Karan Grover

https://krngrvr09.github.io/

Email : kgrover2@wisc.edu
Mobile : +18434530358

Experience

- **Fetch Rewards** — Madison, WI
  Software Engineer — *Jul 2023 - Present*
  - **DynamoDB**: Saved more than $5,000 per month in DynamoDB costs by adding TTL fields to the production table with 16 Billion items. Wrote a massively parallel script using Go Routines to perform the updates at the rate of 40,000 items per second.
  - **Rate-Limiter Service**: Building a multi-region rate-limiter service on AWS using Go, Redis and Elasticache.
- **University of Wisconsin-Madison** — Madison, WI
  Software Engineer — *May 2023 - Jul 2023*
  - **ML Infrastructure**: Productionized a document-annotation ML Model on an HPC cluster using Docker. Details
  - **Dev Instance**: Automated the complete dev-instance setup for all engineers on the team. The setup creates multiple dockerized services and databases with one command and ingests relevant data automatically.
- **Meta (Facebook)** — New York
  Software Engineering Intern — *May 2022 - August 2022*
  - **AI Infrastructure**: Developed CLI tools to help ML Engineers quickly onboard models on the Inference Platform. The command-line tools reduced the model onboarding time from a few days to a couple of hours.
  - **A/B Testing**: Developed the A/B Testing Framework for Facebook's Inference Service using Python and C++. Designed a BI Dashboard to show live results of A/B Testing experiments to ML teams. Details.
- **Microsoft** — Bangalore, India
  Software Engineer — *Jan 2018 - March 2021*
  - **Privacy-Preserving ML**: Developed a cloud-native secure ML service using Intel SGX and Torch. Modified the C++ Torch library to run in a sandboxed cloud environment, profiled ML models inside Intel SGX, and introduced techniques to patch existing security vulnerabilities. Merged into Azure Confidential Computing. Details.
  - **Infrastructure Automation**: Created a VM Orchestration and Monitoring service for the Blockchain team. Scaled it up to handle 2500 VMs on Azure. It was used to identify failures early and improve stability. Details.
  - **Backend Development**: Worked on building a BI Platform for India's largest Non-Profit Organization using React and Flask. Built the following features: milestone tracking, Gantt charts, and progress reports. They were used to provide real-time insights on project progress and highlight areas requiring attention. Details.

PUBLICATIONS

- Sambhav Satija, Apurv Mehra, Sudheesh Singanamalla, **Karan Grover**, Muthian Sivathanu, Nishanth Chandran, Divya Gupta, Satya Lokam. Blockene: A High-throughput Blockchain Over Mobile Devices. In OSDI 2020: USENIX Symposium on Operating Systems Design and Implementation.
- **Karan Grover**, Shruti Tople, Shweta Shinde, Ranjita Bhagwan, Ramachandran Ramjee. Privado: Practical and Secure DNN Inference. Arxiv Preprint 2018
- Sougata Sen, Archan Misra, Vigneshwaran Subbaraju, **Karan Grover**, Meera Radhakrishnan, Rajesh Krishna Balan, Youngki Lee. I4S: capturing shopper's in-store interactions. In ISWC 2018: Proceedings of ACM International Symposium on Wearable Computer.
- Sougata Sen, **Karan Grover**, Vigneshwaran Subbaraju, Archan Misra. Inferring smartphone keypress via smartwatch inertial sensing. In WristSense 2017: IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops).
- **Karan Grover**, Vinayak Naik. Monitoring of Android devices using SNMP. In COMSNETS 2016: 8th International Conference on Communication Systems and Networks.

PROJECTS

- **Knowledge Discovery from Published Literature**: Developing the backend infrastructure of a Knowledge Extraction Framework - an advanced ML-assisted search engine for scientific publications. Supporting over 15 million scientific documents. Details.
- **Fault-Tolerant Distributed Training**: Modified the Pytorch and Gloo library to make the collective communication process fault tolerant. Analyzed the system with various distributed training settings and compared different approaches. Details.
- **Cloud-Native LevelDB**: LevelDB is a key-value store based on LSM Trees. Modified LevelDB to make it cloud-native using Amazon S3 as the backend. Developed a naive implementation and then tuned it to be 3x faster using Distributed Systems concepts. Details.
- **AWS Wordle Backend API**: Developed the Backend for Wordle game using AWS CDK, API Gateway, AWS Lambda, and DynamoDB. Thoroughly documented code, data model, and design decisions. Open-sourced the project. Code.
- **Youtube Watch-Party**: Developed the Backend for a Youtube Watch-Party Application where multiple people can watch youtube videos together. Used sockets, NodeJS, React and TypeScript to build the application. Thoroughly documented code, data model, and design decisions. Open-sourced the project. Code.

Education

- **University of Wisconsin-Madison** Madison, Wisconsin
  MS in Computer Science *Aug 2021 – May 2023*
  - **Courses**: Advanced Systems for ML, Distributed Systems, Introduction to AI, Big Data Systems, Topics in Databases, High Performance Computing, Intro to Information Security, Programming Languages and Compilers
- **Indraprastha Institute of Information Technology** Delhi, India
  Bachelor of Technology in Computer Science *Aug 2013 – Dec 2017*
  - **Courses**: Software Defined Networking, Security Engineering, Network Security, Wireless Networks, Analysis and Design of Algorithms, Operating Systems, Fundamentals of Databases

Certifications

- **NodeJS from Scratch - Educative**: Learned how to build server-side javascript applications using NodeJS, Express and Socket.IO. Certificate
- **Docker & Kubernetes, The Complete Guide - Udemy**: Learned production workflows of deploying Docker apps with Kubernetes. Built CI/CD pipelines from scratch with Github, Travis CI, and AWS. Certificate.
- **Infrastructure as Code with Terraform - Google**: Learned how to build, change, provision, and destroy infrastructure using Terraform via the Google Cloud console. Certificate.

Skills

- **Languages**: C++, Python, Java, Javascript, Bash
- **Technologies**: git, linux, REST, API, Relational Database, noSQL, DynamoDB, AWS CDK, HTML, Azure, Pytorch, Tensorflow, CUDA, Apache Spark, Hadoop, GRPC, Flask, Ruby on Rails, Docker, Kubernetes