

Sameer Chaturvedi

sam.chaturvedi24@gmail.com | github.com/samc24 | linkedin.com/in/sameerc

Boston University, Boston, MA

GPA: 3.75 / 4, Presidential Scholar

Master's: AI + Deep Learning, **Bachelor's:** Computer Science + Astrophysics

Work Experience

- Machine Learning Engineer II at Amazon** (*Prime Video Browse and Discovery*) **2025 - current**
- Working on surfacing the best, most relevant content for viewers while helping creators find their audience.
 - Implemented Mixture of Retrieval ML models to proportionally promote content, based on genre and propensity.
 - Improved Share-of-Voice of genre carousels via custom logistic regression model with stratified sampling and softmax tuning, resolving flat predictions caused by 90%+ channel over-representation in training data.
 - Designed e2e pipeline for 'FreshPicks' - a carousel for new 3P content, driving +134K annualized subscriptions.
 - Architected cross-service content filtering pipeline for new subscription bundle carousels - with realtime entitlement parsing, token decoding, and multi-filter orchestration to deduplicate content for channel subscribers.
- Lead Software Engineer at MIT** (*Plasma Science and Fusion Center*) **2025 - current**
- Architecting DisruptionPy, a critical framework for analyzing plasma fusion data across five Tokamak reactors, to predict and avoid disruptions using anomaly models via an analysis pipeline to build ML-ready datasets.
 - Revamping the backend to decouple data access from specific implementations, enabling pluggable backends (MDSplus, Xarray/Zarr) through abstract interfaces for improved testability and multi-machine scalability.
- Software Engineer II at Microsoft** (*M365 Security Detections Team*) **2021 - 2025**
- Petabyte scale big data platform that leverages ML to detect cyber-attacks in near real-time to protect all of M365.
 - Lead engineer for first ever Microsoft mail anomaly detection to catch malicious mailbox access; response to global security incident where Russian state hackers accessed emails of executives and federal agencies.
 - Lead engineer for GPT based security detection assistant to analyze machine/identity activity for attacks.
 - Designed and developed the production workspace that enables data science and LLM backed security detections in an agile PySpark environment, integrated with AI libraries, CI/CD, and real-time job monitoring.
 - Expedited security detection programming by over 75% on distributed systems for parallel compute and scale.
 - Developed an infinitely scalable event processing system using Spark and Azure Functions with real-time alerting.
 - Supported platform BCDR by designing and developing an auto-failover geo-replicated architecture.
 - Analyzed host logs, atomic detections, random-forest trained entity detections to investigate suspicious activity.
- Lead Machine Learning Engineer at Cerebro Sports** (*Live Gameplay Team*) **2025**
- Developed app to interact with live basketball tournament stats via LLM chat, realtime graphs, and dashboards.
 - Designed pipeline that converts semantic text to SQL, queries stats DB, processes data, handles edge cases.
- Lead Machine Learning Engineer at BU Spark!** (*Computer Vision Team*) **2018-21**
- Worked on "Muse": Computer Vision and ML app for real-time basketball broadcasting, strategy, and training.
 - Partnered with Amateur league in Spain to livestream games and assist coach with film analysis.
 - Custom implementations of Kalman filter, motion energy, Circle Hough Transform, contour detection, frame differencing, and single-shot object detection models to detect + track the ball.
 - Improves shooting skills via pose estimation for shot-mechanics analysis and comparison with NBA players.
 - Implemented Savgol filtering, template matching, tuned model hyperparameters, and used regularization methods
- Machine Learning Engineer at AI.Reverie** (*Computer Vision Team*) **2020**
- Lead engineer on deep learning + iOS app for accurate real-time object detection and instance segmentation.
 - Improved model accuracy through dropout, batch normalization, and training on custom synthetic data.
 - Converted models to work on mobile devices, and used device data logging to plot phone path and trajectory.
 - Productionized the library and ETL pipeline by using custom PyTorch DataLoaders and custom dataset mappers.
- Software Engineer at Red Hat OpenShift** (*Container Runtimes team*) **2019**
- Upstream contributions to CRI-O (Kubernetes container runtime interface) and Libpod (container & pod manager).
 - Unblocked engineers from issues writing to image volumes while inside the container, fixed container hooks monitoring and execution, and enabled users to pass in flags as environment variables from the CLI