charged particle beams and electromagnetic fields. Recently, she has focused on numerical methods in electromagnetism and high-performance scientific computing, including an internship at Nvidia Research. While studying for a bachelor's degree in engineering physics at the University of British Columbia, she developed a highly interdisciplinary background with expertise in areas from theoretical physics and applied mathematics to rapid prototyping and circuit design. She enjoys applying this background to research that balances theory, computation, and hands-on experimental work. She is the APS DPB Early-Career Representative, a Siemann Fellow, and an NSERC Alexander Graham Bell Graduate Scholar. She has won multiple awards, including for the best student poster at the 2018 International Particle Accelerator Conference.

## Hannah Gommerstadt

**PHD STUDENT,**
Carnegie Mellon University

## Diana Cristina Gonzalez Gonzalez

**POSTDOCTORAL FELLOW,**
Massachusetts Institute of Technology

### Session-Typed Concurrent Contracts

Multi-process systems control the behavior of everything from data centers storing our information to banking systems managing money. Each one of these processes has a prescribed role in their contract that governs their behavior during the joint computation. When a single process violates their communication contract, the impact of this misbehavior can rapidly propagate through the system. My research aims to detect and contain this misbehavior by dynamically monitoring violations of a process's contract. In order to model concurrent computation, we use a session type system which allows us to reason about processes that communicate over channels by message-passing. We annotate each channel with a session-type that prescribes the communication contract on that channel. However, many contracts cannot be expressed as a type; for example, consider a factoring process that takes in an positive integer and returns two positive integer factors. We can enforce that this process will return positive integers, but the guarantee we want is that the product of the returned factors is equal to the original number which cannot be expressed as a type. To handle these types of contracts, we have designed a monitoring mechanism that consists of monitoring processes that execute concurrently with the original process. These processes abort computation in the case of a contract violation, but are otherwise transparent. The monitors are able to express a variety of contracts such as determining whether a set of parentheses is matched, a list is sorted, or a tree is serialized correctly. I am currently working on expanding our model to be able to monitor even more exciting contracts and investigating how to monitor information flow in a concurrent setting.

#### BIOGRAPHY

Hannah (Anna) Gommerstadt is a PhD student in the Computer Science Department at Carnegie Mellon University, where she works with Professors Frank Pfenning and Limin Jia. Her research interests are at the intersection of programming languages and security, especially the use of language-based and logic-based methods to provide formal security guarantees. Her current work focuses on dynamic monitoring and contract enforcement in a concurrent setting. She received the Microsoft Graduate Women's Fellowship (2015-2016) and a CMU Presidential Fellowship (2017-2018), and she cofounded CMU's PhD Women Group. She received a bachelor's degree in computer science and mathematics from Harvard University in 2013. Before entering graduate school, she worked as a software engineer at Microsoft.

### On the Performance Limits of Relaying Networks at Low-SNR Regime

We investigate a communication network composed of multiple hops and multiple relays per hop. The relays operate under the decode-and-forward protocol with the set of relays at each hop following a degraded broadcast channel and all the links undergoing multipath fading in the low-SNR regime. Also, the channel state is unknown to all the terminals (non-coherent setting). In this context, we aim to derive capacity bounds for the investigated network, capitalizing on peaky frequency-shift keying (information-theoretically optimum at low SNR) and the theory of network equivalence. Equally important, we propose to consider and devise practical low-SNR schemes that are inspired by information-theoretical insights, while abiding by realistic constraints such as limited peak power and finite bandwidth.

#### BIOGRAPHY

Diana Cristina Gonzalez Gonzalez is currently a Sao Paulo Research Foundation (FAPESP) postdoctoral fellow with the Network Coding and Reliable Communications Group at MIT. She received a bachelor's degree in electronics engineering from the Escuela Colombiana de Ingenieria Julio Garavito in Colombia in 2007 and master's and PhD degrees in electrical engineering from the University of Campinas in Brazil in 2011 and 2015, respectively. In 2015, she served as a visiting researcher with the Vodafone Chair Mobile Communication Systems Technische Universitat in Dresden, Germany. From 2016 to early 2018, she was a postdoctoral fellow with the Wireless Technology Laboratory at the University of Campinas Brazil. In 2016, she received the Best Paper Award from the XXXIV Brazilian Telecommunications Symposium. She was also awarded a postdoctoral scholarship by FAPESP in 2016, an international mobility scholarship by Santander in 2015, and doctoral and master's scholarships by the Brazilian Ministry of Education in 2012 and 2010, respectively.


## Nil Zeynep Gurel

**PHD STUDENT,**
Georgia Institute of Technology



## Sepideh Hassan Moghaddam

**PHD CANDIDATE,**
University of Southern California


### Toward Wearable Sensing Enabled Closed-Loop Non-invasive Vagus Nerve Stimulation

Autonomic nervous system imbalance with elevated sympathetic activity is the hallmark of many mental disorders, and vagus nerve stimulation (VNS) using an implantable device is a proven treatment method to address such imbalance. Non-invasive cervical VNS (n-VNS) offers an inexpensive and low-risk alternative to implantable VNS devices, but may require auxiliary real-time physiological sensing to ensure efficacy and potentially close the loop for stimulation parameter tuning. In this work, in collaboration with Emory University School of Medicine Departments of Cardiology, Radiology, Psychiatry, and Rollins School of Public Health, I investigate the real-time, non-invasively-measured, cardiovascular and peripheral biomarkers of n-VNS that could quantify the stimulation efficacy based on the changes in sympathetic activity. I also evaluate the effects of n-VNS applied after mental stressors on stress-related physiological changes detected through wearable sensing. Healthy subjects and patients with post-traumatic stress disorder (PTSD) are enrolled in a three-day protocol that includes active or sham n-VNS application, bio-signal, blood biomarker, and high-resolution positron emission tomography (HR-PET) collection. Various mental stressors are applied throughout the protocol. Near-term research questions include: (1) understanding the kinetics of n-VNS from bio-signals, and the relevance of wearable sensor data to the gold standard serum biomarkers and imaging, and (2) real-time optimization of n-VNS treatment efficacy through wearable sensing.

#### BIOGRAPHY

Nil Gurel is a PhD student in the School of Electrical and Computer Engineering (ECE) at Georgia Tech. Her research focuses on physiological modulation, monitoring, and active sensing. She works on biomedical instrumentation, signal processing, and machine learning to understand real-time non-invasive biomarkers to close the loop for treatment optimization. Nil received a bachelor's degree in electrical and electronics engineering from Boğaziçi University in Istanbul, Turkey, and a master's degree in ECE from the University of Maryland, College Park. Her master's degree work focused on bio-inspired sensing for micro-aerial vehicles. She has received numerous awards and fellowships, including the Clark School of Engineering Graduate Fellowship (2014-2015), a NextFlex Flexible Hybrid Electronics Best Poster Award (2017), and the IEEE Body Sensor Networks Conference Runner-Up Best Paper Award (2018), and she was the IEEE Engineering in Medicine and Biology Conference Open Finalist (2018). She also received a Teaching and Development Fellowship (University of Maryland, 2015) and the National Science Foundation iREDEFINE Professional Development Award (2018).

### Analysis, Design, and Optimization of Large-Scale Networks of Dynamical Systems

My research focuses on structure identification and optimal control of large-scale networks of dynamical systems. These systems are becoming increasingly important in science and engineering with applications spanning economics, social networks, power systems, and robotics. Consequently, network science has emerged as an interdisciplinary field that draws from diverse disciplines, including graph theory, matrix theory, dynamical systems optimization, and statistical mechanics. Conventional optimal control of distributed systems relies on centralized implementation of control policies. In large systems, centralized information processing imposes a heavy burden on individual nodes and is often infeasible. To overcome this challenge, I combine tools and ideas from control theory, distributed optimization, and compressive sensing to develop distributed estimation and control strategies that require limited information exchange between the individual subsystems.

#### BIOGRAPHY

Sepideh Hassan Moghaddam is currently pursuing a PhD in the Department of Electrical Engineering at the University of Southern California. Her primary research interests are in optimal design analysis and optimization of distributed and large-scale networks. Sepideh received a bachelor's degree in electrical engineering from Sharif University of Technology, Tehran, Iran, in 2013 and a master's degree in electrical engineering from University of Minnesota Twin Cities in 2016. She was the Gold Medalist at the Second National Astronomy Olympiad in 2006 and received the four-year 3M Science and Technology Doctoral Fellowship in 2013.


## Elizabeth Farrell Helbling

**PHD CANDIDATE,**
Harvard University



## Qiuyuan Huang

**POSTDOCTORAL RESEARCHER,**
Microsoft Research


# Design Tradeoffs in the Creation of Autonomous, Insect-Scale Flying Robots

There is a need for small mobile autonomous robots that can move through highly dynamic and complex environments to expand capabilities in a broad range of applications. For example, search and rescue in confined hazardous or inaccessible places; distributed sensor networks with better temporal and spatial resolution; and minimally invasive medical applications. Advances in manufacturing and the proliferation of small electronic devices have paved the way to realizing this vision with centimeter-scale robots. However, there are still significant challenges in making these highly-articulated mechanical devices fully autonomous due to the severe mass and power constraints. I propose that by viewing the system holistically, we can navigate the inherent tradeoffs in each component in terms of size, mass, power, and computation metrics and constraints. During my PhD studies, I developed strategies to create an autonomous insect-scale vehicle, the Harvard RoboBee, with unprecedented mass, power, and computation constraints. I will discuss my work on the analysis of control and power requirements for this vehicle as well as results on the integration of onboard sensors and power electronics.

### BIOGRAPHY

Elizabeth Farrell Helbling is a PhD candidate at Harvard University, where she focuses on the systems-level design of the Harvard RoboBee, an insect-scale flapping wing robot. Her research looks at the integration of the control system sensors and power electronics within the strict weight and power constraints of the vehicle. She received a bachelor's degree in engineering sciences from Smith College in 2012. She received a National Science Foundation (NSF) Graduate Research Fellowship and was co-author on the IROS 2015 Best Student Paper for an insect-scale hybrid aerial-aquatic vehicle. Her work on the RoboBee project is also featured at the Museum of Science in Boston, the London Science Museum, the Te Papa Museum of New Zealand, and has been covered in the media (Science Friday, BBC, and Popular Science). She is interested in the codesign of mechanical and electrical systems for mass-, power-, and computation-constrained robots.

# Deep Learning for Language Intelligence

My research interests are mainly in the areas of machine learning (focusing on deep learning) and natural language processing, or NLP (focusing on (visual) language intelligence). More broadly, my research goal is targeted at research domains in which deep but flexible meaning representation, strong inference, and learnable knowledge have a large payoff: for example, language intelligence robotics. A sample long-range application would be a personal assistant capable of interacting with many application systems to read from and affect the world. I have conducted research at Microsoft Research AI (MSR AI), centered around two themes: (1) neural symbolic approach for NLP, and (2) deep learning for visual-language intelligence. On the neural symbolic approach for NLP, we developed a new network architecture: the Tensor Product Generation Network (TPGN) for NLP, based on the general technique of Tensor Product Representations (TPRs) for encoding and processing symbol structures in distributed neural networks. On deep learning for vision-language intelligence, we proposed a hierarchically structured reinforcement learning approach to addressing the challenges of planning for generating coherent multi-sentence stories for the visual storytelling task. Finally, beyond pushing forward state-of-the-art research, I am a strong proponent of efficient and reproducible research. I also helped to ship the proposed techniques to Microsoft products and create real-world impact. We have been developing two bots, Caption Bot and Drawing Bot, in collaboration with the Microsoft product team.

### BIOGRAPHY

Qiuyuan Huang is a researcher in the Deep Learning group at Microsoft Research AI (MSR AI), having joined the group as a postdoctoral researcher in 2017. She received a bachelor's degree in 2011 from the Department of Computer Science at the University of Science and Technology of China (USTC) and a PhD in 2017 from the Department of Electrical and Computer Engineering at the University of Florida. Her research interests are in the areas of deep learning and natural language processing (NLP). She is committed to the process of bridging the gap between theory and practice and of bringing theoretical results to practical implementations. She also loves writing actual code, building real systems from scratch and making them work simply, efficiently, elegantly, and beautifully.

## Sandy Huang

**PHD CANDIDATE,** University of California, Berkeley

## Rupamathi Jaddivada

**PHD STUDENT,** Massachusetts Institute of Technology

### Enabling Robot Transparency with Informative Actions

As robots become more capable and commonplace, it is increasingly important that the policies they execute are transparent: human end-users need to understand how a robot will act, when it will fail, and and why it failed. Unfortunately, passive familiarization takes awhile—it would take people hours, possibly days, of riding in an autonomous car before they understand its driving style and which situations it can and cannot handle. A lack of understanding is dangerous because passengers may over-trust the car, expecting it to handle situations that it cannot. Thus, we need to speed up this process of making policies transparent. My research leverages insights from cognitive science, reinforcement learning, and optimization to select informative examples that more quickly enable end-users to understand a robot. We found that when human end-users see informative examples of a robot's behavior, they more quickly understand how a robot acts and which situations it can and cannot handle, even when the robot's policy is a complex black box. This leads to safer and more comfortable human-robot interaction compared to the typical approach of relying on passive familiarization.

#### BIOGRAPHY

Sandy Huang is a PhD candidate in the Computer Science Department at UC Berkeley, co-advised by Anca Dragan and Pieter Abbeel. She received a bachelor's degree in computer science (with honors) from Stanford University in 2013. Her research focuses on making robot policies more transparent and robust, with the goal of giving human end-users a better understanding of how a robot will act, when it will fail, and why it failed. Her work was nominated for a best paper award at Human Robot Interaction (HRI) 2018. She received the National Science Foundation GRFP NDSEG Fellowship, the Berkeley Chancellor's Fellowship, and the Google Anita Borg Memorial Scholarship. She was a research intern at DeepMind in summer 2018, working with Raia Hadsell.

### Toward Economically-Efficient and Technically Realizable Electric Energy Service

The proliferation of intermittent, renewable technologies combined with the customer service-oriented trends present new fundamental challenges. Given temporal criticality of operating the physical grid, it is essential to process vast sources of asynchronously gathered data online and use for multi-layered interactive automation at well-understood expected performance. Today's Supervisory and Data Acquisition (SCADA) architecture capable of enabling electricity service is limited to hierarchical temporal spatial assumptions that no longer hold. SCADA must evolve into flexible architecture capable of aligning end-to-end physical and economic processes. Our approach provides a means of multi-layered modeling of interactions in terms of commonly-understood energy and power variables characterizing multi-physical energy conversion processes. This modeling framework allows for zooming in and out of aggregate models as per the spatial and temporal granularity of interest. The most novel finding so far: such multi-layered modeling with intra- and inter-layer interactions not only eases the component-level control design, but also results in a linear interactive model for aggregate-level analysis, thus allowing for end-device participation at value. The approach uses formal mathematical techniques called spatial and temporal lifting that facilitate abstraction of inner component details, revealing only the interface variables (energy and power) to its neighbors and its time-varying limits to its coordinator as a function of economic signals that are learned by the component itself or communicated by its coordinating aggregator. Finally, this framework induces a method for systematic practical composition of an otherwise complex computer platform for simulating technical and economic information exchange within a future electric-energy system. Ultimately, the modeling framework-supported control design demonstrating the effects of end-to-end market transactions on the power grid is being shown using our home-grown scalable electric power system simulator (SEPSS) to result in efficient and technically realizable integration of renewables and near-optimal provision of electricity services.

#### BIOGRAPHY

Rupamathi Jaddivada is a third-year PhD student at MIT, supervised by Professor Marija Ilic. She received a bachelor's degree in electrical and electronics engineering from Jawaharlal Nehru Technological University in India in 2014, and a master's degree in electrical engineering and computer science from Carnegie Mellon University in 2015. Her research interests include modeling control and numerical simulations of complex dynamical systems, in particular for electric-power system applications. She interned at New Electricity Transmission Software Solutions in 2018 and received a gold medal in the Srinivasa Ramanujan Mathematics Competition (SRMC) in 2014.


## Akanksha Jain

**POSTDOCTORAL RESEARCHER,** University of Texas at Austin


### Towards Smarter Hardware Prediction Mechanisms

In today's data-driven world, memory system performance remains critical to the overall performance of many workloads. My research focuses on improving memory system performance by building smarter hardware prediction mechanisms. In particular, I have developed novel caching and prefetching solutions that improve significantly over prior art by exploring larger design spaces. For example, for cache replacement where previous solutions use heuristics targeting specific access patterns, I proposed Hawkeye, a cache-replacement policy which accommodates all access patterns by learning from Belady's optimal solution (ISCA 2016). Similarly, for irregular data prefetchers where existing solutions are limited to a few access patterns, I proposed the Irregular Stream Buffer (ISB) which realizes the previously unattainable combination of two popular learning techniques — namely, address correlation and PC-localization (MICRO 2013). Finally, when considering the interaction between cache replacement and prefetching, I uncovered a new design space that is bounded by Belady's 52-year-old prefetch-agnostic optimal solution on one end and our new prefetch-aware optimal solution on the other end (ISCA 2018). This new design space led to the development of the Harmony cache replacement policy which won the Cache Replacement Championship in 2017. Looking to the future, I believe that many other system optimizations — including both hardware and software optimizations — will benefit from a systematic exploration of larger design spaces. But as systems get more complex, it will be difficult to navigate these design spaces with human insight alone. Therefore, one important goal of my future research is to leverage sophisticated machine learning algorithms to help system designers build highly accurate memory optimizations.

#### BIOGRAPHY

Akanksha Jain is a postdoctoral researcher at the University of Texas at Austin. She received a PhD in computer science from the University of Texas in December 2016. She received bachelor's and master's degrees in computer science and engineering from the Indian Institute of Technology Madras in 2009. Her research interests are in computer architecture with a particular focus on the memory system and on using machine learning techniques to improve the design of memory system optimizations. Her work has been recognized with a best paper honorable mention at MICRO 2013, a Micro Top Picks honorable mention in 2016, and the first-place award at the Cache Replacement Championship in 2017.


## Seyedeh Mahsa Kamali

**PHD CANDIDATE,** California Institute of Technology


### Changing Perceptions in Optics: What Can a Thin Engineered Surface Do?

Over time, our understanding of light has undergone multiple paradigm shifts from confirmation of the wave nature of light to the formation of electromagnetic and quantum light theories. But the general perception of optics, i.e., how we manage the flow of light, has remained almost the same for millennia. Even today, optics reminds us of polished pieces of glasses lumped into systems. Recent advances in nanotechnology have presented a new route to optics using nanopatterned structures that enable precise control of optical wavefronts with subwavelength resolution. My research focuses on thin layers of engineered materials called metasurfaces that can dramatically alter what we perceive of optics and how this can lead to a paradigm shift in the design and fabrication of optical elements and systems. The main principle of optical element design is that the geometrical form of an object is directly correlated to its optical function; e.g., a spherical lens focuses light to a point while a cylindrical lens forms a linear focal spot. I developed flexible metasurfaces for decoupling the geometrical form and optical function of arbitrarily shaped objects. Ultrathin elastic metasurfaces can also vastly shrink the size of tunable optical elements. For example, to make a lens with variable focal length conventionally, multiple optical elements are stacked together, creating a large, heavy device. I demonstrated ultra-thin, ultra-light tunable lenses by radially stretching elastic metasurfaces. I am also interested in developing metasurfaces with enhanced control over light; an important limitation of a thin hologram is that it can only project one 3D image regardless of the angle of illumination. I have developed a new platform to circumvent this limitation by designing metasurfaces that allow for encoding two holographic images under different illumination angles without sacrificing resolution.

#### BIOGRAPHY

Seyedeh Mahsa Kamali is a PhD candidate in the Department of Electrical Engineering at Caltech. She received a master's degree from Caltech and a bachelor's degree from the University of Tehran, both in electrical engineering. She has authored or coauthored more than 50 papers for peer-reviewed journals and conferences and filed for eight patents. Her research has been featured in *Science Physics*, *Tech Briefs*, and *Laser Focus World*, among others. Honors include a second-place best student paper award at the IEEE Photonics Conference in 2016 and selection as an Everhart Lecturer at Caltech in 2018.

## Angjoo Kanazawa

**POSTDOCTORAL RESEARCHER,** University of California, Berkeley

## Mandy Barrett Korpusik

**PHD CANDIDATE,** Massachusetts Institute of Technology

### Perceiving Deformable Shapes: Humans, Animals, and Birds

As we approach a society in which intelligent machines and humans coexist, machines will have to perceive and understand how humans and other embodied agents interact with the world. However, this is challenging because agents act in the world by moving and deforming their bodies, which makes it difficult to establish correspondence between body parts and their changing visual appearance. My research goal is to develop systems that can make sense of deforming bodies in the 3D world. Such systems would enable: robots that learn by imitating people; markerless motion capture; and, ultimately, socially intelligent machines that understand human behavior. To pursue this vision, I have been developing methods for learning generative models of deformable 3D bodies and inferring their internal states from visual data. Specifically, my colleagues and I developed the first fully automatic method for recovering 3D meshes of humans from RGB images by fitting a deformable model. Recently, I combined the power of deep learning with geometric constraints to develop an end-to-end system that can infer 3D meshes of humans in real-time. I demonstrated the utility of recovering such a 3D representation by training a physically simulated agent to perform dynamic skills from YouTube videos. Finally, I also developed methods for other deformable objects, such as animals and birds, where much less data and supervision is available.

#### BIOGRAPHY

Angjoo Kanazawa is a postdoctoral researcher at UC Berkeley, advised by Jitendra Malik, Alyosha Efros, and Trevor Darrell. She received a PhD in computer science from the University of Maryland, advised by David Jacobs, and a bachelor's degree from New York University. Her research is at the intersection of computer vision graphics and machine learning, focusing on 3D understanding of deformable shapes such as humans and animals from monocular images. Her work received the best paper award at Eurographics 2016, and she received a Google Anita Borg Memorial Scholarship. She interned at Columbia University with Peter Belheumer, at NEC Labs America with Manmohan Chandraker, and at Max Planck Institute with Michael Black.

### Deep Learning Models for Spoken Dialogue Systems

Personal digital assistants such as Siri, Cortana, and Alexa must translate a user's natural language query into a semantic representation that the backend can then use to retrieve information from relevant data sources. For example, answering a user's question about the weather requires querying a database with information about the weather at a given time in a specific location. In my work, we are investigating deep learning techniques for performing such a semantic mapping from human natural language to a structured relational database. Specifically, we have explored convolutional neural network architectures for learning a shared latent space where vector representations of natural language queries lie close to embeddings of database entries that have semantically similar meanings. To assess our technology, we have applied these techniques to the nutrition domain and built a full system prototype for diet tracking on the iOS platform.

#### BIOGRAPHY

Mandy Barrett Korpusik received a master's degree from MIT in 2015. She is currently pursuing a PhD in the Spoken Language Systems group at MIT, advised by Dr. Jim Glass. She received a bachelor's degree in electrical and computer engineering from the Franklin W. Olin College of Engineering in 2013. She received the NDSEG Fellowship in 2015 and the best-rated poster presentation award for demos at the Spoken Language Technology (SLT) workshop in 2014. Her primary research interests include natural language processing and spoken language understanding for dialogue systems, and she has previously worked on user intent-detection, semantic tagging, and predicting user-purchase behavior.

## Pan Li

**PHD STUDENT,**
University of Washington

### A Bottom-Up Perspective of Power System Operation: Coordination of Demand Response and Distributed Energy Resources

The new perspective of looking at power system operation is to utilize the flexibility from electricity consumers and distributed energy resources. These flexibilities are starting to play an increasingly important role in energy generation and consumption. For example, with the advances in smart meters, the system operator will have more access towards end consumption profile. The program that promotes the interaction and responsiveness of the customers is called demand response. It offers a broad range of potential benefits on system operation and expansion and on market efficiency. It is, therefore, crucial to estimate its impact on the power system which we treat as an experimental design problem that allows sequential decision. In addition, the installation of renewables, i.e., photovoltaic (PV) panels, across the world has grown exponentially during the past decade. These renewable resources tend to be different from traditional large-scale generators as they are often spatially distributed, leading to many small generation sites across the system. The proliferation of these individual renewable generators (especially PV) has allowed for a much more flexible system but also led to operational complexities because they are often not coordinated. We design a decentralized market to engage the participation of small-scale distributed renewable energy resources. We also use reactive power compensation as demand response to alleviate the problem of fluctuated voltage resulted from introduction of renewables. We show that our proposed methodology achieves a more efficient, flexible, and secure environment to integrate renewables into the power distribution system.

#### BIOGRAPHY

Pan Li is a fourth-year PhD student at the University of Washington, advised by Professor Baosen Zhang. She received bachelor's and master's degrees in electrical engineering from the class of Gifted Young in Xi'an Jiaotong University in China. She also holds the diplome d'ingenieur from Ecole Centrale de Lille, France. Her interests focus on optimization and learning algorithms in complex physical systems with humans in the loop — for example, a power grid. She interned at Pinterest in 2017 and received the Keith and Nancy Rattie endowed fellowship and graduate fellowship from the Clean Energy Institute at the University of Washington. Her work was featured in the Electrical Engineering Kaleidoscope at the Department of Electrical Engineering, University of Washington.

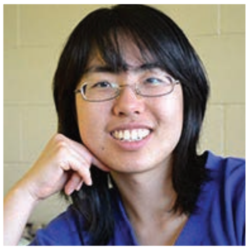

## Qi Li

**POSTDOCTORAL RESEARCHER,**
University of Illinois at Urbana-Champaign

### Truth Discovery from Multi-Sourced Data

Quality is a major challenge in this big-data era. Data can come from various sources, but those sources are not all equally reliable. It is intuitive to trust reliable sources more, but there is often little or no prior knowledge on which piece of information is accurate and which data source is more reliable. For example, in crowdsourcing, human workers may provide conflicting answers, but it is usually unknown who are more reliable and which answers are correct. This is the key challenge in detecting the veracity of information from multi-sourced data. Furthermore, the existence of heterogeneous data types, long-tail distributions, streaming data, unstructured data, and so on further increases the difficulty of discovering the true facts. My research focuses on developing novel unsupervised methods to detect the truths by integrating source reliability estimation and truth finding. Many of my methods have been successfully applied to numerous application domains (for example, health care, transportation, and the environment). They can be incorporated into various research fields (for example, web mining, crowdsourcing, mobile sensing, and information extraction) to improve data quality and achieve better performance.

#### BIOGRAPHY

Qi Li is currently a postdoctoral researcher in the Department of Computer Science at UIUC, working with Professor Jiawei Han. Qi obtained a PhD in computer science and engineering from the University at Buffalo in 2017, advised by Professor Jing Gao, and a master's degree in statistics from UIUC in 2012. Her research interests lie in the area of data mining with a focus on the collection and aggregation of information from multiple data sources. She has received several awards, including the Presidential Fellowship of the University at Buffalo, the Best CSE Graduate Research Award, and the CSE Best Dissertation Award from that university's Department of Computer Science and Engineering.


## Mengxia Liu

**PHD CANDIDATE,**
University of Toronto


### Colloidal Quantum Dots for Solar Energy Conversion

Increasing global energy demand drives the development of clean energy sources that will help reduce the consumption of fossil fuels. Solar energy, the most abundant renewable source, is converted to electricity using photovoltaic devices. The photovoltaics market has witnessed rapid growth in the past decade, and today, many photovoltaic strategies aim at low-cost solution-processed manufacture. Colloidal quantum dots (CQDs), emerging semiconductors, have attracted attention in view of their spectral tunability. The bandgap of CQDs is readily tuned to harvest infrared solar energy. This could enable both full-spectrum devices and also tandem solar cells that can be integrated with wider-bandgap semiconductors. Unfortunately, a high density of surface-associated trap states, low carrier mobilities, and an inhomogeneous energy landscape have previously limited CQD photovoltaic performance. Three strategies were developed to address these problems. Through new materials processing approaches, I succeeded in increasing charge extraction, reducing band tail states, and lowering the barrier to carrier hopping in CQD solids. The benefits from enhanced charge extraction were demonstrated in a double-sided junction architecture enabled by the engineering of an electron-accepting layer. A solution-phase ligand-exchange method was then developed to create CQD inks that can be deposited as an active layer in a single step. The resulting CQD films exhibited a flattened energy landscape that increased the carrier diffusion length and contributed to solar cells having certified efficiencies of 11.3 percent. After this, a hybrid material system was designed through combining CQDs with epitaxially-grown inorganic metal halide perovskites. The matrix-passivated CQD films showcased a two-fold increase in carrier mobility and superior thermal stability compared to pristine CQDs. My work provides promising pathways to achieve more fully the potential of CQD solids and to showcase these advances in improved performance for CQD solar cells.

#### BIOGRAPHY

Mengxia Liu is a PhD candidate in the Department of Electrical and Computer Engineering at the University of Toronto, working in Professor Edward H. Sargent's group. She received a bachelor's degree in materials science and engineering from Tianjin University, China, in 2014. Her research focuses on colloidal quantum dot optoelectronic devices, which have shown promise for low-cost, high-efficiency solar energy harvesting. After completing her PhD, she will move to the Cavendish Laboratory in University of Cambridge as a research associate. She has received the Chinese Government Award for Outstanding Self-Financed Students Award (2018), the Hatch Graduate Scholarship (2015-2018), and a China National Scholarship (2014).


## Seville Zeynep Lulec

**PHD CANDIDATE,**
University of Toronto


### Innovative Circuit Design Techniques for Wireless RF Transceivers with High Configurability and Low Power

The wireless industry has witnessed vast subscriber growth in the past decade. The diffusion of smartphones has contributed to increased data revenues which are expected to become even more important in coming years in light of 5G platforms. To satisfy the growing demands, a smartphone needs to be compatible with the latest as well as the previous standards, resulting in the accommodation of multiple transceivers. A software-defined radio (SDR) is a hypothetical concept that envisions a single transceiver for the entire cellular system and thus greatly reduces the overall required power area and cost. An SDR should be highly configurable, can tune to any band, and receive/transmit any modulation across a wide range of standards. In line with these trends, my research focuses on discrete-time implementations for wireless transceivers and mainly on the passive-switched-capacitor (PSC) circuits. PSC circuits do not require an active element and can achieve high dynamic range with accurate configurable characteristics with low power. However, PSCs have a few limitations: first, only real poles can be realized, limiting the minimum power and area consumption; second, their analysis is not intuitive and requires tedious charge-balance equations. My research addresses these limitations by introducing a continuous-time model for easy analysis and design of PSC circuits. By using the model, known structures can be analyzed in a simplified way that allows intuitive understanding and advanced designs that were not possible before. For example, I demonstrated a filter prototype that pushes the state of the art. This work can also branch out into new directions by employing machine learning techniques that can potentially replace filters in radios, ADCs, and sensor front-ends with high dynamic range, configurable, and technology-scaling friendly PSC counterparts.

#### BIOGRAPHY

S. Zeynep Lulec is a PhD candidate in the Department of Electrical and Computer Engineering at the University of Toronto, supervised by Professor Antonio Liscidini and Professor David Johns. Her studies focus on innovative, highly configurable discrete-time circuits that can help the realization of software-defined radio architectures for the next generation wireless systems. She received a bachelor's degree from the Middle East Technical University (METU) in Turkey in 2008 and a master's degree from École Polytechnique Fédérale de Lausanne (EPFL) in Switzerland in 2010, both in electrical and electronics engineering. She worked as a research engineer in Optical Microsystems Laboratory at Koc University and at the METU-MEMS Research and Applications Center, both in Turkey. She received an EPFL excellence scholarship and an ADI outstanding student-designer award. She is a member of IEEE SSCS YP Committee and founding board member of METU Alumni Association of Canada.


## Derya Malak

**POSTDOCTORAL RESEARCH ASSOCIATE,** Massachusetts Institute of Technology


### Coordinating Caching and Computation in Networks

To provide robustness to degradation of service and effective and reliable networked communications with consideration of the physical layer, it is required to have an understanding of the fundamental limits of reliable communications in networks. Using multi-hop Wi-Fi links for long backhaul connections can possibly be a cost-effective solution for rural connectivity. To this end, I have been developing reliable routing protocols, capacity-achieving recoding with delay guarantees for multi-hop wireless networks. Currently, my research focuses on caching for computation, using some of the most recent results in information theory to devise a theoretically motivated look at the central problem of coordinating computation and caching in networks. I have also investigated the trade-off between delay and rate in wireless networks with imperfect feedback. This approach has brought together signal flow techniques to the area of coding, melding two areas that have hitherto been quite separate. The results permit analysis of the perennially vexing problem of accounting accurately for delay when coding. I also worked on content caching in device-to-device networks and explored topics such as the Internet of Things and biology, and I am the lead inventor of a patent-pending technology.

#### BIOGRAPHY

Derya Malak is a postdoctoral research associate at MIT and Northeastern University, working with Professors Muriel Medard and Edmund Yeh. She received a PhD in electrical and computer engineering from the University of Texas at Austin in 2017, where she was supervised by Professor Jeffrey G. Andrews and affiliated with the Wireless Networking and Communications Group. Previously, she received a master's degree in electrical and electronics engineering from Koc University in Turkey in 2013, and a bachelor's degree in electrical and electronics engineering (with a minor in physics) from Middle East Technical University in Turkey in 2010. She held summer internships at Huawei Technologies and Bell Laboratories. She received the Graduate School Fellowship from the University of Texas at Austin and her postdoctoral appointment is funded by a grant from a DARPA program with MIT, Northeastern, and BBN on wireless networking.


## Zelda Mariet

**PHD STUDENT,** Massachusetts Institute of Technology


### Negatively Dependent Measures for Machine Learning

As machine learning becomes ubiquitous, so do the problems that plague it: inefficient use of data and computational power, time-consuming brute-force search through hyper-parameter space, lack of strong generalization guarantees, and more. My research focuses on analyzing negatively dependent measures as a tool to approach these fundamental problems. These measures, among which Determinantal Point Processes have already proven to be of significant interest to the machine learning community, encode negative dependence between items in subsets of a ground set: they endow the space with repulsive forces between similar points, enabling a careful balancing of the quality and diversity of a subset. My research focuses on two key problems: (1) developing scalable learning and sampling for negatively dependent measures over large datasets and (2) leveraging their properties to guide the search in machine learning design and analysis.

#### BIOGRAPHY

Zelda Mariet is beginning her fifth year as a PhD student in the Department of Electrical Engineering and Computer Science (EECS) at MIT, where she studies the theory and application of negatively dependent measures for machine learning model design and optimization. She received the 2018 Google PhD Fellowship in machine learning and was previously funded by the Criteo research faculty fellowship award. She has also interned at Google (Brain Research & Machine Intelligence), where she studied problems related to time series prediction and Determinantal Point Processes. Previously, she received bachelor's and master's degrees from Ecole Polytechnique in France, where she was accepted to the Corps des Mines after graduation. She was also a silver medalist at the SWERC programming competition.

## Negar Mehr

**PHD CANDIDATE,**
University of California, Berkeley

## Aastha Ketan Mehta

**PHD STUDENT,** Max Planck Institute for Software Systems

### Analysis and Control of the Current and Future Traffic Networks

The rapid increase of vehicular traffic congestion, delays, and emissions in metropolitan areas points to the importance of traffic management and control. As autonomous vehicles become tangible technologies, it will be crucial to investigate the impacts of autonomy deployment on such costs. My research has two themes: developing control and planning strategies for mitigating congestion in current traffic networks, and studying the potential impacts of autonomy deployment on future transportation networks. Because infrastructure expansion is a prolonged and high-priced process, development of traffic control strategies for increasing the throughput of the existing traffic networks is important. We have worked on development of such strategies and studied their effectiveness in various case studies. We have developed ramp-metering strategies and traffic-signal controls for reducing travel time of vehicles in freeway and urban traffic networks. To derive effective controls, we have also developed traffic models capable of capturing network-wide impact of driver behavior, such as lane-change maneuvers. It is envisioned that autonomy deployment will boost the network mobility. We have studied the validity of this impact under selfish routing behavior of the drivers in networks with mixed autonomy. We derived the conditions under which the network mobility is guaranteed to increase as a result of autonomy increase. We have shown that when these conditions do not hold, deployment of autonomous vehicles can worsen the total delay in traffic networks. Currently, we are studying routing and pricing mechanisms that can mitigate the inefficiencies of the coexistence of autonomous and regular vehicles.

#### BIOGRAPHY

Negar Mehr is a PhD candidate in the Department of Mechanical Engineering at UC Berkeley. She received a bachelor's in mechanical engineering from Sharif University of Technology in 2013. Her research interests include controls, cyber-physical systems, transportation engineering, and robotics. Her current focus is on analyzing transportation networks with mixed autonomy. She is developing algorithms and scalable tools for efficient management of future traffic networks. She was the co-recipient of the first prize for the best student paper award at the International Conference on Intelligent Transportation Systems in 2016. She was the graduate winner of the 2017 WTS-OC (Women Transportation Seminars) scholarship. She received the Chang-Lin Tien Graduate Fellowship in 2015 and 2017. She also received the Graduate Division Block Grant Award in 2015 and 2018, and the Eltoukhy East-West Gateway Fellowship in 2013.

### Pacer: Efficient I/O Side-Channel Mitigation in the Cloud

An important concern for cloud customers is the confidentiality of their cloud-hosted data and computation. Of particular concern are leaks through side channels that arise due to the sharing of cloud resources among distrusting tenants. For instance, by passively observing contention at a shared network link, one can infer the timing and size of a tenant's network packets, which in turn can reveal the tenant's secrets. To mitigate such network side channels, we are building Pacer, a system to shape a tenant's network traffic such that any publicly observable aspect of the traffic (e.g., packet sizes timing) is independent of the tenant's secrets. To make traffic shaping efficient, Pacer can dynamically adapt the shape of the tenant's traffic as long as adaptation is only based on public information in the tenant's requests. For instance, Pacer can allow the shape of HTML traffic to differ from that of video traffic if the type of the traffic requested (i.e., HTML or video) is public information. Pacer's design addresses several challenges in designing a secure, efficient, and practical traffic shaping system. Pacer relies on dynamic profiling of tenant applications to determine public input dependent traffic shapes, which helps to reduce end-user latency and bandwidth wastage during shaping. It relies on a small isolated component in the hypervisor to transmit packets at precise times specified in a tenant's transmission schedule. If the tenant fails to provide packets that the schedule expects, Pacer transmits dummy packets which are indistinguishable to a passive observer. Finally, Pacer's shaping is designed to be TCP-friendly and not cause a congestive collapse of the network. Preliminary experiments show that Pacer can protect private aspects of the tenants' traffic while incurring moderate overheads.

#### BIOGRAPHY

Aastha Ketan Mehta is a final-year PhD student at Max Planck Institute for Software Systems (MPI-SWS), advised by Peter Druschel and Deepak Garg. Her research interests involve operating systems, networking, and security. She received a bachelor's degree in computer science from Birla Institute of Technology and Science (BITS) in India in 2011. She spent a year as a technical staffer at NetApp Inc. in Bangalore, and she interned at Microsoft Research in the U.K. in 2015, working on SGX-enabled secure multiparty computing protocol and side channels in SGX. She was among 200 young researchers selected to attend the Fourth Heidelberg Laureate Forum.


## Baharan Mirzasoleiman

**POSTDOCTORAL RESEARCHER,**
Stanford University


### Big Data Summarization Using Submodular Functions

Data summarization, a central challenge in machine learning, is the task of finding a representative subset of manageable size out of a large dataset. It has found numerous applications, including image summarization, document and corpus summarization, recommender systems, and non-parametric learning, to name a few. A general recipe to obtain a faithful summary is to turn the problem into selecting a subset of data elements optimizing a utility function that quantifies "representativeness" of the selected set. Often times, the choice of utility functions used for summarization exhibit submodularity, a natural diminishing returns property. In other words, submodularity implies that the added value of any element from the dataset decreases as we include more data points to the summary. Thus, the data summarization problem can be naturally reduced to that of a constrained submodular maximization. Although there are efficient centralized algorithms for the aforementioned problems, they are highly impractical for massive datasets, as sequentially selecting elements on a single machine is heavily constrained in terms of speed and memory. Hence, in order to solve the above submodular optimization problems at scale, we need to make use of MapReduce-style parallel computation models or resort to streaming algorithms. We briefly overview the recent methods and theoretical analyses for submodular maximization in distributed and streaming settings and show the effectiveness of these techniques on several real-world applications, including summarizing 80 million Tiny Images and summarizing 45 million user visits from the Featured Tab of the Today Module on Yahoo! Front Page.

#### BIOGRAPHY

Baharan Mirzasoleiman is a postdoctoral researcher at Stanford University, working with Jure Leskovec. She is interested in designing and building large-scale machine learning and data mining algorithms. She received a PhD from ETH Zurich (2017), where she worked on large-scale summarization using submodular maximization, advised by Andreas Krause. She received a bachelor's degree from University of Tehran and a master's degree from Sharif University of Technology, where she worked on influence and information propagation through social networks. She held research internships at Google Research MTV (2016) and Google Zurich (2017). She received the ETH medal for an outstanding doctoral dissertation (2017), SNSF Early Postdoc.Mobility Fellowship (2016), and the Google Anita Borg Memorial Scholarship (2014).


## Nasim Mohammadi Estakhri

**POSTDOCTORAL RESEARCHER,**
University of Pennsylvania


### From Solving Equations to Compact Optical Devices: The Fascinating World of Nanoscale Light-Matter Interactions

Metamaterials hold the potential to bring innovative ideas into the emerging fields of nano-optics and nano-photonics. The physical properties of these artificially engineered structures can be designed through their subwavelength constitutive meta-molecules, providing an interesting platform to develop the next generation of optical devices with enhanced capabilities and smaller footprints. Within this context here, we present a fully-integrable metamaterial platform to solve complicated mathematical equations in the form of linear integral equations, using electromagnetic waves. An equation is intrinsically different from a forward problem (such as taking derivative of a signal) and entails a feedback mechanism. To address this point, our design employs a network of parallel waveguides as an external or internal feedback path for the optical wave which, along with the specially designed metamaterial operator, closes the signal loop and generates the solution to the integral equation in the steady states (in collaboration with Dr. Edwards and Professor Engheta at UPenn). This proposal enables the implementation of conceptually any linear integral equation, even with a translationally variant kernel. In the long term, this work is expected to provide the foundation for high-speed, ultra-small, and low-power analog optical processors. In the same context, yet for different applications, we also propose a new technique to achieve free-space optical modulation and interferometry with compact dielectric metasurfaces. Metasurfaces are two-dimensional counterparts of metamaterials that can be designed to mimic the functionality of conventional optical elements over orders of magnitude smaller footprints. In this project we introduce metasurface doublets for applications in modulators, waveplates, and pulse compression, suitable for mass-produced, compact, and cheap optical devices.

#### BIOGRAPHY

Nasim Mohammadi Estakhri is a postdoctoral researcher in Professor Nader Engheta's group at UPenn. She received a PhD in electrical engineering from the University of Texas at Austin in 2016 and a master's degree (summa cum laude) in 2011 and a bachelor's degree in 2007 in electrical engineering fields and waves, both from the University of Tehran. Her work has appeared in several prestigious journals. She received the George J. Heuer Jr. PhD Endowed Graduate Fellowship from UT at Austin, honorable mention at IEEE AP-S International Symposium student paper competition, and the IEEE Photonics Society Graduate Student Fellowship in 2016. She received the 2018 DMP Post-Doctoral Travel Award from APS and was invited to Stanford's EECS Rising Stars workshop in 2017.


## Rajalakshmi Nandakumar

**PHD CANDIDATE,**
University of Washington



## Shadi Noghabi

**PHD STUDENT,** University of Illinois at Urbana-Champaign


## Physiological Computing at Scale

Today's smartphones have multiple wireless sensors, such as Wi-Fi and acoustic sensors, and processing power equal to that of a computer. These wireless sensors can be leveraged to enable novel applications, specifically in the field of medical diagnosis. Further, the ubiquity of these devices enables deployment at a larger scale, creating an impact in the quality of life. To this end, I developed and deployed two applications: (1) ApneaApp to diagnose sleep apnea, and (2) a system that detects opioid overdose. Sleep apnea is a common medical disorder that occurs when breathing is disrupted during sleep. Today, diagnosing sleep apnea requires the polysomnography test which is an expensive and labor-intensive process. We designed ApneaApp which is the first contactless system that detects sleep apnea using just an off-the-shelf smartphone. It enables this by tracking the fine-grained breathing movements from multiple subjects without requiring any contact with the phone. We deployed the system with 37 patients at Harborview sleep lab and demonstrated that ApneaApp can detect apnea events with a correlation >0.95 compared to Polysomnography. Our work was the Best Paper Nominee at Mobisys 2015 and was exclusively licensed by ResMed Inc. Fatal opioid overdose remains a public health epidemic in the United States. Each day, 115 Americans die from opioid overdose. Unlike many life-threatening medical emergencies, opioid toxicity is readily reversible. Thus, a fundamental challenge of fatal opioid overdose events is that victims die alone with no or insufficiently timely diagnosis and treatment. To help connect potential overdose victims with life-saving interventions, we developed algorithms for smartphones that unobtrusively recognize opioid overdose by its physiologic precursors. We deployed our system in an approved supervised injection facility (SIF) where users self-inject illicit opioids. Our system had 96 percent sensitivity and 97 percent specificity (n=206) for identifying opioid-induced central apnea, a key event commonly preceding fatal opioid overdose.

### BIOGRAPHY

Rajalakshmi Nandakumar is a PhD candidate in the Paul G. Allen School of Computer Science and Engineering at the University of Washington. She works in the Networks and Mobile Systems lab, advised by Professor Shyamnath Gollakota. Her research focuses on leveraging the ubiquity of smart devices with wireless sensors to enable novel applications that can improve quality of life. She received the UW CoMotion Graduate Student Innovator award for her work on ApneaApp. Previously, she was a research assistant in the Mobility Networks and Systems lab at Microsoft Research India.

## Building Large-Scale Production Systems for Latency-Sensitive Applications

In today's era of increased engagement with technology, the myriad interactive and latency-sensitive applications around us necessitate handling large-scale data quickly and efficiently. My research focuses on designing and developing production-quality systems with particular attention to improving end-to-end latency and building massive-scale solutions. At large scale, providing low latency becomes increasingly challenging with many issues around distribution of data and computation, providing load balance, handling failures, and continuous scaling. We explore these issues on a wide range of systems, from a large-scale geo-distributed blob storage system that is running in production serving 450 million users (Ambry), to a stateful stream processing system handling 100s of TBs for a single job (Samza), and a real-time edge computing framework transparently running jobs in an edge-cloud environment (Fluid-Edge).

### BIOGRAPHY

Shadi Noghabi is a fifth-year PhD student in the Department of Computer Science at UIUC, co-advised by Professor Indranil Gupta and Professor Roy Campbell. Her research interests are in distributed systems, cloud computing, big data, and edge computing. She has published in SIGMOD and VLDB, and has collaborated extensively with industry (in particular with LinkedIn and Microsoft Research). Her research career has led and contributed to many production open-source projects, including Ambry (LinkedIn's geo-distributed object store) and Samza (a top-level Apache project used by more than 15 companies). She received the Microsoft Research Dissertation Grant and the Mavis Future Faculty Fellowship.


## Jelena Notaros

**PHD STUDENT,** Massachusetts Institute of Technology


# Integrated Optical Phased Array Architectures and Applications

Radio-frequency (RF) phased antenna arrays have enabled the development and advancement of numerous applications, such as radio transmitters and receivers, RADAR, television broadcasting, and radio astronomy. However, due to the relatively long wavelengths of RF frequencies, large-scale implementation of these arrays — necessary for generation of ultra-high resolution and arbitrary radiation patterns — is cumbersome and expensive. As a solution, integrated optical phased arrays, operating at much shorter optical wavelengths and fabricated in state-of-the-art integrated silicon photonics platforms, open up possibilities for large-scale applications in a variety of areas, including light detection and ranging (LIDAR) for autonomous vehicles, three-dimensional (3D) holography for visible-light displays, free-space optical communications, and trapped ion quantum computing. Here, a review of recent integrated optical phased array architectures, results, and applications will be presented. First, a novel CMOS-compatible silicon photonics platform with monolithically-integrated rare-earth-doped lasers and 3D-integrated driving electronics will be presented. Next, a variety of one-dimensional and two-dimensional optical phased array architectures and corresponding integrated photonic devices, including on-chip lasers, splitters, phase shifters, and antennas, will be discussed. Finally, recent experimental results showing far-field beam steering, near-field focusing, and arbitrary pattern forming will be presented and applied to a number of applications, including LIDAR and holography.

### BIOGRAPHY

Jelena Notaros is a PhD student and researcher in the Photonic Microsystems Group at MIT. She received a bachelor's degree from the University of Colorado at Boulder in 2015 and a master's degree from MIT in 2017. Her research interests are in integrated silicon photonics devices, systems, and applications. Her work has been published in top-tier Nature, OSA, and IEEE journals and conference proceedings. She is a DARPA Riser (one of 50 top emerging leaders in science and technology), an MIT Presidential Fellow, and a National Science Foundation Graduate Research Fellow. She received the 2014 IEEE Region 5 Student Paper Competition Award, 2018 and 2014 OSA Incubic Milton Chang Student Travel Grants, the 2014 Sigma Xi Undergraduate Research Award, the 2015 Chancellor's Recognition Award, the 2015 College of Engineering Outstanding Graduate for Academic Achievement Award, and the 2015 Electrical Engineering Distinguished Senior Award.


## Shijia Pan

**POSTDOCTORAL RESEARCHER,** Carnegie Mellon University


# Indoor Human Information Acquisition from Physical Vibrations

The number of everyday smart devices (such as smart TV, Samsung SmartThings, Nest, Google Home, etc.) is projected to grow to the billions in the coming decade. The Cyber-Physical Systems or Internet of Things (IoT) systems that consist of these devices are used to obtain human information for various smart building applications. Different sensing approaches have been explored, including vision-, sound-, RF-, mobile-, and load-based methods to obtain various indoor human information. From the system perspective, general problems faced by these existing technologies are their sensing requirements (e.g., line-of-sight, high deployment density, carrying a device) and intrusiveness (e.g., privacy concerns). My research focuses on non-intrusive indoor human information acquisition through ambient structural vibration, which I call "structures as sensors." People's interaction with structures in the ambient environment (e.g., floor, table, door) induces those structures to vibrate. By capturing and analyzing the vibration response of structures, we can indirectly infer information about the people causing it. However, challenges remain. Due to the complexity of the physical world (in this case, structures and people), sensing data distributions can change significantly under different sensing conditions. Therefore, from the data perspective, accurate information learning through a pure data-driven approach requires a large amount of labeled data, which is costly and difficult if not impossible to obtain in sensing applications. My research addresses these challenges by using physical insights to guide the sensing process. My system can robustly learn human information from limited labeled data distributions by iteratively expanding the labeled dataset. With insights into the relationship between changes of sensing data distributions and measurable physical attributes, the expansion order is guided by measured physical attributes to ensure a high learning accuracy in each iteration.

### BIOGRAPHY

Shijia Pan is a postdoctoral researcher at Carnegie Mellon. She received a PhD degree in electrical and computer engineering from Carnegie Mellon and a bachelor's degree in computer science and technology from the University of Science and Technology of China. Her research interests include cyber-physical systems, the IoT, and ubiquitous computing. She has published at computer science ACM/IEEE conferences (IPSN, Ubicomp) and in civil engineering journals (*Journal of Sound and Vibration, Frontiers in Built Environment*). Awards include a Nick G. Valais Graduate Fellowship, a Google Anita Borg Memorial Scholarship, best poster awards (SenSys, IPSN), a best demo and presentation awards (Ubicomp, SenSys Doctoral Colloquium), and audience-choice award (BuildSys) from ACM/IEEE conferences.


## Emma Pierson

**PHD STUDENT,**
Stanford University



## Neelam Prabhu Gaunkar

**PHD CANDIDATE,**
Iowa State University


## Fast Threshold Tests for Detecting Discrimination

Threshold tests have recently been proposed as a useful method for detecting bias in lending, hiring, and policing decisions. For example, in the case of credit extensions, these tests aim to estimate the bar for granting loans to white and minority applicants with a higher inferred threshold for minorities indicative of discrimination. This technique, however, requires fitting a complex Bayesian latent variable model for which inference is often computationally challenging. Here, we develop a method for fitting threshold tests that is two orders of magnitude faster than the existing approach, reducing computation from hours to minutes. To achieve these performance gains, we introduce and analyze a flexible family of probability distributions on the interval [0 1]—which we call discriminant distributions—that is computationally efficient to work with. We demonstrate our technique by analyzing 2.7 million police stops of pedestrians in New York City.

### BIOGRAPHY

Emma Pierson is a fourth-year PhD student in the Leskovec lab in the Stanford Department of Computer Science. Previously, she completed a master's degree in statistics at Oxford University on a Rhodes Scholarship. Her research applies statistics and machine learning to two application areas: (1) computational health and (2) discrimination and inequality. She also writes about these topics for broader audiences in publications including The New York Times, The Washington Post, The Atlantic, FiveThirtyEight, and Wired. Her work has received a best paper award (AISTATS 2018), a best poster award (ICML Workshop on Computational Biology 2016), and a best talk award (HitSeq 2015 at ISMB). She is supported by Hertz and NDSEG Fellowships.

## Design of Low Non-Uniform Field Sensors for Portable Diagnostic Systems

The evolution of wearable technologies and, more importantly, consumer awareness demands advancements in sensing mechanisms, analysis, and diagnostics. Magnetic field-based technologies such as current sensing, switching, navigation, and data recording have also evolved with the demand of lower operational power and long-term stability. However, spatial variations in magnetic fields produced by magnets have traditionally been a deterrent towards magnetic sensor utility in non-localized detection. In this work, the possibility of nuclear magnetic resonance (NMR) signal detection in low, non-uniform magnetic fields is explored. Non-uniform magnetic fields produced by permanent magnets were studied, modeled, and tested with portable pulsed field generation systems to validate NMR measurements for test samples placed in non-uniform fields. In this work, finite element simulations were initially used to estimate the spatial magnetic field in the exterior of a cylindrical ring magnet. A small region of uniform magnetic field, commonly called the saddle point, occurs at distances equal to the inner radius of the ring magnet. Test samples may be placed at saddle points that occur at locations where the field is maximal in one direction and minimal in another. The static flux density (B) at the saddle point was used to estimate the NMR voltage signal from hydrogen protons placed at the saddle point. From prior literature, it is known that the estimated NMR voltage is proportional to the square of B and the coil sensitivity, which depends on the pulsed field strength. Therefore, an efficient combination of the two interacting fields is required to detect the highest NMR voltage. Through this work, a systematic analysis and development towards designing portable NMR systems for examining localized targets will be achieved. Incorporating the estimations from the simulation results, NMR measurement results, and design considerations for the system will be presented.

### BIOGRAPHY

Neelam Prabhu Gaunkar is a PhD candidate in the Department of Electrical and Computer Engineering at ISU. She received a master's degree in electrical engineering from ISU in 2014. Her research interests include magnetic sensors, high-speed systems, and the phenomena of resonance. Awards include an IBM PhD Fellowship, an American Society for Non-Destructive Testing (ASNT) Fellowship, ISU's Zaffarano Award for excellence in research, and ISU Graduate Research and Teaching Excellence Awards. Since 2014, she has published 13 peer-reviewed articles, four articles in education journals, and numerous conference presentations.


## Maithra Raghu

**PHD STUDENT,**
Cornell University and Google Brain



## Niranjini Rajagopal

**PHD CANDIDATE,**
Carnegie Mellon University


# Insights from Deep Representations with Healthcare Applications

To continue the successes of deep learning, it becomes increasingly important to better understand the phenomena exhibited by these models, ideally through a combination of systematic experiments and theory. Central to this challenge is a better understanding of deep representations. I overview adapting Canonical Correlation Analysis (SVCCA) as a tool to directly compare latent representations across layers, training steps, and even different networks. The results highlight the difference between signal and noise stability in the representations and show us how we can representationally compress networks. We also gain insights into per-layer and per-sequence step convergence, along with differences between generalizing and memorizing networks. I then overview some of the applications of this style of analysis in helping us design better, more tractable tasks for uncertainty estimation with noisy health care data. Finally, I highlight some applications in the context of reinforcement learning — a new testbed that lets us study different RL algorithms, single agent, multiagent, and self-play settings and evaluate generalization in a systematic way.

### BIOGRAPHY

Maithra Raghu is a PhD candidate in computer science at Cornell University, working with Jon Kleinberg. She also collaborates extensively with Google Brain, particularly with Samy Bengio. She is interested in developing a science of deep learning and deep representations and applying these insights to health care.

# A Sensor Fusion Approach to Indoor Localization

Being able to locate people and things inside a building accurately, instantly, and cost-effectively will revolutionize how we interact with our indoor surroundings and open new application domains. Indoor localization at scale is not yet a reality due to technical barriers in terms of what sensors are available on commodity devices and gaps in scaling these systems from labs to realistic building environments. My work takes a principled approach to solving the problems faced by emerging ranging and localization technologies while scaling up with regard to sensor placement, mapping, and noisy sensor measurements. I apply tools from estimation theory and statistical signal processing to embedded sensing systems. My work addresses three major localization problems by fusing sensor data and floor plan information. First, we explore quantitatively how beacon placement impacts localization performance in complex indoor environments. Second, we design a floor-plan-aware localization solving technique that can accurately localize in the presence of inaccurate non-line-of-sight signals and insufficient line-of-sight signals. Third, we design a crowdsourced pedestrian-based mapping system on mobile phones for rapidly configuring beacon infrastructures. We have demonstrated our techniques with multiple real-world systems. Our work with ultrasonic beacons and unmodified phones won the Microsoft Indoor Localization Competition in 2015 and spun out into a start-up. Applied to ultrawide-band radios, our system won the same competition in 2018 with the ability to support 3D tracking. We also built a multi-user augmented reality application, which won the IPSN Best Demo Award in 2018. We are now working with NIST on localizing firefighters in operation without dedicated infrastructure in a building.

### BIOGRAPHY

Niranjini Rajagopal is a PhD candidate in electrical and computer engineering at CMU, advised by Anthony Rowe and Bruno Sinopoli. She received a master's degree from CMU and a bachelor's degree from the National Institute of Technology, Trichy, in India. Her research is at the intersection of estimation theory, statistical signal processing, and embedded sensing systems. She has interned at Texas Instruments and with Apple's location team. She is a graduate teaching fellow at the Eberly Center for Teaching Excellence and Educational Innovation. She has been part of the multi-university TerraSwarm and CONIX research centers. She has published at several ACM/IEEE conferences (IPSN, SenSys, RTAS, ICCPS, RTSS). She received the CMU William S. Dietrich II Presidential PhD Fellowship in 2015 and the Samsung PhD Fellowship in 2016.

## Alena Rodionova

**PHD CANDIDATE,**
University of Pennsylvania

### Foundations of Safe Autonomy: On-Board Verification and Formally-Constrained Machine Learning

Autonomous vehicles (AVs) have driven millions of miles on public roads, but even the simplest scenarios such as a lane change maneuver, have not been certified for safety. As there is no systematic method to bound and minimize the risk of decisions made by the vehicle's decision controller, the insurance liability of autonomous vehicles currently is entirely on the manufacturer. Because no single test can determine the safety of self-driving cars, there is a need for fundamental research on methods for assessing the safety of automated driving systems. The main objective of my research is to establish the foundations of Safe Autonomy — trustworthy and fair autonomy where people can safely trust, depend upon, and co-exist with autonomous systems in transportation.

#### BIOGRAPHY

Alena Rodionova is currently a PhD candidate in the Department of Electrical and Systems Engineering at the University of Pennsylvania. Her research is focused on formal analysis and verification of safety-critical systems such as medical devices and risk assessment verification and control of cyber-physical systems such as autonomous vehicles. She received bachelor's and master's degrees in mathematics from the Siberian Federal University in Russia in 2012 and 2014, respectively. Before joining UPenn in 2017, she was with the Cyber-Physical Systems Group at TU Wien in Austria.

## Rachel Rudinger

**PHD CANDIDATE,**
Johns Hopkins University

### Neural Models for Decompositional Semantics

Central to building systems that can understand and perform high-level reasoning about natural language texts is the ability to map surface level linguistic expressions (e.g., a sentence) to deeper semantic representations that can support inference. To train a system to perform such a mapping with modern statistical techniques requires both a designated semantic ontology and sufficient training data to fit the statistical models. Traditionally, these ontologies contain large numbers of exclusive categorical labels that require experts to annotate. Our work on decompositional semantics aims to replace these cumbersome ontologies with bundles of atomic semantic features that can be annotated by posing simple commonsensical questions to non-expert crowdsource workers. Using this approach, we have developed several datasets targeting semantic tasks such as Event Factuality Prediction (EFP) and Semantic (Proto-)Role Labeling (SPRL). EFP is the task of determining whether an event mentioned in text actually happened. For example, it entails processing a sentence such as "Pat watered the plants" or "Pat failed to water the plants" and answering the question "Did the watering happen?" SPRL involves predicting simple properties that distinguish participants in an event. For example, in processing a sentence such as "The cat ate the rat," we'd aim to predict which participant (the cat or the rat) "acted with volition" (the cat), "changed state" (the rat), and so on. For both tasks, we have developed neural models that achieve state-of-the-art performance in predicting the target semantic labels. These models make use of bidirectional Recurrent Neural Networks (RNNs) with both linear and tree structures. We employ multi-task training strategies, i.e., simultaneously training the networks on closely related semantic resources to yield further gains. Through targeted analysis, we can investigate the systematic improvements and errors these models make.

#### BIOGRAPHY

Rachel Rudinger is a PhD candidate in computer science at the JHU Center for Language and Speech Processing, working with Benjamin Van Durme. She received a bachelor's degree in computer science from Yale in 2013. Her research interests in natural language processing include commonsense reasoning, knowledge acquisition from text, and event semantics. She is a member of the Decompositional Semantics Initiative, and her recent work focuses on developing state-of-the-art neural models for fine-grained labeling of semantic properties in text. She interned at the Allen Institute for Artificial Intelligence and received the National Science Foundation Graduate Research Fellowship and the JHU Whiting School of Engineering Fellowship.

## Emma R. Schmidgall

**RESEARCH ASSOCIATE,**
University of Washington

### Integrated Photonic Circuits for Single Quantum Emitters

An entangled graph state of qubits is a valuable resource for both universal quantum computation and quantum communication. To date, entanglement generation rates are too low to realize these multi-qubit networks due to photon emission into unwanted spatial and spectral modes. The integration of crystal, defect-based qubits with photonic circuits can significantly enhance photon collection efficiency, albeit at the cost of degrading the defect's optical properties, such as an increase in inhomogeneous emission energies (linewidth broadening of GHz vs a few tens of MHz) and decreased spectral stability (spectral diffusion of tens of GHz vs a few hundred MHz). Compensating for this static and dynamic spread in emission energies is of critical importance for scalable on-chip graph state generation. We demonstrate the ability to tune the emission energy of photonic device-coupled near-surface NV centers over a large (200 GHz) tuning range with applied bias voltage. This is larger than the inhomogeneity of implanted NV centers, suggesting a viable route for indistinguishable photons from separate emitters. However, measurements on many single waveguide-coupled NV centers highlight the variability in response to an applied bias voltage. Despite this variability, we are able to apply real-time voltage feedback control to partially stabilize the emission energy of a device-coupled NV center.

#### BIOGRAPHY

Emma R. Schmidgall is a research associate in Professor Kai-Mei Fu's group at the University of Washington. She received a PhD in physics from the Technion-Israel Institute of Technology, a master's degree in nanomaterials from Imperial College London, a master's degree in physics from the University of Cambridge, and bachelor's degrees in physics and history from the California Institute of Technology (Caltech). She was a Marshall Scholar (2007–2009) and an Intelligence Community Postdoctoral Fellow (2016–2018).

## Alexandra Kathryn Schofield

**PHD CANDIDATE,**
Cornell University

### Local Privacy for Text Data to Train Distributional Semantic Models

Distributional semantic models such as topic models and word embeddings use word frequency and co-occurrence information to provide new insights for researchers investigating phenomena in large bodies of text, such as how concepts and subjects of study arise or change over time. Sometimes, to protect individual privacy or copyright, text data for training these models cannot be distributed freely. Existing mechanisms to provide privacy guarantees aim to make documents statistically indistinguishable, often by adding random noise proportional to variance in the data. While these processes may protect the identity of a patient in a medical dataset, they not only add more noise than distributional semantic models can tolerate, but also give too little information to users about which texts are used in the model and what topics arise in their documents. My work in this area focuses on anonymizing phrases and sentences instead of documents; for example, it allows a user to determine that a book uses the word "desert" more frequently than others, but makes it statistically impossible to reconstruct any sentence from that book. My strategy also uses data compression designed for text models in which one adds random noise to a compressed representation of a document's statistics, rather than the total document data, to reduce the privacy budget cost.

#### BIOGRAPHY

Alexandra Schofield is a PhD candidate in computer science at Cornell University, advised by Professor David Mimno. Her research focuses on natural language processing, specifically understanding and improving latent variable models for analysis of real-world datasets by humanist and social science researchers. This work ranges from experimental efforts to understand corpus preprocessing and front-end development of tools to interact with topic models to theoretical work on the limits of privacy for these models. Alexandra received a bachelor's degree in computer science and mathematics from Harvey Mudd College and spent one year as a software engineer at Yelp Inc. before joining her doctoral program. She has been awarded the NDSEG and Nationals Science Foundation GRFP Fellowships, the Microsoft Graduate Women's Scholarship, and the Google Anita Borg Memorial Scholarship.


## Sivaranjani Seetharaman

**PHD CANDIDATE**
University of Notre Dame



## Yuanyuan Shi

**PHD STUDENT,**
University of Washington


## Congestion in Large-Scale Transportation Networks: Analysis and Control Perspectives

Fluid-like models and their discretizations, such as the Cell Transmission Model, have proven successful in modeling traffic networks. However, given the complexity of the dynamics, it is not surprising that the dynamical properties of these models, especially in congested regimes, are not yet well characterized. My research addresses this gap by proposing a new modeling paradigm where an analogy between discretized fluid-like traffic flow models and a class of chemical reaction networks is constructed by suitable relaxations of key conservation laws. This framework allows us to draw upon powerful structural results from chemical reaction network theory to study the existence and stability of congested steady states in large-scale transportation networks. Using these models, we also propose control techniques that optimally utilize the vast quantities of real-time data such as density estimates and measurements obtained from mobile devices and smart infrastructure to design and implement both macroscopic (network-level) and microscopic (vehicle-level) real-time control actions that provide provable guarantees on congestion mitigation in large-scale traffic networks.

### BIOGRAPHY

Sivaranjani Seetharaman is a doctoral candidate in the Department of Electrical Engineering at the University of Notre Dame. She received undergraduate and master's degrees in electrical engineering from the PES Institute of Technology and the Indian Institute of Science in 2011 and 2013, respectively. Her research interests are in the area of distributed control for large-scale infrastructure networks with an emphasis on transportation networks and power grids. She is a recipient of the prestigious international Schlumberger Foundation Faculty for the Future Fellowship (2015–present) and the Zonta International Amelia Earhart Fellowship (2015–2016). She was also a Notre Dame (National Science Foundation) Ethical Leaders in STEM Fellow (2016–2017).

## Real-Time Control for Complex Physical Systems: A Tractable Data-Driven Approach

Decisions on how to best operate large-scale complex physical systems such as power systems, large commercial buildings, and transportation networks are becoming increasingly challenging because of the growing system complexity and uncertainty. For example, controlling a large building currently requires one to have a model describing its dynamics. However, the cost, time, and effort associated with-learning an accurate first-principles dynamical model of a large building would be enormous. Therefore, for simplicity and convenience, many control frameworks adopt a linear model that can be far away from the true system dynamics. To overcome the drawbacks of real-time control with first-principles models-or oversimplified linear models, we propose a new data-driven predictive control framework. We leverage the data obtained from sensor measurements to construct an input convex neural network for the physical system dynamics where the weights between neurons are constrained to be positive, and some direct "passthrough" layers are added for better representation power. Without loss of generality, we proved that input convex neural network can represent all-convex system dynamics and achieve comparable performance in fitting non-convex systems. Also, we proposed a projected gradient method in solving the constrained optimal control problem-where the system model is described by neural networks. Experiment results show that our data-driven control framework can achieve significantly more energy reduction compared with previous methods for large building systems and reduce the modeling time from years to minutes. Currently, a prototype implementation of the input convex neural-network control framework is being implemented in the University of Washington Electrical Engineering Building.

### BIOGRAPHY

Yuanyuan Shi is a fourth-year PhD student in the Department of Electrical and Computer Engineering at the University of Washington, advised by Baosen Zhang. She is also pursuing a master's degree in statistics. Previously, she received a bachelor's degree in control and system engineering from Nanjing University in 2015. Her research focuses on relating machine learning with control theory to solve emerging and open problems in complex physical system control (e.g., power system transportation networks).

## Misha Sra

**RESEARCH AFFILIATE,**
Massachusetts Institute of Technology

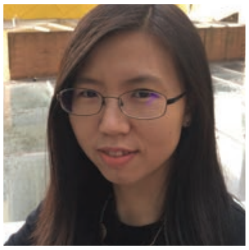

## Lili Su

**POSTDOCTORAL ASSOCIATE,**
Massachusetts Institute of Technology

## Perceptual Engineering

How can new technologies allow users to interact with digital content in the most direct and personally meaningful way? This question has driven interaction design for decades. In the evolution of computing, interaction has gone from punch cards to mouse and keyboard to touch and voice. Similarly, devices have become smaller and closer to users. Wearables and head-mounted devices bring computing into constant physical contact with the users' bodies. With every transition, interaction becomes more direct and the things people can do become more personal. The main question driving my research is: how can interactions and devices become even more direct and personal? My goal is to create systems that use the entire body for input and output and automatically adapt to each user's unique state and context. I call my concept "perceptual engineering"—i.e., immersive systems that alter the user's perception and influence or manipulate it in subtle ways. For example, they modify a user's sense of space, place, balance, and orientation, manipulate their visual attention, and more, all without the user's explicit awareness, to guide the interactive experience effortlessly. My research explores the use of cognitive illusions to manage a user's attention, breathing actions for direct-interaction, machine learning for automatic virtual-world generation, embodiment for remote collaboration, tangible interactions for play augmentation, and galvanic vestibular stimulation for reducing nausea and guiding users in immersive experiences. My perceptual engineering approach has been shown to (1) increase the user's sense of presence in VR/MR, (2) provide a novel eyes-, ears-, and hands-free way to communicate with the user through proprioception and other senses, and (3) serve as a platform to question the boundaries of our sense of agency and trust.

### BIOGRAPHY

Misha Sra is a research affiliate in the MIT Media Lab. She received a PhD from the Media Lab in 2018, working with Professor Pattie Maes. Her research asks the question: what if the whole body were an interface? Her work has been published at top HCI and VR conferences. She has received an ACM CHI best paper award, an ACM VRST best paper award, a Silver Award in the annual Edison Awards global competition, and other honorable mentions. Her work was featured in MIT Technology Review, UploadVR, and Business Insider, and on the Discovery Channel, and she was a Robert Wood Johnson Foundation well-being research fellow at the Media Lab.

## Defending Distributed Learning Against Arbitrarily Malicious Attacks

A distributed system consists of networked components that interact with each other to achieve a common goal. Given the ubiquity of distributed systems and their vulnerability to adversarial attacks, it is crucial to design systems that are provably secured. I have been exploring designing robust distributed learning algorithms that are provably resilient to Byzantine attacks. Two particular areas of focus: (1) Distributed statistical learning in the presence of arbitrarily malicious workers. Here, we focus on a map-reduce type of architecture and algorithm; we capture the distributed learning system using a server-client model where the clients are prone to Byzantine attacks. The adversary can adaptively choose the set of clients to attack. On the other hand, each worker only keeps a small sample. To train a model over this learning system, close interaction between the network components is necessary. We believe we are among the first to study this problem. We proposed methods that are provably resilient to Byzantine attacks even in the high dimensional setting. (2) Byzantine-resilient distributed inference over multi-agent networks. We studied Consensus-Based Multi-Agent Optimization and Consensus-Based Distributed Hypothesis Testing. For the former, we characterized the performance degradation caused by the Byzantine attacks and designed efficient algorithms that can achieve the optimal fault-tolerance performance. For the latter, we propose what we believe is the first learning algorithm under which the good agents can collaboratively identify the underlying truth.

### BIOGRAPHY

Lili Su is a postdoctoral researcher in the Computer Science and Artificial Intelligence Laboratory at MIT, hosted by Professor Nancy Lynch. She received a PhD from the University of Illinois at Urbana-Champaign in 2017, supervised by Professor Nitin H. Vaidya. Her master's work was on ordinal data processing at the CSL Communication Group from 2012 to 2014. She received a bachelor's degree from Nankai University in China in 2011. Her research intersects distributed systems, brain computing, security optimization, and learning. She was among three nominees for the Best Student Paper Award at the 2016 International Symposium on Distributed Computing, and she received the 2015 Best Student Paper Award at the International Symposium on Stabilization, Safety, and Security of Distributed Systems. She received UIUC's Sundaram Seshu International Student Fellowship for 2016. She has also served on program committees for several conferences.


## Wenjing Su

**HARDWARE ENGINEER,**
Google


## Additively Manufactured Reconfigurable Microwave Components Based on Microfluidics for Wireless Sensing and Internet-of-Things Applications

This research combines additive manufacturing and microfluidics with microwave and radio frequency (RF) electronics to provide a novel low-cost flexible and reconfigurable solution for the Internet of Things (IoT). Microfluidics, an emerging technology that allows the precise control of an extremely small amount of liquid within tiny channels, can be used in IoT applications to achieve Lab-on-Chip (LoC) functionality and an extraordinary reconfigurability. Microwave structures are very sensitive to the surrounding environment and thus, excellent sensors while passive radio frequency identification (RFIDs) provide low-cost zero-power solutions for wireless liquid sensing. This work has developed various proof-of-concept disposable wireless liquid sensors and RFID-based sensing platforms for LoC applications. To realize excellent reconfigurability inexpensively and compactly, this research also has studied new reconfigurable RF integration topologies by integrating microfluidic channels and dielectric or conducting liquids. To significantly decrease the production time and cost and enhance the ubiquitous smart items, this research studied additive manufacturing (AM) in IoT applications. This research has been exploring the possibility of replacing the conventional electronics and microfluidics fabrication methods with cost-effective additive manufacturing methods such as inkjet printing and 3D printing. This work presents first-of-its-kind, cost-effective, rapid, low-temperature, and environmentally friendly AM fabrication methods for various reconfigurable antennas or microwave components, wearable sensors, and sensing platforms. In summary, this research focuses on utilizing new AM fabrication techniques and novel microfluidics topologies to provide a low-cost, flexible, and scalable solution for wireless sensing and IoT applications.

### BIOGRAPHY

Wenjing Su currently works at Google as a hardware engineer focusing on wearable antenna designs. She received a PhD in electrical and computer engineering from the Georgia Institute of Technology, advised by Professor Manos Tentzeris. Her research interfaces advance novel fabrication techniques (e.g., inkjet printing, 3D printing), special mechanical structures (e.g., microfluidics), and high-performance microwave components/antennas to solve existing problems and extend the application of smart health wearable electronics in IoT applications. Her research interests include applied electromagnetics, RFID, additively manufactured electronics, wearable antennas, reconfigurable antennas, wireless sensing, machine learning aid sensing, flexible electronics, and green electronics.

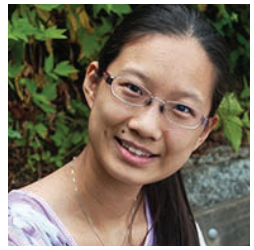


## Yihan Sun

**PHD STUDENT,**
Carnegie Mellon University


## Parallel Algorithms and Data Structures in Theory and Practice

Today, the explosion of data has raised new challenges necessitating considering parallelism and/or concurrency in algorithm design and engineering. On the other hand, multi-core processors are now used from smartphones to big database servers, making it possible to bring parallel algorithms to practice. Unfortunately, classic algorithm design and engineering mainly focus on the sequential setting, and implementations of parallel/concurrent algorithms are usually complicated, less portable and scalable. My research has been mostly focusing on designing parallel and concurrent algorithms that are simple, generic, and practically efficient and meanwhile have good theoretical worst-case guarantees. Some research topics that I have been working on include: (1) Parallel tree structures. I have a series of work on parallel and concurrent tree algorithms implementations and applications that have tight theoretical bounds and good practical performance. The most essential components include a theoretically-efficient algorithmic framework for parallel balance binary tree, as well as a parallel general-purpose C++ library implementing all these algorithms using fully-persistent (functional) tree structures. We apply this tree structure to many applications, e.g., range query, segment query, inverted indices, transactional memory with multi-version concurrency control (MVCC) and efficient garbage collection (GC) and an HTAP (hybrid transactional/analytical processing) database system with snapshot isolation. (2) Write-efficient algorithms. Motivated by the significantly higher cost of writing than reading in emerging NVM technologies, I have been seeking sequential and parallel write-efficient algorithms with an emphasis on tree algorithms and geometric algorithms. (3) Parallel algorithms in computational geometry. I have worked on both theoretical and experimental work in computational geometry. They either use our framework of the parallel tree structures or a group of randomized incremental algorithms. (4) Assorted parallel algorithms on other topics including parallel sort and semi-sort algorithms, parallel shortest path algorithms, and probabilistic tree embeddings.

### BIOGRAPHY

Yihan Sun is currently a fifth-year PhD student in the Computer Science Department at Carnegie Mellon, advised by Professor Guy Blelloch. She received a bachelor's degree in computer science from Tsinghua University in 2010. She received the Best Bachelor Thesis Award from Tsinghua (first place in the Computer Science Department) in 2014. Her research interests lie in parallel and concurrent algorithms and data structures and their applications that are simple, generic, and practically efficient and that, meanwhile, have good theoretical guarantees for worst-case performance.

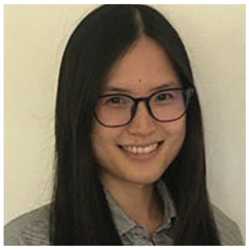


## Yixin Sun

**PHD STUDENT,**
Princeton University



## Mina Tahmasbi Arashloo

**PHD CANDIDATE,**
Princeton University


### Routing Attacks on Internet Services

My research focuses on exploring the interdependencies between the underlying Internet infrastructure and the overlay Internet services to uncover new security/privacy flaws and design robust countermeasures. Previous works are typically limited to analyzing individual network layers in isolation, missing critical interactions across different layers that adversaries can exploit. We have revealed new attacks that adversaries can exploit by manipulating Internet routing to compromise the security/privacy guarantees of Internet services such as the Tor network and the Public Key Infrastructure (PKI). We successfully perform real-world routing attacks on the live Tor network in an ethical manner with over 90 percent accuracy to deanonymize Tor users. We then develop proactive and reactive countermeasures that have impacted two Tor design specifications and are currently being integrated into the Tor Metrics Portal. We also successfully perform real-world routing attacks and obtain digital certificates of a victim domain from five top Certificate Authorities (CAs) compromising the security of PKI. We then develop two countermeasures to protect domain owners. Our countermeasure, the Multiple Vantage Point Verification, has been deployed by Let's Encrypt, the world's largest CA, which has issued hundreds of millions of certificates.

#### BIOGRAPHY

Yixin Sun is a PhD student at Princeton. She received a bachelor's degree in computer science (with highest distinction) and mathematics from the University of Virginia in 2013. Her research interests include network security and privacy. She received the Information Controls Fellowship from the Open Technology Fund in 2015 and the SEAS Award for Excellence from Princeton in 2017. In the past, she has interned at Verisign Labs NEC Labs and the International Computer Science Institute (ICSI).

### Programmable Network Monitoring and Control

Computer networks are crucial to online services. They transfer data between end users and data centers and between servers in data centers to process user requests. To ensure that these networks are efficient, secure, and highly available, it is essential to continuously (1) monitor them for congestion anomalies and failures and (2) take control actions accordingly on network devices — i.e., network interface cards (NICs) — at the end hosts and switches in the network. To that end, programmable network devices have emerged to provide operators with fine-grained control over how to process packets, what information to maintain across packets, and what control actions to take based on it. The programming interfaces of these devices, however, are quite low-level and only suitable for programming a single device, making them cumbersome to use in today's large-scale and complex networks. Therefore, my research focuses on designing programming platforms that facilitate the use of programmable network devices for large-scale and real-time network monitoring and control both at the end hosts and inside the network. More specifically, I have designed (1) domain-specific languages that are expressive enough for high-level specification of a broad set of network monitoring and control tasks while being amenable to efficient implementation and (2) compilers that use efficient intermediate data structures to automatically distribute and implement these specifications on programmable devices. So far, I have focused on programming platforms for end-to-end network transport, network-wide stateful monitoring and control, and path-based network monitoring. Moving forward, I plan to explore the design of a holistic monitoring and control platform using the rich but local and per-flow control actions on the NICs and the aggregate monitoring information observed by the switches to realize even more sophisticated network-wide policies.

#### BIOGRAPHY

Mina Tahmasbi Arashloo is a fifth-year PhD student in the Computer Science Department at Princeton, working with Professor Jennifer Rexford. She is primarily interested in networked systems with a focus on software defined networking (SDN) and programmable data planes. Previously, she received a bachelor's degree in computer engineering from Sharif University of Technology in Iran.


## Caroline June Trippel

**PHD CANDIDATE,**
Princeton University


### Concurrency and Security Verification in Heterogeneous, Parallel Systems

Modern computer systems achieve high performance by employing parallelism and high degrees of specialization and heterogeneity. This scenario presents challenges to reliable and secure systems design. For example, the heterogeneous compute elements make a variety of assumptions about the ordering of synchronization and memory accesses — i.e., they have different memory consistency models (MCMs). Additionally, with more and more functionality being put into hardware, security exploits increasingly attack the implementations on which seemingly secure algorithms are run. My dissertation creates formal automated tools and techniques for specifying and verifying the correct execution of software on heterogeneous parallel systems to improve reliability and security. To address reliability, my work has produced techniques and a novel tool-flow TriCheck for full-stack MCM verification. TriCheck leverages SMT-based analysis of hardware-aware program execution graphs to determine whether a specific combination of high-level language compiler ISA and microarchitecture could ever produce an illegal program outcome. In TriCheck's evaluation of the RISC-V ISA's support for C11 programs, TriCheck identified fundamental flaws in the RISC-V ISA's draft MCM specification. This work also identified bugs in two previously "proven-correct" compiler mappings from C11 to IBM Power and ARMv7, leading to the discovery of flaws in C11's MCM itself. To tackle security, my research adapts and extends the graph-based MCM analysis techniques above. Specifically, my work has produced a methodology and tool for evaluating a microarchitecture's susceptibility to formally-specified classes of security exploits and automatically synthesizing proof-of-concept exploit code when applicable. Using this tool to evaluate a standard out-of-order processor on its susceptibility to Flush+Reload cache side-channel attacks resulted in automatically synthesized programs representative of those that were publicly disclosed for conducting Meltdown and Spectre attacks. Evaluating the same processor on susceptibility to Prime+Probe attacks generated two new exploits, highlighting the importance of automated verification techniques for identifying hardware security vulnerabilities.

#### BIOGRAPHY

Caroline Trippel is a PhD candidate in the Computer Science Department at Princeton, advised by Professor Margaret Martonosi. Her work has resulted in formal, automated tools and techniques for specifying and verifying the correct and secure execution of software running on such systems. She has influenced the design of the RISC-V ISA memory consistency model (MCM). She received a bachelor's degree in computer engineering from Purdue University in 2013, a master's in computer science from Princeton University in 2015, and was a 2017-2018 NVIDIA Graduate Fellow.


## Ellen Vitercik

**PHD STUDENT,**
Carnegie Mellon University


### Foundations of Data-Driven Algorithm Design

My research spans artificial intelligence, algorithm design, and computational economics. I am particularly interested in developing machine learning tools for automated algorithm design. Practitioners often use machine learning techniques to optimize over families of parametrized algorithms, tuning parameters based on typical problems from their domains. Ideally, this results in algorithms with high performance on future problems. Unfortunately, parameter optimization procedures come with few guarantees. To counteract this trend, I am building a theoretical framework for application-specific algorithm design. My work involves characterizing problem-specific structure that I can exploit to provide computational tools with strong statistical guarantees. This task is challenging because across many domains, a small tweak in parameters can cause a cascade of changes in an algorithm's behavior. As a result, the algorithm's performance is a non-convex and volatile function of its parameters. In a related direction, my research also includes data-driven auction design, a special case of algorithm design with tremendous real-world impact. The rise of the internet has led to worldwide participation in electronic marketplaces, which in turn has generated a deluge of consumer data. In my work, I design algorithms that learn consumers' valuations from a small amount of data. In addition, I develop tools that utilize this data to design auctions and optimize prices, with revenue maximization as the goal. Finally, I am passionate about protecting against the societal risks that come with the exploitation of data, such as privacy loss. For example, large internet companies regularly tune prices according to consumer data. Without proper precautions, the resulting prices may leak information about consumers' purchase histories. In my work, I develop algorithms that utilize data while protecting private information contained therein.

#### BIOGRAPHY

Ellen Vitercik is a PhD student in computer science at Carnegie Mellon University, advised by Nina Balcan and Tuomas Sandholm. Her primary research interests are artificial intelligence, machine learning, theoretical computer science, and computational economics. Her honors include a National Science Foundation Graduate Research Fellowship and a Microsoft Research Women's Fellowship. She received a bachelor's degree in mathematics from Columbia University, graduating summa cum laude and Phi Beta Kappa.

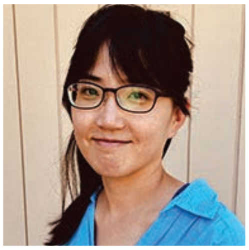


## Ching-Hua Wang

**PHD STUDENT,**
Stanford University


### 2D Materials for CMOS and Memory Integration

Modern CPU performance is reaching a plateau due to the scaling limitations of transistor technology. Additionally, a "memory wall" between CPU and memory has become a significant limitation as data usage has expanded exponentially requiring future computer technology to directly integrate memory and transistor devices. Layered two-dimensional (2D) materials have shown promise for overcoming key CPU limitations due to their intrinsically thin body that is preferable for device scaling. Also, the weak layer allows for low-temperature transfer processing, which is suitable for 3D monolithic CPU integration. I first studied the high-mobility 2D material black phosphorous (BP) transistors for Complementary Metal-Oxide-Semiconductor (CMOS) low-power digital circuits. It is required to develop both n- and p-type devices using the same semiconductor material to create CMOS technology. My work demonstrated that BP can be used as an effective n-type transistor by using ultra-low work function metals, achieving record high n-type current. I discovered the physical mechanisms of controlling doping and de-pinning effects for n- and p-type BP transistors, pushing BP closer to CMOS implementations. I have also focused on the integration of memory and transistors for memory arrays with a complete 2D material system using hexagonal-Boron Nitride (hBN) as a memory cell and monolayer Molybdenum Disulfide (MoS2) as the transistor selector. Utilizing the transferrable feature of 2D material, I demonstrated the very low processing temperature for a transistor switching memory structure and achieved the first two-floor stacking of transistor switching memory cell in a monolithic 3D structure. Besides using MoS2 transistor as a selector, I also investigated using nonlinear electrical transport in 2D material out-of-plane direction for two terminal selectors, which can greatly reduce the unit cell area. This 2D system can be easily integrated for monolithic 3D CPUs, overcoming the "memory wall."

#### BIOGRAPHY

Ching-Hua (Fiona) Wang is a PhD student in the Electrical Engineering Department at Stanford University, supervised by Professor H.-S. Philip Wong. She received a master's degree in electronic engineering from the National Tsing-Hua University in China in 2009, then participated in the 20 nm CMOS front-end R&D group at TSMC for two years. She also worked on nonvolatile memory in the Ya-Chin King and Chrong-Jung Lin research group, published twice at IEDM, and won the Taiwanese Government Scholarship for PhD programs. She received a best-in-session award at SRC Techcon 2017 for her work on Scandium contacts to black phosphorus.


## Zi Wang

**PHD CANDIDATE,**
Massachusetts Institute of Technology


### Towards Integrated Intelligent Systems

Throughout the history of artificial intelligence (AI), it has been a common theme to fragment the grand goal of building intelligent agents into many in-depth but often incompatible subproblems, including statistical machine learning, optimization, motion planning, computer vision, and so on. My research stems from a different perspective that emphasizes the breadth: how to integrate those individual modules into a cohesive intelligent system that can act and adapt itself in complex environments. My work focuses on one of the most challenging aspects of the integration of AI modules: learning and planning. Learning aims to represent knowledge as a model, while a planner makes decisions. This problem is a variant of the reinforcement learning problem, but the practicality of classic methods is doubtful for long-horizon high-dimension problems. A central question of integrating learning and planning is how to decide which actions to select from a potentially infinite action space. The choice of actions determines both the performance of the planner and the data to be gathered for learning.To this end, I have worked on the mathematical formulation of action selection as Bayesian optimization, active model learning for structured actions, and task and motion planning that enables long-horizon planning with learned models. Currently, I am investigating how to integrate perception, learning, and planning, and control modules to demonstrate intelligent behaviors on a physical robotic system. In particular, I would like to understand how to "wire" the existing modules together so that desired properties such as tractability and analyzability can be induced.

#### BIOGRAPHY

Zi Wang is a PhD candidate in the MIT Computer Science and Artificial Intelligence Laboratory, advised by Professors Stefanie Jegelka, Leslie Kaelbling, and Tomas Lozano-Perez. Her research interests lie broadly in machine learning and artificial intelligence, currently with applications to robotics problems. She received a master's degree in EECS from MIT in 2016, and a bachelor's degree in computer science and technology from Tsinghua University in 2014. She has received many awards, including the Greater China Computer Science Fellowship and the Google Anita Borg Memorial Scholarship. While at MIT, she served as co-president of Graduate Women in Course 6 (EECS) and as a reviewer of many leading conferences and journals in her field, such as the Annual Conference on Neural Information Processing Systems (NIPS) and the Journal of Machine Learning Research (JMLR).

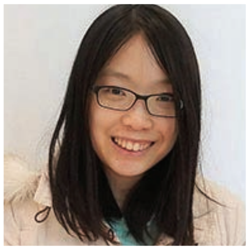


## Fan Wei

**PHD CANDIDATE,**
Stanford University



## Nora S. Willett

**PHD STUDENT,**
Princeton University


### Fast Algorithms for Property Testing

My research area is graph theory and combinatorics and their application to theoretical computer science. The goal of property testing is to design an algorithm to efficiently test whether an object such as a function or a graph has a certain property such as containing triangles or linearity, respectively. Typically, the input size of the object such as a graph is enormous and even polynomial time complexity is impractical in real-life applications. We want to design a randomized algorithm which only queries a small number of bits of the input independent of the input size but can still make a correct decision with good precision. Although constant time queries have been found in some cases, a common drawback and open question in the field has been that the constant is typically enormous. We resolved this question for testing permutations, i.e., strings of length n of distinct integers 1 to n, with respect to several commonly used distances in permutations defined from the perspective of geometry, algebra, and analysis. Not only did we show that a property testing algorithm with constant query complexity exists but also that the constant is mild. We managed to overcome the common obstacle in testing general properties. There are many other open questions in property testing related to graphs, sequences, and permutations depending on different metrics we use. I am actively working on these areas. Improving the bounds in these topics is closely connected to other fields of mathematics such as number theory. Besides property testing, I also work in other fields of theoretical computer science related to smooth analysis and social network. I apply methods in combinatorics to explain why certain natural algorithms on graphs and social networks with certain properties have small time complexity.

#### BIOGRAPHY

Fan Wei is a PhD candidate in the Department of Mathematics at Stanford University, supervised by Professor Jacob Fox. She received a bachelor's degree in mathematics from MIT and a master's (advanced study with distinction) from Cambridge University in the U.K. Her research interests are graph theory combinatorics and their application to theoretical computer science. She has held internships in the Microsoft Research New England and Microsoft Research Redmond Theory groups, and she received the Cambridge Overseas Trust Scholarship from Cambridge University. She also won the AMS-MAA-SIAM Frank and Brennie Morgan Prize for Outstanding Research in Mathematics and the MIT John. A. Bucsela Prize in Mathematics.

### Animation Authoring Tools

My research explores methods to provide animation novices with authoring tools for creating live 2D animations and dynamic illustrations. One process leverages simulation to automatically create secondary animation for live performances. Another designs an interface for triggering artwork swaps during a performance. To add expressiveness to live 2D animated characters, I developed behaviors that allow artists to apply secondary motion to parts of characters (Willett UIST'17). Secondary animation is the subtle movement of parts like hair, foliage, or cloth that complements and emphasizes the primary motion—the animation of the character's main parts. My work introduces physically-inspired rigs—the Follow rig, the Rest pose rig, and the Collision rig—that propagate the primary motion of layered illustrated characters to produce plausible secondary motion. Another method to improve the live performance of illustrated characters is through triggering discrete artwork changes on a character. In 2D animation, characters are represented by a set of artwork layers that typically depict different body parts or accessories. Within this context, discrete transitions are realized by swapping artwork layers to produce large changes in pose and appearance. I designed and evaluated a multi-touch interface for triggering artwork swaps during performances (Willett UIST'17). This trigger layout design enables users to quickly recognize and tap visual triggers without looking away from a live preview of the character. In addition, since animators typically practice before live performances, common patterns from practice sessions are encoded in a predictive model that is used to highlight suggested triggers during performances. Users of this system found it helpful in reducing the cognitive load of choosing artwork swaps during a performance.

#### BIOGRAPHY

Nora S. Willett is a fourth-year PhD student in computer science at Princeton, working with Adam Finkelstein. She received a best paper award at the ACM Symposium on User Interface Software and Technology (UIST) in 2017. She interned twice at Adobe Research and once at Autodesk Research. Previously, she received a bachelor's degree from Stanford University and worked on the film "Mr. Peabody and Sherman" at DreamWorks Animation.


## Anqi Wu

**PHD STUDENT,**
Princeton University


# Gaussian Process Based Nonlinear Latent Structure Discovery in Multivariate Spike Train Data

Much recent work in computational neuroscience focuses on methods for extracting low-dimensional, latent structure from multi-neuron spike train data. Most of these methods have used linear dimensionality reduction of firing rates, linear models of latent dynamics, or both. We propose a nonlinear latent variable model that can identify low-dimensional highly-nonlinear structure underlying high-dimensional spike train data, which we call the Poisson Gaussian-Process Latent Variable Model (P-GPLVM). This model specifies the joint distribution over multi-neuron spike trains in terms of conditionally Poisson spiking with firing rates driven by the composition of two underlying Gaussian processes (GPs)—one governing the trajectory of a low-dimensional temporal latent variable and another governing a set of tuning curves that map the latent variable to high-dimensional firing rates. Using nonlinear tuning curves allows the model to discover low-dimensional latent structure even when spike responses are themselves high-dimensional (e.g., as in hippocampal place cell or entorhinal grid cell codes). To infer the model parameters and GP hyperparameters from data, we introduce the decoupled Laplace approximation, a fast approximate inference method that allows us to efficiently optimize the latent path while marginalizing over latent tuning curves. We show that this method outperforms previous Laplace-based inference methods in both the speed and accuracy. We apply the model to spike trains recorded from hippocampal place cells and show that it outperforms a variety of previous methods for latent structure discovery, including variational auto-encoder (VAE) based methods that parametrize the nonlinear mapping from latent space to spike rates with a deep neural network.

### BIOGRAPHY

Anqi Wu is a PhD student in the Princeton Neuroscience Institute at Princeton, advised by Professor Jonathan W. Pillow. She received a master's degree in computer science at the University of Southern California and a bachelor's degree in electrical engineering from the Harbin Institute of Technology in China. Her research interest is to develop Bayesian statistical models to characterize structure in neural data in the interdisciplinary field of machine learning and computational neuroscience. She has worked as a summer research associate at the University of Texas at Austin and as an intern at at Microsoft Research Cambridge in the U.K. She has published in several conference proceedings at top machine learning conferences. She was nominated as one of two candidates for a Google PhD fellowship representing Princeton in 2018, and she received a Chevron Fellowship while at USC.

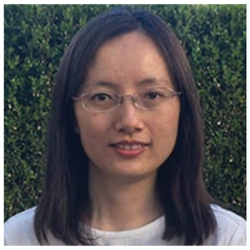


## Huijuan Xu

**PHD STUDENT,**
Boston University


# Video Understanding with Localization

Millions of videos come out every day, and the videos tend to be long and untrimmed with sparse events inside the video, such as surveillance videos. Automatically identifying events in the video would be very helpful for monitoring such large amounts of video data. However, the event localizations in previous temporal activity detection models take the sliding window-based approaches, which produce inflexible activity boundaries and are time-consuming. We propose Region Convolutional 3D Network (R-C3D) for temporal activity detection, which can detect arbitrary length activities and run in fast speed with proposal stage filtering out irrelevant background segments. Furthermore, we combine the event localization in video with natural language components, which enhances video understanding with richer language description via working in two topics: (1) dense video captioning which localizes distinct events in video stream and generates captions for the localized events and (2) natural language localization in video which requires temporally localizing the input query sentence in the video.

### BIOGRAPHY

Huijuan Xu is a PhD student in the Computer Science Department at Boston University, advised by Professor Kate Saenko. Her research focuses on deep learning, computer vision and natural language processing, and video understanding with localization. Specifically, her work explores visual question answering, video language description, cross-modal retrieval, and temporal activity detection. Her work "Region Convolutional 3D Network (R-C3D) for Temporal Activity Detection" won the Most Innovative Award in the ActivityNet Challenge 2017.

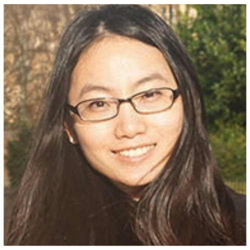

## Mengjia Yan

**PHD STUDENT,**
University of Illinois at Urbana-Champaign

### Secure Cache Hierarchies

In today's hyper-digitalized world, security and privacy have become primary concerns. Sensitive information must be processed in trustworthy ways on various computation platforms, ranging from mobile devices to public clouds. Unfortunately, state-of-the-art hardware technology today is vulnerable to sophisticated attacks. In fact, several key performance features of modern processors have demonstrated vulnerabilities to security attacks such as the recent Meltdown and Spectre attacks. Hardware exploitations break the assumptions underpinning numerous software security mechanisms, and thus represent serious threats to current systems. My research focuses on an important class of security threats: cache-based covert-side channel attacks. In these attacks, an attacker is able to stealthily leak sensitive information from a system without violating existing security policies enforced by the software layer. These attacks have been gaining popularity and currently exist for major vendor processors. My research makes contributions in two directions: designing new attacks and proposing new architectural defense techniques. First, I worked on designing new types of attacks targeting emerging applications and architectures to help the community to identify unexploited security vulnerabilities. For example, modern cache hierarchy designs are moving away from inclusive caches. Non-inclusive caches provide an illusion to researchers that they are resistant to cache attacks. However, I have demonstrated the first cross-core Prime+Probe attacks on state-of-the-art non-inclusive sliced caches showing that cache attacks are still a large security problem. Second, I worked on designing practical detection and defense mechanisms to combat cache-based side channels by leveraging architecture innovations which attain much better trade-offs between performance effectiveness and implementation complexity. I designed "ReplayConfusion," an efficient detection methodology against cache-based covert channel attacks. I also proposed "SHARP," a new secure hierarchy-aware cache replacement policy to effectively defend against cross-core cache-based side channel attacks.

#### BIOGRAPHY

Mengjia Yan is a PhD student at UIUC, advised by Professor Josep Torrellas. She received a master's degree from UIUC in 2016 and a bachelor's degree from Zhejiang University in 2013. She was selected as a Mavis Future Faculty Fellow for 2018, and she received a W.J. Poppelbaum Memorial Award and an ACM SIGARCH Student Scholarship for Celebration of 50 Years of the ACM A.M. Turing Award in 2017.

## Hazal Yuksel

**POSTDOCTORAL RESEARCH ASSOCIATE,**
Columbia University

### Widely Tunable FDD Wireless: Challenges and Opportunities

The recent trend in proliferation of wireless devices coupled with upcoming 5G standards requiring rates of 100Gb/s has created an increasing need for high throughput wireless data. Spectrum has become an expensive resource: an FCC auction for a 4MHz bandwidth around the 700MHz generated $19B in 2008. There are only two ways to get around this scarcity: (1) exploring whitespace, i.e., geographically underutilized parts of the spectrum, or (2) increasing spectral efficiency for a particular frequency bandwidth. Regarding (1), the FCC has adopted various strategies such as reverse-auctioning to buy unused TV bands, opening up spectrum in a fragmented fashion. Most current transceivers use fixed-frequency elements such as high order filters or circulators, preventing them from taking advantage of frequency hopping. The way to take advantage of (2) involves duplex communication. While time division duplexing (TDD) is still widely in use, it is not optimal in terms of maximizing throughput. Frequency division duplexing (FDD) enables full data capacity for both transmit and receive, and FDD deployments can provide greater coverage and achieve cell edge rates further from the base station. However, an FDD capable radio has its own challenges. Self-interference cancellation (SIC) is required since the transmitter signal (>20 dBm) is trillions of times larger than the receiver sensitivity level (<-90 dBm), which can saturate or damage the receiver. We want to be able take advantage of both of these methods (frequency agility and duplexing): our goal is to build a software-defined, frequency tunable, and FD/FDD capable transceiver. We use a distributed transmission line-based amplifier that has baseband gain and phase control to achieve SIC. Furthermore, we use a frequency selective source degeneration on the sub-PAs to mitigate the TX noise in the RX band.

#### BIOGRAPHY

Hazal Yuksel is currently a postdoctoral research associate in the CoSMIC Lab in the Electrical and Computer Engineering Department at Columbia. She received a master's degree and a PhD in electrical and computer engineering (ECE) from Cornell in 2015 and 2018, respectively, and a bachelor's degree in ECE from Duke in 2008. Her research focuses on biomedical applications of nonreciprocal duplexing networks. Her graduate research focused on the design and theoretical widely tunable, ultra-flexible, software-defined duplexing transceivers. She also investigated the consequences of discontinuities of IV curves in device modeling. She was a Facebook PhD Research Fellowship Finalist in 2016 and received the Outstanding Teaching Assistant Award from Cornell ECE in 2018.


## Amy Xian Zhang

**PHD STUDENT,** Massachusetts Institute of Technology



## Shanghang Zhang

**POSTDOCTORAL RESEARCHER,** Carnegie Mellon University


# Building Systems to Improve Online Discussion

Little has changed in the design of online discussion systems such as email forums and chat in the decades they have been available, even as discussions grow in size and scope. As a result, problems within online social platforms and communities have intensified, with issues ranging from scale and a lack of context leading to information overload to misinformation and malicious actors exacerbating polarization and harassment. My research involves designing and building tools to give users direct control over their information and experiences in online discussion systems. I focus on three major techniques and explore their effectiveness through the deployment of novel systems to real communities: (1) I develop summarization tools to help make sense of large discussions, such as Wikum, a collaborative wiki environment for summarizing threaded forums in a recursive fashion, and Tilda, a tool for collective note-taking while participating in group chat. (2) I build annotation tools to situate discussions in the context of what is being discussed. Tools include Eyebrowse, a system to bring social experiences to web browsing using in-page annotations and chat, and Nota Bene, a tool for students to mark up textbooks with comments conveying their emotional state. (3) I work on moderation tools to give users more control over the governance and delivery of messages, so that messages only go to those who want them. Tools include Murmur, a novel mailing list system that introduces fine-grained personal filtering, and Squadbox, a tool designed to help prevent online harassment.

### BIOGRAPHY

Amy X. Zhang is a fifth-year PhD student at MIT's Computer Science and Artificial Intelligence Laboratory (CSAIL), focusing on human-computer interaction and social computing. She is also a 2018-19 Fellow at the Berkman-Klein Center at Harvard. She received a master's degree in advanced computer science at the University of Cambridge in the U.K. on a Gates Fellowship. She received a bachelor's degree in computer science at Rutgers, where she captained the Division I women's tennis team. She interned twice at Microsoft Research and twice at Google Research and was a software engineer at a news startup. Her work received a best paper honorable mention at the CHI Conference on Human Factors in Computing Systems and has been featured by ABC News, BBC, CBC, The Verge, and *Business Insider*. Her research is supported by a Google Fellowship and a National Science Foundation Graduate Research Fellowship.

# Deep Understanding of Urban Mobility from Cityscape Webcams

Deep understanding of urban mobility is of great significance for many real-world applications such as urban traffic management and autonomous driving. We develop deep learning methodologies to extract vehicle counts from streaming real-time video captured by multiple low-resolution web cameras and construct maps of traffic density in a city environment. In particular, we focus on cameras installed in the Borough of Manhattan in New York City. The large-scale videos from these web cameras have low spatial and temporal resolution, high occlusion, large perspective, and variable environment conditions that cause most existing methods to lose their efficacy. To overcome these challenges, the thesis develops several techniques: a block-level regression model with a rank constraint to map the dense image feature into vehicle densities; a deep multi-task learning framework based on fully convolutional neural networks to jointly learn vehicle density and vehicle count; deep spatio-temporal networks for vehicle counting to incorporate temporal information of the traffic flow; and multi-source domain adaptation mechanisms with adversarial learning to adapt the deep-counting model to multiple cameras. To train and validate the proposed system, we have collected a large-scale webcam traffic dataset, CityCam, which contains 60 million frames from 212 webcams installed in key intersections. Of these, 60,000 frames have been annotated with rich information, leading to about 900,000 annotated objects. We believe this is the first and largest webcam traffic dataset with such large numbers of elaborate annotations. The proposed methods are integrated into the CityScapeEye system, which has been extensively evaluated and compared to existing techniques on different counting tasks and datasets with experimental results demonstrating the effectiveness and robustness of CityScapeEye.

### BIOGRAPHY

Shanghang Zhang received a PhD from Carnegie Mellon and a master's degree from Peking University. Her research interests include deep learning and computer vision. She has been working on large-scale traffic video analysis, vehicle detection and counting, salient object segmentation, domain adaptation, and Image Synthesis with GAN. She has received the Adobe Academic Collaboration Funding Qualcomm Innovation Fellowship (QInF), the Competition Finalist Award, and the Chiang Chen Overseas Graduate Fellowship. She serves as a reviewer for numerous publications, and has been an intern at Adobe Research.


## Yiying I. Zhu

**POSTDOCTORAL FELLOW,**
Georgia Institute of Technology



## Marinka Zitnik

**POSTDOCTORAL FELLOW,**
Stanford University


### Emerging Techniques in Contrast-Enhanced Ultrasound Imaging

Ultrasound — a real-time, high-resolution, portable, non-ionizing, and cost-effective imaging technique — continues to proliferate in the clinical environment. Unfortunately, contrast and consequently sensitivity and specificity of ultrasound imaging is relatively low compared to other imaging modalities. Therefore, ultrasound imaging is often combined with contrast agents, such as microbubbles, for more effective diagnosis and image-guided therapy. Previously, we developed a noninvasive "microsurgery" strategy using ultrasonic microbubble cavitation to reduce cardiac tissue for hypertrophic cardiomyopathy treatment. However, one major limitation of traditional microbubbles is their micrometer size, which constrains them inside vessels and prevents them from reaching their extravascular targets. To address this limitation, we introduced a combined ultrasound-photoacoustic (USPA) imaging system and dual-modality imaging contrast agent — optically triggered perfluorohexane nanodroplets (PFHnDs). Under laser irradiation, these liquid nanodroplets can vaporize into gaseous microbubbles. In liquid form, nanometer size PFHnDs can be delivered to compartments outside of vasculature, and once delivered, optically activated PFHnDs convert to microbubbles serving as contrast agents. The laser-induced vaporization of PFHnDs also provides photoacoustic contrast, which makes PFHnDs a dual USPA contrast agent. A unique feature of PFHnDs is that after activation, they stochastically condense to liquid droplets within a few milliseconds. The vaporization followed by recondensation can be triggered repeatedly upon each laser pulse, creating "blinking" ultrasound (and photoacoustic) signals. This enables super-resolution and contrast-enhanced ultrasound imaging. Furthermore, external manipulation of vaporization-recondensation dynamics can be used to increase imaging contrast or to assess the properties of tissue surrounding PFHnDs. As an example, I will present my recent work on using ultrasound imaging transmit pulse to manipulate dynamic behavior of optically vaporized PFHnDs. Finally, I will discuss the role of ultrasound together with light and nanotechnology in future research.

#### BIOGRAPHY

Yiying Zhu is a postdoctoral fellow in Georgia Tech's School of Electrical and Computer Engineering, where she works in the Ultrasound Imaging and Therapeutics Research Laboratory. Her research interest is to develop advanced imaging techniques using ultrasound with light and nanotechnology. She received a PhD in biomedical engineering (2017) and a master's degree in electrical engineering (2013), both from the University of Michigan, where she received an Endowment for the Development of Graduate Education Award. She received a bachelor's degree in electrical engineering from the University of Birmingham in the U.K. and a bachelor's degree in telecommunication engineering from Huazhong University of Science and Technology in China.

### Next-Generation Algorithms for Complex and Heterogeneous Biomedical Networks

Complex systems are ubiquitous in biology and medicine, where they encode connectivity patterns at all scales. For the first time, technology has enabled large-scale, objective measurements of these systems, such as the genetic fingerprinting of each of 37 trillion cells in the human body. Massive datasets from such technologies can accelerate discovery and create a new era of precision medicine with highly targeted drugs and treatments. We study large complex interconnected systems at all scales, from interactions of proteins in a cell to interactions between patients in a society. Mathematical machinery that is central to these complex systems is machine learning on graphs. We define novel graph convolutional deep networks that learn data transformation functions, which we use to map nodes to embeddings, points in low-dimensional vector spaces. These deep models can embed nodes, subgraphs, and highly multi-relational graphs. We show that these methods enable practical biomedical applications of high impact, such as predicting which diseases drugs would treat, identifying which side effects patients will get from a specific drug combination, and predicting which roles a protein has in a particular human tissue. We also present a mathematical understanding of graphs in the context of biology and medicine, where such algorithms can uncover new biology and be used to accelerate knowledge discovery. We have developed methods that generate actionable insights in the form of hypotheses derived from biomedical data. We show that these network science algorithms can be used to enhance networks and automatically identify hypotheses with the highest prospect, ranging from drug repurposing in medical drug networks to cell-type discovery in cell interaction networks.

#### BIOGRAPHY

Marinka Zitnik is a postdoctoral fellow in computer science at Stanford University, working with Professor Jure Leskovec. She received a PhD in computer science from the University of Ljubljana in 2015 while also conducting research at Imperial College London, University of Toronto, Baylor College of Medicine, and Stanford University. She received a bachelor's degree in computer science and mathematics in 2012. Her research interests are in computational biology, machine learning, and network science. Her work has received several best paper, poster, and research awards at major computational biology conferences, including ISMB, CAMDA, RECOMB, and BC2. She received a Google Anita Borg Memorial Scholarship, was selected as a Young Fellow at the Heidelberg Laureate Forum, and received the Jozef Stefan Golden Emblem Prize.

RISING STARS IN EECS ALUMNAE:

# WHERE ARE THEY NOW?

**Between 2012 and 2017, more than 300 early-career women in EECS attended Rising Stars career-development workshops at MIT, Carnegie Mellon, UC Berkeley, and Stanford.**

As of June 2018, fully 30 percent of those participants were in faculty positions. Slightly more than 20 percent were working in industry; about the same percentage were in postdoctoral positions. Slightly less than 20 percent were students, while the rest were in other academic roles.

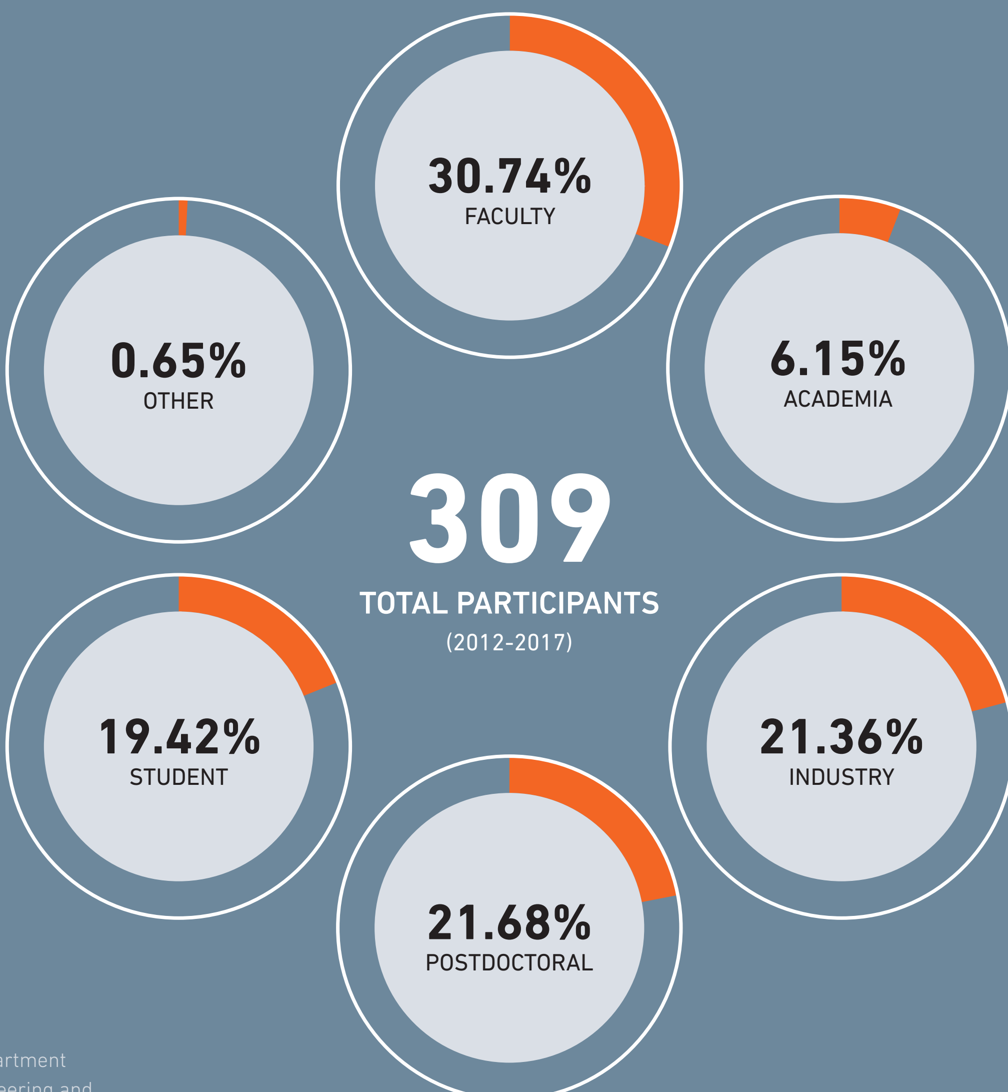


**SOURCE:** MIT Department of Electrical Engineering and Computer Science

## RISING STARS IN EECS ADMINISTRATIVE TEAM

### Dorothy Curtis
SENIOR RESEARCH SCIENTIST
Computer Science and Artificial Intelligence Laboratory, MIT

### Myron "Fletch" Freeman
MANAGER OF DEPARTMENTAL COMPUTING, EECS

### Irene Yong Rong Huang
ADMINISTRATIVE OFFICER, EECS

### Rachida Kernis
ADMINISTRATIVE ASSISTANT, EECS

### Anne Stuart
COMMUNICATIONS OFFICER, EECS

### William Tilden
EVENTS COORDINATOR, EECS

### Myung-Hee Vabulas
PROGRAM COORDINATOR, EECS

Special thanks to **SUZANA LISANTI** for assistance with this year's Rising Stars website.

Design: Opus
Printing: Puritan Capital

---

## RISING STARS IN EECS WORKSHOPS

**2018:** Massachusetts Institute of Technology

**2017:** Stanford University

**2016:** Carnegie Mellon University

**2015:** Massachusetts Institute of Technology

**2014:** University of California Berkeley

**2013:** Massachusetts Institute of Technology

**2012:** Massachusetts Institute of Technology

For links to earlier Rising Stars workshop websites, please visit risingstars18-eecs.mit.edu.

## RISING STARS IN EECS SPONSORS

MIT School of Engineering

MIT EECS
ELECTRICAL ENGINEERING AND COMPUTER SCIENCE

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

LIDS

MTL MICROSYSTEMS TECHNOLOGY LABORATORIES

rle AT MIT

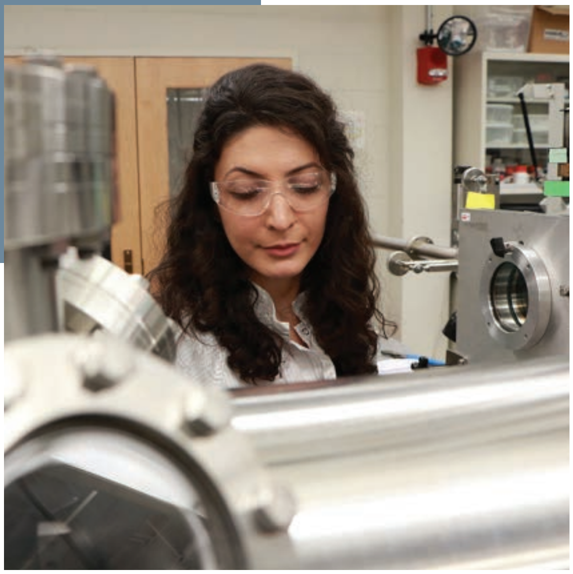

"The Rising Stars Workshop was very educational and inspiring for me—**a unique platform to learn from the experiences of the leading researchers** with diverse and successful academic careers, and to form connections with a community of peers who aspire to be the future leaders in various fields of EECS."

**FARNAZ NIROUI**
**Assistant Professor of EECS, MIT**
**Rising Stars, 2015**

"The Rising Stars workshop was an amazing opportunity to **chat with top professors** about my research, to **learn from insiders** about how to thrive in an academic career, and to **meet the next wave** of world-class researchers in EECS."

**TAMARA BRODERICK**
**ITT Career Development Assistant Professor of EECS, MIT**
**Rising Stars, 2013**

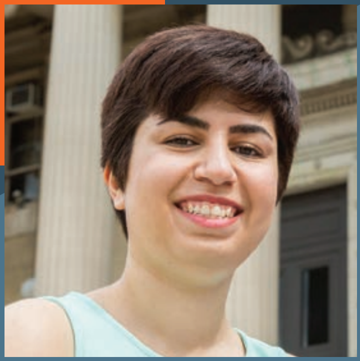

"The Rising Stars workshop gave me a unique opportunity to meet and network with professors as well as future colleagues. **Informative and insightful talks and panels have already helped me a lot to find my path** and I keep going back to what I learned from the workshop as I prepare for my new job."

**NEGAR REISKARIMIAN**
**PhD Candidate in Electrical Engineering, Columbia University**
**Assistant Professor of EECS, MIT (effective July 2019)**
**Rising Stars, 2017**

PHOTO OF FARNAZ NIROUI, PAGE 46:
Denis Paiste, MIT Materials Research Center

PHOTO OF TAMARA BRODERICK, ABOVE:
Lillie Paquette, MIT School of Engineering

BACK COVER PHOTO FROM RISING STARS 2015:
Gretchen Ertl

# CONTACT

Asu Ozdaglar

DEPARTMENT HEAD
MIT DEPARTMENT OF ELECTRICAL
ENGINEERING AND COMPUTER SCIENCE

MIT SCHOOL OF ENGINEERING
DISTINGUISHED PROFESSOR OF
ENGINEERING

OFFICE: 617-253-4600
DIRECT: 617-253-4602
asuman@mit.edu

DEPARTMENT OF ELECTRICAL
ENGINEERING AND COMPUTER SCIENCE

Massachusetts Institute of Technology
77 Massachusetts Avenue, 38-401
Cambridge, MA 02139

eecs.mit.edu
risingstars18-eecs.mit.edu