

Amey Agrawal

Systems Engineer, Microsoft Research

ameya.info [@ agrawalamey12@gmail.com](mailto:agrawalamey12@gmail.com) github.com/agrawalamey [Google Scholar](https://scholar.google.com/citations?user=ameya)

Education

Jul 2018 | **Birla Institute of Technology and Science Pilani** | **Pilani, India**
Aug 2014 | B.E. (Hons.), Computer Science

Experience

Present | **Microsoft Research** | **Bangalore, India**
Jan 2021 | *Systems Software Engineer-II* | Mentors: [Dr. Muthian Sivathanu](#), [Dr. Bhargav Gulavani](#)
Building low-level systems to support large-scale workloads in Azure's globally distributed AI infrastructure service.

Nov 2020 | **Qubole Inc.** | **Bangalore, India**
Jul 2018 | *Member of Technical Staff-II* | Mentors: [Rohit Karlupia](#), [Joydeep Sen Sarma](#)
Worked on various applied machine learning and software engineering problems to enhance Qubole's data science platform. Published research in several top-tier venues.

Dec 2017 | **Norah.ai** | **Bangalore, India**
Jul 2017 | *Software Engineering Intern* | Mentor: [Bharath Bhushan](#)
Built core data-plane components for Qubole's Deep Learning clusters based on TensorFlow and Apache Spark.

Jun 2017 | **Norah.ai** | **Bangalore, India**
May 2017 | *Deep Learning Research Intern*
Developed various deep learning models for animation synthesis for Norah.ai's AI-powered game engine.

Publications

Singularity: Planet-Scale, Preemptible, Elastic Scheduling of AI Workloads [pdf]
Singularity Team, Microsoft
Preprint: [arXiv:2202.07848](https://arxiv.org/abs/2202.07848) (2022) [CoRR]

Learning Digital Circuits: A Journey Through Weight Invariant Self-Pruning Neural Networks [pdf][code]
[Amey Agrawal](#), and [Rohit Karlupia](#)
Sparsity in Neural Networks Workshop 2021; New in ML Workshop, NeurIPS, 2019, Vancouver [SNN'21]

Delog: A Privacy Preserving Log Filtering Framework for Online Compute Platforms [pdf] [dataset]
[Amey Agrawal](#), [Abhishek Dixit](#), [Namrata Shettar](#), [Darshil Kapadia](#),
[Rohit Karlupia](#), [Vikram Agrawal](#), and [Rajat Gupta](#)
Proceedings of IEEE International Conference on Big Data, 2019, Los Angeles [BigData'19]

Logan: A Distributed Online Log Parser [pdf]
[Amey Agrawal](#), [Rajat Gupta](#), and [Rohit Karlupia](#)
Proceedings of IEEE International Conference on Data Engineering, 2019, Macau [ICDE'19]

Select Research Projects

Efficient Device Sharing in Distributed Deep Learning Training Jobs | Jan'21 - Present
Mentors: [Dr. Muthian Sivathanu](#), [Dr. Bhargav Gulavani](#)

- > Currently, working on a proxy layer for GPU drivers that enables transparent checkpointing and time slicing for distributed deep learning training workloads with minimal overhead.
- > Efficient device sharing between data-parallel peers enabled by this system would power efficient job scheduling and resource management on Microsoft's next generation deep learning platform.

Log Parsing & Anomaly Detection System [blog] | Aug'19 - Mar'20
Mentors: [Rohit Karlupia](#)

- > Developed a distributed log parsing algorithm that provides **39x** speed-up over the previous state-of-the-art systems.
- > Created new metric to evaluate performance of log parsing systems. Developed a privacy-preserving federated learning framework to perform anomaly detection on sensitive log data.
- > The system is used to perform real-time anomaly detection on **250 million log lines every day** at Qubole. This work has been published in parts at **ICDE'19** and **BigData'19**.

Learning Sparse Binary Networks

Sep'19 - Present

Mentors: [Rohit Karlupia](#)

- > Designed a framework to learn self-pruning binary neural networks. The learned networks completely encode the knowledge in topology alone and as a result are weight agnostic.
- > Identified special characteristics of batch normalization that allows such networks to virtually act like a digital circuit composed of NOR gates. Based on this finding, proposed a new normalization method that allows for binarization of activations along with network weights and inputs. Presented this work at **New in ML Workshop at NeurIPS'19**.
- > Currently, developing new initialization methods and regularization strategies to extend the framework to be compatible with large convolutional neural networks.

Learning Efficient Job Placement Policy for ETL jobs on Big Data Platforms

Aug'20 - Nov'20

Mentors: [Joydeep Sen Sarma](#), [Rohit Karlupia](#)

- > Designed a hybrid system that utilized a custom machine text parsing algorithm in conjunction with learned features to fingerprint semantically unique workloads.
- > Developed a learning based scheduling algorithm that utilized job fingerprints to provide optimized node-level packing. Demonstrated that the system could lead to **15% saving** in Qubole's control plane operation costs based on past job submission trends.

Deep Reinforcement Learning for Autonomous Warehouse Robots [\[code\]](#)

Jan'17 - Apr'17

Mentor: [Prof. Surekha Bhanot](#)

- > Developed an environment for training reinforcement learning agents using a popular robotics simulation package.
- > Designed Q-learning based agents to perform various navigation tasks for autonomous operation in warehouses. Introduced various regularization techniques to help model generalize better when transferred from simulation to the real world.

Select Engineering Projects

Managed MLflow Service in Qubole

Apr'20 - Nov'20

Mentors: [Rajat Gupta](#), [Vipul Modi](#)

- > Lead the development efforts to create a managed MLflow service in Qubole's data science platform, was awarded **Spotlight Award** for delivering under tight time constraints.
- > Designed and implemented several modules in MLflow, Apache Spark, and Jupyterlab to facilitate seamless integration across the MLOps stack. Worked with the product team to conceptualize a MLflow extension to manage versioned datasets.

RStudio Cloud Workbench [\[blog\]](#) [\[demo\]](#)

Aug'19 - Apr'20

Mentors: [Rajat Gupta](#), [Vipul Modi](#)

- > Coordinated engineering aspects of partnership between RStudio and Qubole to develop a fully managed RStudio workbench on cloud that integrated with Apache Spark to enable machine learning use cases.
- > Designed and built storage layer for RStudio workbench on top of Amazon S3 using a FUSE-based file-system, S3FS. Contributed several performance and feature enhancements to S3FS, including, improved metadata cache and support for persistent ACLs.

Secure Application Proxy for Data Plane Applications

Sep'20

Mentors: [Rajat Gupta](#)

- > Redesigned the application proxy layer at Qubole, which led to a **5x** speed-up in page load times for various web user interfaces hosted in Qubole's data plane.
- > Built an authentication caching mechanism which resulted in **20x** reduction in latency of data plane API calls.

Callisto: Bringing Jupyter notebooks to classroom [\[blog\]](#) [\[demo\]](#) [\[code\]](#)

Jan'18

Mentors: [Prof. Surekha Bhanot](#)

- > Developed cross-platform desktop application to host and grade assignments designed in Jupyter notebook.
- > The platform helped lower the barrier to entry in scientific Python ecosystem for newcomers by providing one-click setup of development environment.
- > Hosted live programming contents on the platform with over 250 concurrent users. This work was later presented at **PyCon India, 2020**.

Teaching and Leadership Roles

- Covid Central, Yavatmal** *Technical Lead* [demo] Apr'21 - Jul'21
> Led a group of four developers to create a platform to manage Covid-19 patients in hospitals. Provided the system to regional hospitals for no-charge.
- Introduction to Neural Networks & Fuzzy Logic** *Lead Teaching Assistant* [assignments] Aug'17 - May'18
> Introduced Python programming assignments along with a new custom-built evaluation platform. Other responsibilities included coordinating the team of seven teaching assistants to conduct labs, designing assignments and helping students with the term project.
- Introduction to Machine Learning** *Teaching Assistant* Jan'18 - May'18
> Conducted introductory sessions on scientific Python ecosystem, organized tests and programming assignments for over 100 students in the class.
- Department of Visual Media, BITS Pilani** *Design Team Lead* [behance] May'15 - Feb'16
> Headed a team of eight graphics designers to create websites and video games for cultural and technical festivals at BITS Pilani.
- Concepticon Initiative** *Founding Member/Lead Coordinator* Oct'14 - Feb'15
> Led a team of fifty volunteers to organize career awareness events among high school students. These events witnessed participation of over seventeen hundred students across three cities.

References

- | | |
|-------------------------|---|
| > Dr. Muthian Sivathanu | <i>Partner Research Manager, Microsoft Research</i> [🌐] |
| > Dr. Bhargav Gulavani | <i>Principle Research Engineer, Microsoft Research</i> [🌐] |
| > Prof. Surekha Bhanot | <i>Professor, BITS Pilani</i> [🌐] |
| > Joydeep Sen Sarma | <i>Founder & CEO Clearfeed.ai/Co-author Apache Hive</i> [🌐] |
| > Rohit Karlupia | <i>Founder & CEO BigFlip.in/Author Sparklens</i> [🌐] |