# RAHUL ANAND SHARMA

Email: rahulanandsharma@gmail.com
Website: https://rahul-anand.github.io/
Phone: 412-708-9085
Address: Sunnyvale, USA

**Research Interests:** Large Scale Systems and Networking, Data Analysis, Large Language Models and Generative AI

## PROFESSIONAL SUMMARY

Applied Scientist with expertise in Machine Learning, Security, and Large-Scale Distributed Systems. Proven track record of developing innovative solutions in AI-powered debugging, network security, and machine learning. Published researcher with significant contributions to academic and industry.

## APPOINTMENTS

| | | |
|---|---|---|
| AWS, Amazon | **Applied Scientist** | Jan 2023 - now |
| Microsoft Research, Redmond | **Research Intern** | May 2020 - July 2021 |
| Microsoft Research, India | **Research Fellow** | July 2016 - July 2018 |

## EDUCATION

**Carnegie Mellon University** *August 2018 - Jan 2023*
Ph.D. CyLab
*Thesis:* Practical network layer machine learning for IoT security
*Advisors:* Prof. Vyas Sekar & Prof. Anthony Rowe

**International Institute of Information Technology** *July 2011 - July 2016*
B.Tech and M.S. By Research in Computer Science

## HONORS AND AWARDS

- Our work on "Lumos" covered by Hackernews, CyLab, and others
- Carnegie Institute of Technology, Deans Fellow, Carnegie Mellon University
- Won outstanding technical achievement award for FarmBeats as part of Microsoft A.I. School
- Our work on "Annotation of Cricket Videos" covered by the Washington Post, NDTV, and others
- Made into Deans list of Academic and Research excellence at IIIT Hyderabad

## PATENTS

1. USO10942767B2: "Deep neural network workload scheduling" *March 2021*
2. US20200150640A: "Sensor Fall Curve Identification" *May 2022*

## PUBLICATIONS

**CoNEXT:** "Lumen: A Framework for Developing and Evaluating ML-Based IoT Network Anomaly Detection"
**Rahul Anand Sharma**, Ishan Sabane, Maria Apostolaki, Anthony Rowe & Vyas Sekar
CoNEXT, 2022

**USENIX Security:** "Lumos: Identifying and Localizing Diverse Hidden IoT Devices in an Unfamiliar Environment"
**Rahul Anand Sharma**, Elahe Soltanaghaei, Anthony Rowe & Vyas Sekar
USENIX Security, 2022

**USENIX Security:** "Accurately Measuring Global Risk of Amplification Attacks using AmpMap"
Soo-Jin Moon, Yucheng Yin, **Rahul Anand Sharma**, Yifei Yuan, Jonathan M. Spring & Vyas Sekar
USENIX Security, 2021

**BuildSys:** "Robust and Practical WiFi Human Sensing Using On-device Learning with a Domain Adaptive Model "
Elahe Soltanaghaei, **Rahul Anand Sharma**, Zehao Wang, Adarsh Chittilappilly, Anh Luong, Eric Giler, Katie Hall, Steve Elias & Anthony Rowe
ACM Conference on Systems for Energy-Efficient Built Environments (BuildSys), 2020

**SIGCOMM:** "Contention-Aware Performance Prediction for Virtualized Network Functions"
Antonis Manousis, **Rahul Anand Sharma**, Vyas Sekar & Justine Sherry
ACM Special Interest Group on Data Communication (SIGCOMM), 2020

**IPSN:** "All that GLITTERs: Low-Power Spoof-Resilient Light Anchors for Augmented Reality"
**Rahul Anand Sharma**, Adwait Dongare, John Miller, Nicholas Wilkerson, Daniel Cohen, Vyas Sekar, Prabal Dutta & Anthony Rowe
ACM/IEEE Conference on Information Processing in Sensor Networks (IPSN), 2020

**NSDI Poster:** "DeepEdge: A Network Edge for Deep Learning Workloads"'
**Rahul Anand Sharma** & Ranveer Chandra
USENIX Symposium on Networked Systems Design and Implementation Poster (NSDI Poster), 2019

**COMPASS:** "Low-Cost Aerial Imaging for Small Holder Farmers"
Vasuki Narasimha Swamy, Deepak Vasisht, Zerina Kapetanovic, **Rahul Anand Sharma**, Ranveer Chandra, Manohar Swaminathan, Anirudh Badam, Gireeja Ranade & Sudipta Sinha
ACM SIGCAS Conference on Computing & Sustainable Societies (COMPASS), 2019

**Arxiv:** "Learnability of Learned Neural Networks"
**Rahul Anand Sharma**, Praneeth Netrapalli, Navin Goyal, & Monojit Choudhary

**Sensys:** "Fall-curve: A novel primitive for IoT Fault Detection and Isolation"
Tusher Chakraborty, Akshay Nambi, Ranveer Chandra, **Rahul Anand Sharma**, Manohar Swaminathan, Jonathan Appavoo & Zerina Kapetanovic
ACM Conference on Embedded Networked Sensor Systems (SenSys), 2018

**WACV:** "Automated top view registration of broadcast football videos"
**Rahul Anand Sharma**, Bharath Bhat, Vineet Gandhi & C.V. Jawahar
IEEE Winter Conference on Applications of Computer Vision (WACV), 2018

**JOSIVP:** "Automatic analysis of broadcast football videos"
**Rahul Anand Sharma**, Visesh Chari, Vineet Gandhi & C.V. Jawahar
Springer Journal on Signal, Image and Video Processing (JOSIVP), 2016

**ACPR:** "Fine-Grain Annotation of Cricket Videos"
**Rahul Anand Sharma**, Pramod Kompalli & C.V. Jawahar
Asian Conference on Pattern Recognition (ACPR), 2015

**ICVGIP:** "Event Recognition in Broadcast Soccer Videos"
Himangi Saraogi, **Rahul Anand Sharma** & Vijay Kumar
Indian Conference on Computer Vision, Graphics and Image Processing (ICVGIP), 2016

## PROJECTS

- **AWS, Amazon**
  *Applied Scientist* January 2023 - Now

  – Q Spark Auto-debugging: Led the development of a generative AI-powered auto-debugging system for Apache Spark applications running on AWS services (MaxDome, AWS Glue, and Amazon EMR). Leveraging machine learning and generative AI techniques, the system provided automated root cause analysis (RCA) for failed Spark applications, significantly reducing mean time to resolution from days to minutes. Additionally, designed and integrated the auto-debugging feature with AWS Glue, Amazon EMR, and MaxDome consoles, ensuring a seamless user experience across these platforms.

  – Active network testing: Developed an innovative service for validating AWS Networking's Access Control Lists (ACLs) using live network test traffic. This service implemented an active testing methodology to verify traffic correctness, avoiding reliance on network simulation, heuristics, and models. By ensuring ACLs allow and drop traffic as expected in the live network, this solution significantly improved network security and reliability for AWS customers.

- **Microsoft Research** Bangalore, India
  *Research Fellow* Aug 2016 - July 2018

  – FarmBeats: Worked with Dr. Ranveer Chandra & Dr. Manohar Swaminathan on FarmBeats: IoT for agriculture. Our mission is to improve farming through data-driven practices. We are inventing new ways to gather data, store them in the cloud, and provide analytics to farmers.

  – Learnability of Neural Networks: Together with Dr. Praneeth Netrapalli & Dr. Navin Goyal, we are exploring the simplicity of learned neural networks under various settings: learned on real vs. random data, varying size/architecture, and using large minibatch size vs. small minibatch size. The notion of simplicity used here is that of learnability, i.e., how accurately can the prediction function of a neural network be learned from labeled samples?

  – DeepEdge: In collaboration with Dr. Ranveer Chandra, we propose a new Internet Edge architecture, called DeepEdge, for handling Deep Learning workloads on a network edge. It leverages insights about the structure of Deep Neural Networks (DNN) to schedule tasks in a way that can achieve good performance for small workloads and gracefully degrades when the workload increases

- **HackerRank** Bangalore, India
  *Software Engineering Intern* May 2015 - Aug 2015

  – Code Checker: HackerRank organizes various online programming contests, and Codechecker is the module that makes sure that each submitted solution passes the required test cases while following the specified resource constraints. As part of my work, I improved the Codechecker by adding support for Image Processing related problems. My other work as an intern led to a ten-fold reduction in deployment time by improvising the asset pre-compilation pipeline.

  – Static Code Analysis: As part of the internal hackathon organized by the company, we developed a static code analysis tool, which is currently under development to be offered as a service to end-users

- **IIIT Hyderabad** Hyderabad, India
  *Research assistant* May 2013 - Aug 2016

  – Automated top view registration of broadcast football videos: Developed a system for automatic registration of a broadcast soccer frame to its corresponding top view. For the task of image registration, we propose an alternate approach exploiting the edge information and

demonstrating its success in a specific scenario of registering football broadcast video frames on the static top view model of the playing surface

- Automatic analysis of broadcast football videos using contextual priors: Demonstrated that contextual information can be harnessed for automatic analysis of sports videos. Here, the proposed algorithm allows us to extract salient events, such as Goals, fouls, corners, substitutions, etc., automatically from a given broadcast soccer video.
- Fine Grain Annotation of Cricket Videos: Presented a solution that enables rich semantic annotation of Cricket videos at a fine temporal scale. Our approach circumvents technical challenges in visual recognition by utilizing information from online text commentaries. We obtain a high annotation accuracy, as evaluated over a large video collection.
- Event Recognition in Broadcast Soccer Videos: Proposed an approach for soccer event recognition using deep convolutional features combined with domain-specific cues. This approach uses the deep convolution feature (TDD) in combination with our proposed algorithm of Automated top view registration to improve upon the task of Event recognition.

- **IIIT Hyderabad** Hyderabad, India
  *Teaching Assistant*
  - Artificial Intelligence, Principle of Programming Languages, Digital Image Processing and Digital Signal Analysis and Applications

## INVITED TALKS AND CONFERENCES

**Lumen: A Framework for Developing and Evaluating ML-Based IoT Network Anomaly Detection**

- at CONIX Annual Review, CyLab, CMU, Pittsburgh, PA *Oct 2022*
- at CyLab Annual Review, CyLab, CMU, Pittsburgh, PA *Oct 2022*

**Lumos: Identifying and Localizing Diverse Hidden IoT Devices in an Unfamiliar Environment**

- at USENIX Security, Boston, MA *Aug 2022*
- at CONIX Student-Liaison Seminar, CMU, Pittsburgh, PA *Aug 2021*
- at IoT Reading Group, CMU, Pittsburgh, PA *July 2021*
- at CONIX Annual Review, CyLab, CMU, Pittsburgh, PA *Oct 2020*
- at CyLab Annual Review, CyLab, CMU, Pittsburgh, PA *Oct 2020*

**All that GLITTERs: Low-Power Spoof-Resilient Light Anchors for Augmented Reality**

- at CONIX Annual Review, CyLab, CMU, Pittsburgh, PA *Oct 2020*
- at IPSN, CPS-IoT Week 2020 (virtual) *April 2020*

**FarmBeats: AI, Edge & IoT for Agriculture**

- at Digital Platforms for improving Rural Livelihoods, Patna, India *April 2018*
- at Microsoft Techfest, Redmond, WA *Mar 2018*

**Automated top view registration of broadcast football video**

- at WACV, Lake Tahoe, NV *Mar 2018*

**Fine-Grain Annotation of Cricket Videos**

– at R & D Showcase, IIIT Hyderabad, India *Mar 2017*