

Karan Chaudhari

Jersey City, New Jersey | (201) 920-8302 | karan.chaudhari@nyu.edu | krnch.me

WORK EXPERIENCE

- **Software Engineer 2, Microsoft, NY** (Oct 2022 - Present)
 - Optimizing Azure File storage at global scale. Built and deployed features providing latency, throughput and iops improvement for smb, nfs and rest protocols. Improved smb write performance by 30% and rest write by 15% which includes over 200 million of revenue for file storage
 - Led the revamp of Azure Files AI infrastructure to support accelerating OpenAI and enterprise AI workloads. Engineered application-layer enhancements across Azure Storage (Files, Table, Queue, Blobs) and optimized network protocols. Delivered a 35%–40% reduction in server latency and a significant performance boost for all AI inference workloads globally.
 - Partnered with engineering teams globally to identify and resolve flow bottlenecks and enhance developer productivity. Contributed on tools to streamline developer workflows and performance metrics reporting
 - Successfully deployed first party application pipelines globally enabling msgraph for entirety of azure files. Involved in multiple security initiatives at microsoft and part of the additional team which brought entra id authentication for azure files
- **Software Engineer, Morgan Stanley, NY** (Aug 2018 - Oct 2022)
 - Maintained and improved real-time/batch risk calculation platforms and margin engines for prime brokerage and wealth management
 - Built and operated Jenkins CI/CD pipelines for microservices handling global exchange data
 - Enhanced the NetClient C++ diagnostics library, building Python wrappers for network-level monitoring and developer automation
- **SDE Intern, Amazon Web Services, WA** (May 2017 - Aug 2017)
 - Worked on the AWS Codecommit Team. Developed three new API's for merging commits-Squash merge, Two branch merge, and Rebase merge. Worked on Coral, Brazil, AWS S3 and AWS DynamoDb
- **Research Assistant, NYU, NY** (Mar 2017 - May 2017)
 - Extracted all the data on kiva.org using Multiple EC2 instances running in parallel. Built models in python and apache spark to predict loan behavior and analyze all loan and lending data

PROJECTS

- **XFiles AI Assistant** Built an ai assistant using incident related data , automated building a vector store based of existing azure files tsgs, internal and public documentation in addition to last two years of incident data .
This project has led to a successful and rapid resolution of customer incidents at microsoft
- **Etrade Calc engine** Worked as a lead on the scalability and optimization squad for integrating all Etrade accounts due to recent merger. Added new calculation mechanisms for margin, improved platform resiliency, and account calc time by 8x for this project
- **Regression framework at morgan stanley** Built the regression framework , a separate microservice from scratch at morgan stanley , this service can be triggered using apis or as a batch run 3 times a day for ny, london and hong kong stock exchanges to analyze and calculate differences between realtime and end of day margins for all daily trades at morgan stanley

- **GraffitiMatch** Extracted all the 30000 graffiti images on data.sf.gov. Used this dataset to train a model using neural networks to predict if the image is graffiti or not. This project was built and presented at Techcrunch Disrupt hackathon and our submission won an award
- **Recommendation System** Used the yelp Dataset to predict business ratings using various methods as clustering, regression and Matrix Factorization. Published a paper titled "Predicting business success using Predictive data analytics" based on this project
- Contributed to **Open Source** projects. You can view my work on **my Github page** - github.com/krnch

EDUCATION, RELEVANT COURSES AND TECHNICAL SKILLS

- **M.S. Computer Science, New York University** (May 2018)
- **B.E. Computer Science, PICT, Pune** (May 2016)
- Languages: C, C++, Java, Python (8+ years; backend focus). Additional experience in Angular, Ruby, Python (Flask) for frontend/web.
- Relevant graduate courses: Cloud Computing, Big Data Analytics, Computer Vision, Artificial Intelligence, Machine Learning, Information Security and Privacy
- DevOps: Use GitHub Actions for CI/CD and personal website deployment

LEADERSHIP ACTIVITIES AND VOLUNTEERING

- Volunteered for the fresh air fund to teach high school physics and computer science to high school students.
- Actively involved in recruiting at morgan stanley and microsoft . Conducting interviews for new hires at morgan stanley and microsoft
- Involved in mentoring engineers at employee resource groups at morgan stanley and microsoft