Ramchandran Muthukumar, Ph.D Student,
Department of Computer Science, Johns Hopkins University
rmuthuk1@jhu.edu | (+1) 443-541-1140 | LinkedIn

Research Interests

Recent research has indicated the fragility of state-of-the-art machine learning models in the presence of malicious adversaries. Imperceptible perturbations/corruptions on a test image are often sufficient to foil prediction. The success of an adversary is determined by the threat model, the trained deep network, the network architecture and finally the data distribution. My work seeks to answer the two central questions of Adversarial Robustness - generalization in the presence of adversary and certified robustness.

I am also interested in exploiting structure in data to develop efficient algorithms (eg. sparsity, low rank). My first research project exploited low rank structure in PDE solutions and current state-of-the-art sketching techniques to design memory efficient optimization algorithms. My latest work exploits sparsity in neural activation patterns to demonstrate a tighter notion of local Lipschitzness of neural networks.

Education & Research Experience

Johns Hopkins University July 2019 - present
Ph.D Student, Department of Computer Science
Advisors : Jeremias Sulam

Cornell University July 2017 - May 2019
Research Assistant, Department of Operations Research and Information Engineering
Advisors : Madeleine Udell, Drew Kouri (Sandia National Lab.)

Birla Institute of Technology and Science Pilani, Goa Campus July 2013 - May 2018
Master of Science in Mathematics,
Bachelor of Engineering in Computer Science

Preprints

- Adversarial robustness of sparse local Lipschitz predictors (Arxiv 2022)
  R. Muthukumar, J.Sulam
  - Input-specific pruning of neural networks helps provides tighter estimation of local Lipschtz constant.
  - Provides tighter robustness certificates leveraging the sparse activation patterns of ReLu feedforward networks.
  - Improved robust generalization bound that scale favourably with depth.

Published Manuscripts

- Adversarial Robustness of Supervised Sparse Coding (NeurIPS 2020)
  J. Sulam, **R. Muthukumar**, R. Arora
  - Demonstrates bounds on robust generalization error for supervised dictionary learning
  - Provides tighter deterministic robustness certificates leveraging sparse structure.
- Randomized Sketching Algorithms for Low-Memory Dynamic Optimization (SIAM-OPT *2021*)
  R. Muthukumar, D. Kouri and M.Udell
  - Builds on cutting edge randomized sketching algorithms to perform low memory PDE optimization
  - Provably convergent optimization algorithms based on Inexact Trust Region methods.

Software Projects

- Presolve Routines for LP and SDP @ Google Summer of Code 2016
  Mentor - Dr. Madeleine Udell
  - Implemented fast Presolving algorithms to speed up LP optimization in scientific computing language - Julia
  - Benchmarked against existing solvers for speed and efficiency
- Python Wrapper for LowRankModels.jl @ Data-Driven Discovery of Models DARPA Program 2018
  Mentor - Dr. Madeleine Udell
  - Implementing Python wrapper for the julia package LowRankModels.jl within the D3M framework
  - Testing for imputation problems on standard datasets.

- Machine Translation : Natural Language to SQL @ Fall 2020
  Mentor - Dr. Philip Koehn
  - Reproduce SOTA models (such as Seq2Sql) on WikiSQL data set.
  - Explore data augmentation techniques for improved performance : generating equivalent synthetic queries using synonyms sourced from WordNet (such as in Tweet2Vec)
  - Benchmarked against existing models for accuracy and efficiency

Awards

Summer 2021	MINDS Summer Fellowship in Data Science
Spring 2022	MINDS Spring Fellowship in Data Science

Talks and Conferences

Summer 2022	Adversarial Robustness of sparse local Lipschitz predictors, Poster @ Sparsity in Neural Networks Workshop 2022
Spring 2022	Adversarial Robustness of sparse local Lipschitz predictors, Invited Talk @ MINDS Retreat 2022
Fall 2020	Adversarial Robustness of Supervised Sparse Coding, Poster @ NeuRIPS 2020
Spring 2020	Randomized Sketching Algorithms for Low Memory Dynamic Optimization, Invited Student Guest speaker @ CSL Student Conference, UIUC Invited Talk @ East Coast Optimization Meeting 2020, GMU (cancelled due to covid)
Summer 2019	Sketching Algorithms for Approximate Gradients @ JuliaCon 2019
Summer 2016	Presolving Algorithms for Optimization @ JuliaCon 2016

Teaching Assistantship

Fall 2021	Sparse Representations in Computer Vision and Machine Learning
Fall 2020	Sparse Representations in Computer Vision and Machine Learning

References

Jeremias Sulam, Madeleine Udell, Drew Kouri