Research Statement

Research overview

My work focuses on *algorithmic* aspects of machine learning (ML). Concretely:
(1) I design algorithms for acquiring useful information from massive amounts of data, and
(2) I develop theoretical tools to analyze why, when, and how well these algorithms work.

Machine learning will continue to dramatically impact society in the coming decades. ML algorithms already influence how we receive our information and how we interact with each other. In the very near future, ML will power our cars, trucks, and bikes; will control our power grid; and will determine what types and doses of medicine to give to every patient.

However, numerous foundational questions in machine learning remain unanswered:
(1) Certifying whether a given ML algorithm makes a good enough decision (according to any reasonable definition of "goodness") is beyond the reach of most existing analysis methods.
(2) Designing a machine learning algorithm to operate under resource constraints (such as computation time, or energy, or memory) is a difficult task.
(3) Ensuring that a ML system is robust in the presence of adversaries is increasingly a major obstacle to its deployment in everyday life.
(4) Finally, scaling up ML to the big data regime remains a significant outstanding challenge.

As an assistant professor (starting at Iowa State in August 2015, and subsequently at NYU Tandon since November 2019), I have commenced upon an ambitious research program that aims to resolve the challenges listed above. It is clear by now that a broad interdisciplinary effort is required to get impactful solutions. Therefore, my research fuses ideas from computer science, electrical engineering, and applied mathematics. The specific techniques I use to address the above foundational ML problems blend insights from *signal processing* (combining Fourier/wavelet analysis and compressive sensing) and *discrete algorithm design* (combining approximation and combinatorial algorithms along with first-order iterative methods).

While my core focus is on theoretical and algorithmic aspects of ML, the work of my research group is greatly enriched by a diverse set of collaborators. I have worked with optics experts on imaging problems; with mechanical engineers on problems in controls and robotics; with material scientists on microstructure modeling; with transportation engineers on roadway analytics and autonomous navigation; with neuroscientists on cognitive modeling; and with security researchers on privacy and robustness of machine learning systems. My research has been supported by two prestigious NSF early career awards: the CRII Award in 2016, and the CAREER Award in 2018. I have also been supported by other grants/gifts from the NSF, the NGA, the NIH, DARPA, ARPA-E, Department of Energy, Air Force Research Lab, Toyota Research, and the Black and Veatch Foundation.

Research themes

Much of my work has focused upon designing *fast, provable, and practical algorithms* for various machine learning and signal processing problems. I organize my contributions by topic below.

Learning discrete structures. At the heart of most modern data processing algorithms lies a key insight: despite its apparent size, the aggregate data obeys some notion of concise, *low-dimensional structure*. (A popular structural assumption is *sparsity*, where data is well-modeled by basis representations that contain only a few nonzero coefficients.)

In collaboration with Piotr Indyk (MIT) and Ludwig Schmidt (UW), I designed the fastest algorithms for discovering *discrete* low-dimensional structures in high-dimensional data for a variety of special cases. Examples include combinatorial structures such as trees [1,2], paths [3], and connected subgraphs [4]. Our algorithms were the first to achieve (near) linear running time, which is the gold standard for algorithmic efficiency in big data applications. Our paper [5] subsumed all the above methods into a single graph-structured learning framework, and was awarded the **Best Paper Award** at ICML 2015.

Our motivation for these algorithms came from two application areas. The first area where my algorithms can be applied is the well-known signal processing paradigm of *compressive sensing*, which states that if a signal admits a sparse representation, then it can be robustly acquired by only taking a small number of randomized measurements (or samples). In collaboration with Richard Baraniuk (Rice), Volkan Cevher (EPFL), and Marco Duarte (UMass), I had previously developed a framework called *model-based compressive sensing,* which enables structured signal acquisition at *statistically optimal* sampling rates while simultaneously guaranteeing *near-linear* recovery time [6,7]. My work on recovering neural spike trains at optimal rates was awarded the **Best Student Paper Award** at SPARS 2009, and our combined work on this topic has been cited over 2000 times (according to Google Scholar). This paradigm has been fruitfully applied in diverse applications such as hyperspectral imaging, MRI, high-speed video, and biosensing.

The second application area arose from *inference over graphs*. This corresponds to situations where we reliably know the edge structure of a given graph and the challenge is to leverage this knowledge to enable improved decision- making. This formulation impacts a variety of applications, including: finding small communities in a social network [5]; rapid identification of congestion in a roadway network [22]; or even segmenting high-resolution seismic image data [26,27]. I describe additional applications of these algorithmic techniques in greater detail below.

Learning matrix representations. Linear models for data (such as sparsity defined for a fixed basis) is a somewhat simplistic, first-order assumption. It is often more fruitful to leverage second-order *nonlinear* structures, such as low-rankness in matrix data. Again, an interesting tension arises: while matrix models are often statistically beneficial, discovering them requires significantly more *computational* resources. To alleviate this tradeoff, I have developed provably fast (near-linear time) algorithms for learning low-rank representations [13]; the first provably accurate methods for double-sparse coding [14]; and the first provably polynomial-time algorithm for learning sparse factor models from incomplete samples [15]. I presented the latter algorithm in a **keynote talk** at the Midwest Machine Learning Symposium in 2018. In a related stream of work, I have worked on accelerating matrix estimation problems using techniques from numerical linear algebra, such as learning graphical models [23], rank-constrained optimization [24], and distributed matrix completion [25]. All these theoretical results are building blocks towards understanding how much training data is required to learn reliable ML models for recommender systems.

Learning neural networks, provably. While the basic ideas underpinning neural networks are by now classical, the emergence of deep neural networks has resulted in dramatic advances in numerous applications including computer vision, natural language processing, and robotics, to name a few. Nevertheless, most neural learning algorithms lack correctness and efficiency guarantees, and from a theoretical standpoint, important open questions remain. What are the precise statistical and computational requirements for training a neural network model? And once we have trained a model, can we certify that we have succeeded?

I have established a series of results that address neural network learning algorithms from a theoretical lens. In particular, I developed the fastest provably accurate algorithms for learning (shallow) polynomial neural networks [16]. The high level idea is to use a particular linearization trick that enables us to borrow our previous results on fast matrix estimation algorithms. This work was given the **Best Poster Award** at the Midwest Machine Learning Symposium in 2018. Using a similar linearization trick, I also developed an early set of results for provably learning shallow networks equipped with Rectified Linear Unit (ReLU) activations [17]. Similar insights lead to provable upper bounds for meta-learning problems [35] and distributed learning [19].

The last several years have witnessed a large number of interesting theoretical advances in provable network learning, but the large majority of these advances have focused on the supervised setting. A particular research focus of mine has been to develop analogous provable guarantees for training neural networks in the *unsupervised* setting. In collaboration with Raymond Wong (Texas A&M), leveraging our previous techniques on sparse coding, I provided the first set of provably accurate algorithms for learning (shallow, narrow) auto-encoder networks [18]. More recently, we provided the first analysis of training shallow (but wide) auto-encoder networks [29]. An expansion of our results to more general unsupervised learning models would have significant impact.

Security aspects of neural networks. Despite their tremendous promise in diverse applications, deep neural networks have been shown to be rather brittle when subjected to malicious inputs. In collaborations with Soumik Sarkar (Iowa State) and Siddharth Garg (NYU), I have expanded the study of this vulnerability. I first showed that even *benign* alterations to the input (such as changing the illumination of an outdoors scene, or adding facial hair to a mugshot) can have catastrophic implications for deep learning-based image classifiers [30]. A more thorough understanding of this effect may affect how we design vision systems for autonomous navigation. I also exhibited similar brittle behavior for feedback-based neural network controllers used in reinforcement learning [33].

Deep neural networks in materials design. The large majority of machine learning research has focused on developing models tested on pristine benchmark datasets acquired for computer vision and text analytics applications. However, scientific and engineering applications tend to be far more messy. In particular, design and manufacturing applications pose unique barriers to widespread deployment of machine learning, specifically related to the paucity - as well as heterogeneity - of training data. A major component of my recent research has been to explore ML innovations in this area, and my guiding hypothesis has been that *knowledge of the underlying physics* may help alleviate these barriers. I have begun to validate this hypothesis in computational *materials design,* specifically those encountered in the context of manufacturing photovoltaic devices.

In collaborations with several domain experts such as Baskar Ganapathysubramanian (Iowa State), Duane Johnson (Iowa State), Zhenan Bao (Stanford), and Ross Larsen (NREL), I have commenced upon developing ML-based methods for (computational) materials design. We have introduced novel physics-informed techniques that use generative adversarial networks (GANs) to simulate microstructures that satisfy specified target properties. I presented our approach at a **keynote presentation** in an NIH Workshop on Machine Learning and Multiscale Modeling in 2019. Overall, our results indicate that the models that we learn have considerable promise to not just capture salient microstructure features, but also discover the physics underlying their generation [32,34]. As a result of our work, we developed benchmark microstructure datasets that won the 2nd prize at the **Open Data Challenge at the 2019 MRS Annual Meeting** in 2019.

Transportation analytics and driver behavior modeling. I am very interested in building and studying ML models that *continuously* interact with large scale human-engineered systems. From the engineering perspective, how best should ML models assimilate continuous streams of raw human-generated data? But more importantly, what is the societal impact of the numerous feedback loops that may arise due to the various levels of predictions? The area of transportation analytics offers a unique, alternative application domain for exploring such questions. In collaboration with Anuj Sharma (Iowa State), I have designed novel inference methods for incident detection problems in transportation analytics. The goal is to achieve rapid detection of incidents and/or congestion events in highways [28]. A key innovation of ours has been to extend the scope of neural tracking-based algorithms (such as DeepSORT) to the semi-supervised setting [22].

Similar problems in the intersection of transportation and human behavior modeling can be addressed via ML models trained on physiological data. In collaboration with Anuj Sharma, Soumik Sarkar, and Matthew Rizzo (University of Nebraska), I developed ML models that both correlate driver behavior features (such as control actions taken at a busy street intersection) with physiological measurements (such as glucose levels and heart rate), conditioned on whether the driver under study exhibited had been pre-diagnosed with pathologies such as type-1 diabetes [36].

Inverse problems and computational imaging. Many of my algorithmic contributions can be immediately applicable to various forms of computational imaging and other inverse problems. My work on graph-based inference, in collaboration with researchers at Shell International E&P, has enabled quasi-linear time algorithms for detecting anomalies in pre- and post-migrated seismic image data [27,28]. Along with Richard Baraniuk (Rice), Aswin Sankaranarayanan (CMU), Piotr Indyk (MIT), and Kevin Kelly (Rice), I showed that techniques from convex and combinatorial optimization lead to the design of novel sample-efficient mechanisms for image acquisition [37].

The role of efficient and provably accurate ML-based approaches becomes even more important in the context of *nonlinear* inverse problems [8]. A particularly interesting class of imaging problems, including X-ray crystallography and ptychography, can be unified into the broad category of "phase retrieval". In collaboration with Namrata Vaswani (Iowa State) and Salman Asif (UC Riverside), I have developed novel, sample-efficient phase retrieval algorithms using sparsity [9,10], low-rank [11,41], and neural network priors [38]. The latter work builds upon my generic approach for provably solving inverse problems using generative models [20,21].

The above algorithms require that enough training samples are available to pre-train a useful generative model; however, in applications such as medical imaging, training data might be scarce. I have developed provable inverse imaging approaches in the *data-free* regime, showing that *untrained* neural network priors (with suitable inductive bias encoded in the architecture) can also be fruitfully used to solve nonlinear inverse imaging problems [39]. In the first result of this kind, I have demonstrated the efficacy of the untrained network prior for solving high-dynamic range (HDR) image reconstruction problems encountered in low-light imaging scenarios [40].

Future Directions

As the impact of ML systems on society continues to grow, so too will the importance of resolving its foundational aspects, understanding how ML systems work, and making ML reliable and secure. At NYU, I have already begun exploring new avenues within the above broad research agenda. One direction of interest has been to examine ML security at the systems level and to propose novel ML that offer privacy guarantees (in the cryptographic sense) *by design*. In an emerging collaboration with Siddharth Garg and Brandon Reagan (NYU), by blending new ideas from combinatorial optimization, we have been successful in devising deep learning architectures that support private inference but with an order-of-magnitude improvement in latency.

A second direction of interest is to continue pushing the boundaries of scientific machine learning. In separate emerging collaborations with Christopher Musco (NYU) and Nidhi Rajput (Stony Brook) we have begun to address challenges of machine learning models in energy-related design applications (both from the theoretical as well as applied perspectives). The key in both projects is to take advantage of the physics of the underlying problem in order to perform better predictions with fewer (either computational or experimental) resource costs.

Finally, while my primary research focus has been in engineering domains, I am also interested in leveraging my core ML research for educational innovation. In an emerging collaboration with Nasir Memon (NYU), we have begun exploring ML-based approaches for content curation and recommender systems in the context of learning management systems (LMS) design. The goal is to leverage ML advances for improving student outcomes as well as course design for instructors. Preliminary results on the NYU Tandon CS Bridge program have shown considerable promise.

References

1. C. Hegde, P. Indyk, and L. Schmidt. "Nearly linear-time model-based compressive sensing". In Proc. Intl. Colloquium on Automata, Languages, and Programming (ICALP), July 2014.
2. C. Hegde, P. Indyk, and L. Schmidt. "A fast approximation algorithm for tree-sparse recovery". In Proc. IEEE Int. Symp. Inform. Theory (ISIT), June 2014.
3. C. Hegde, P. Indyk, and L. Schmidt. "Approximation-tolerant model-based compressive sensing". In Proc. ACM Symp. Discrete Alg. (SODA), Jan. 2014.
4. C. Hegde, P. Indyk, and L. Schmidt. A fast adaptive variant of the GW algorithm for the Prize-Collecting Steiner Tree problem. DIMACS Workshop, Dec. 2014.
5. C. Hegde, P. Indyk, and L. Schmidt. A nearly linear-time framework for graph-structured sparsity. In Proc. Int. Conf. Machine Learning (ICML), July 2015.
6. C. Hegde, P. Indyk, and L. Schmidt. Approximation algorithms for model-based compressive sensing. IEEE Trans. Inform. Theory, 61(9):5129–5147, Sept. 2015.
7. R. Baraniuk, V. Cevher, M. Duarte, and C. Hegde. Model-based compressive sensing. IEEE Trans. Inform. Theory, 56(4):1982–2001, Apr. 2010.
8. M. Soltani and C. Hegde. "Fast algorithms for demixing signals from nonlinear observations". IEEE Trans. Sig. Proc., 65(16):4209–4222, Aug. 2017.
9. G. Jagatap and C. Hegde. "Fast sample-efficient algorithms for structured phase retrieval". In Adv. Neural Inf. Proc. Sys. (NeurIPS), Dec. 2017.
10. G. Jagatap and C. Hegde. "Towards sample-optimal methods for solving random quadratic equations with structure". In Proc. IEEE Int. Symp. Inform. Theory (ISIT), June 2018.
11. Z. Chen, G. Jagatap, S. Nayer, C. Hegde, and N. Vaswani. "Low rank Fourier ptychography". In Proc. IEEE Int. Conf. Acoust., Speech, and Signal Processing (ICASSP), Apr. 2018.
12. V. Shah, M. Soltani, and C. Hegde. "Reconstruction from periodic nonlinearities, with applications to HDR imaging". In Proc. Asilomar Conf. Sig. Sys. Comput., Nov. 2017.
13. C. Hegde, P. Indyk, and L. Schmidt. "Fast recovery from a union of subspaces". In Adv. Neural Inf. Proc. Sys. (NeurIPS), Dec. 2016.
14. T. Nguyen, R. Wong, and C. Hegde. "A provable approach for double-sparse coding". In Proc. AAAI Conf. Artificial Intelligence, Feb. 2018.
15. T. Nguyen, A. Soni, and C. Hegde. "On learning sparsely used dictionaries from incomplete samples". In Proc. Int. Conf. Machine Learning (ICML), Jul. 2018.
16. M. Soltani and C. Hegde. "Towards provable learning of polynomial neural networks using low-rank matrix estimation". In AISTATS, Apr. 2018.
17. G. Jagatap, C. Hegde. "Linearly Convergent Algorithms for Learning Two-Layer ReLU Networks". In Proc. Int. Symp. Inform. Theory (ISIT), June 2018.
18. T. Nguyen, R. Wong, and C. Hegde. "On the Dynamics of Gradient Descent for Autoencoders". In Proc. Artificial Intelligence and Statistics, April 2019.
19. Z. Jiang, A. Balu, C. Hegde, and S. Sarkar. "Collaborative deep learning over fixed topology networks". In Adv. Neural Inf. Proc. Sys. (NIPS), Dec. 2017.
20. V. Shah and C. Hegde. "Solving linear inverse problems using GAN priors: An algorithm with provable guarantees". In ICASSP, Apr. 2018.

21. C. Hegde, "Algorithmic Aspects of Inverse Problems Using Generative Models", Proc. Allerton Conf. Comm. Control, and Computing, Oct. 2018.
22. P. Chakraborty, C. Hegde, and A. Sharma. "Freeway incident detection from cameras: A semi-supervised learning approach". Intelligent Transportation Systems (ITSC), Nov. 2018.
23. M. Soltani and C. Hegde. "Fast algorithms for learning latent variables in graphical models". In Proc. ACM KDD Workshop on Mining and Learning with Graphs (KDD MLG), Aug. 2017.
24. M. Soltani and C. Hegde. "Fast low-rank matrix estimation for ill-conditioned matrices". In Proc. IEEE Int. Symp. Inform. Theory (ISIT), June 2018.
25. C. Hubbard and C. Hegde. "Parallel computing heuristics for matrix completion". In Proc. IEEE Global Conf. Signal and Image Processing (GlobalSIP), Nov. 2017.
26. L. Schmidt, C. Hegde, P. Indyk, L. Lu, X. Chi, and D. Hohl. "Seismic feature extraction using Steiner tree methods". In ICASSP, Apr. 2015.
27. L. Schmidt, C. Hegde, P. Indyk, J. Kane, L. Lu, and D. Hohl. "Automatic fault localization using the Generalized Earth Movers Distance". In ICASSP, May 2014.
28. P. Chakraborty, C. Hegde, and A. Sharma. "Trend filtering in network time series with applications to traffic incident detection". NIPS Time Series Workshop (TSW), Dec. 2017.
29. T. Nguyen, R. Wong, and C. Hegde. "Benefits of Jointly Training Autoencoders: A Neural Tangent Kernel Approach", IEEE Transactions on Information Theory, 2021.
30. A. Joshi, A. Mukherjee, S. Sarkar, and C. Hegde. "Semantic Adversarial Attacks: Parametric Transformations That Fool Deep Classifiers", ICCV,, November 2019.
31. G. Jagatap and C. Hegde, "Algorithmic Guarantees for Inverse Imaging with Untrained Network Priors", Neural Information Processing Systems (NeurIPS), 2019.
32. A. Joshi, M. Cho, V. Shah, S. Sarkar, and C. Hegde, "InvNets: Encoding geometric and statistical invariances in deep generative models." AAAI, Feb. 2020.
33. X. Lee, S. Ghadai, K. Tan, C. Hegde, S. Sarkar, "Spatiotemporally Constrained Action Space Attacks on Deep Reinforcement Learning Agents", AAAI, Feb. 2020.
34. X. Lee, et al. "Fast inverse design of microstructures via generative invariance networks", Nature Computational Science, Apr. 2021.
35. M. Cho, M. Soltani, and C. Hegde. "One-Shot Neural Architecture Search via Compressive Sensing", ICLR Workshop on Neural Architecture Search, May 2021.
36. P. Chakraborty et al, "Quantifying vehicle control from physiology in Type-1 diabetes", Traffic Injury Prevention (TIP), Nov. 2019
37. C. Hegde, A. Sankaranarayanan, W. Yin, R. Baraniuk, "A convex approach for learning near-isometric linear embeddings", IEEE Transactions on Signal Processing, 2015.
38. R. Hyder, V. Shah, S. Asif, C. Hegde, "Alternating minimization for phase retrieval using generative priors", ICASSP 2019.
39. G. Jagatap and C. Hegde, "Algorithmic Guarantees for Inverse Problems Using Untrained Network Priors", NeurIPS 2019.
40. G. Jagatap and C. Hegde, "HDR Imaging using Deep Image Priors", ICASSP 2020.
41. G. Jagatap, Z. Chen, S. Nayer, N. Vaswani, and C. Hegde, "Sub-diffraction super-resolution imaging for structured data", IEEE Trans. Computational Imaging, Oct. 2019.

Teaching Statement

Traditional barriers in scientific education are quickly breaking down, and the modern engineering student needs to acquire a diverse array of skills in order to effectively compete and excel in a professional career. For engineers, these include a broad set of core technical skills in related fields such as computer science, software engineering, mathematical optimization, and statistical reasoning. These also include more general soft skills such as the ability to work in a large team, the ability to realize conceptual ideas into tangible prototypes and/or software, and the development of effective skills in technical communication

The core elements of my pedagogical approach mirror the basic principles of my research: *attention to a **breadth** of technical topics*, and *an emphasis on inter-disciplinary **multi-dimensional aspects***. I use the principles throughout my teaching process, detailed as follows.

***Breadth of teaching contributions*.** I have taught several courses across MIT, Iowa State, and NYU. I have taught 4 undergraduate courses (EE, CE, and CS) and three graduate courses (in ECE and CSE).

At NYU, I have made the following teaching contributions:
- *ECE-GY 6143: Introduction to Machine Learning.* In Spring and Summer 2020, I taught the graduate introductory course to Machine Learning in the ECE department. I created a new set of lecture notes[5] as well as several new lab exercises. In all, I instructed more than 300 students.
- *CSE-GY 9223/ECE-GY 9123: Deep Learning.* In Fall 2020 and Spring 2021, I revived the special-topics graduate course on Deep Learning at NYU Tandon. This course has recently been approved to become permanent, and has served more than 250 students.

For my contributions, I have been nominated for the **Jacobs Award** in Teaching Excellence in 2021.

At Iowa State, I made the following teaching contributions:
- *EE/CPRE 529: Data Analytics for Electrical and Computer Engineering*. I designed a new course focusing on foundational principles of massive data analysis, and applications of these principles to solve engineering challenges. This course has now become a staple in the ECE grad curriculum.
- *CPRE 310: Theoretical Foundations of Computer Engineering*. This is a core junior-level course in the CPRE and SE undergraduate curriculum. I revamped this course significantly, and initiated a new weekly recitation section which reaffirms the lecture material with hands-on problem solving.

For my contributions at Iowa State, I was awarded two teaching awards: the **Warren Boast Award** for Undergraduate Teaching in 2016, and the **Boast-Nilsson Award** for Educational Impact in 2018.

Prior to this, as an instructor in Computer Science at MIT, I co-taught the following courses:
- *MIT 6.042: Mathematics of Computer Science.* In Spring 2014, along with Albert Meyer I co-taught a core sophomore level course that provides an overview of several concepts in discrete mathematics pertinent to computer science and engineering that are not covered in a mathematics curriculum.

- *MIT 6.006: Introduction to Algorithms*. In Spring 2015, along with Piotr Indyk and Yuan Zhou I co-taught a core junior level course that covered common algorithmic paradigms, data structures, and hands-on problem solving.

Multi-dimensional aspects. As part of my teaching philosophy, I try to target a few simple objectives. These include:

Adapting to my students' inter-disciplinary backgrounds. Classes can be rather heterogeneous. Students tend to forget what they have learned previously. I start all my lectures with a complete recap of what we have done so far; this helps bring everyone on to the same page. I strictly use pen-and-whiteboard (nowadays, an Apple Pencil and an iPad whiteboard) in my classes. This helps me adapt my teaching pace to my students' rate of understanding far easier than, say, a rigid format such as Powerpoint slides. This also helps me draw pictures on the fly, which can be helpful for students who think visually like I do.

Preparation. I try to be meticulous in my preparation for each lecture. Most of what I say in class, I prepare in advance in the form of typeset lecture notes. This helps in slowing down and clarifying my thinking; reduces errors; and leads to easy dissemination of class material to the students. I also post my notes (along with any associated source code) early, and students in my classes have consistently appreciated the presence of these notes as useful learning aids during lecture. All my course materials can be found freely available online [Refs 1-7].

Inclusivity and respect. I cherish interaction with *all* students. I am continuously available to everyone in my class via message boards; Piazza claims that average instructor response time to questions/posts during my summer Machine Learning course was less than 30 minutes. I constantly seek verbal feedback from students during conversations and office hours. I try my best to learn my students' names and backgrounds even in very large classes. I consciously seek out opportunities to engage with, and mentor, students from minority backgrounds [Refs 7-9]. In the end, my job as an instructor is not merely to achieve a set of educational objectives. It is to engage with smart individuals and motivate them to become better thinkers, scientists, and engineers.

References

[1] C. Hegde and A. Kamal, "Theoretical Foundations of Computer Engineering", Lecture Notes..
[2] C. Hegde, "Principles of Data Analytics", Lecture Notes.
[3] OpenStax CNX, http://www.cnx.org.
[4] Mini-course on Adversarial Machine Learning, notes online at Fall 2018 Talks.
[5] C. Hegde, "Introduction to Machine Learning", Lecture Notes. Available online since July 2020.
[6] C. Hegde, "Deep Learning", Lecture Notes. Available online since July 2020.
[7] I was a mentor for the ECSEL program: https://ecsel.ece.iastate.edu/.
[8] I was a mentor for the McNair program: https://www.mcnair.iastate.edu/.
[9] I was a mentor for the ESPI NYC weekend program: https://espi.nyc/