

Adithya Samavedhi

Seattle, WA | Ph: +1 (858)-241-4810 | adithya.samavedhi@gmail.com
[linkedin.com/in/adithyasamavedhi/](https://www.linkedin.com/in/adithyasamavedhi/) | [adithya-samavedhi.github.io/](https://github.com/adithya-samavedhi)

WORK EXPERIENCE

The Trade Desk

Oct 2025 – Present

Seattle, USA

Software Engineer II

- Authored and open-sourced **ttd-data**, a PyPI SDK for The Trade Desk's Data APIs, adopted by clients driving 5,000+ requests/second of production traffic; standardized authentication, retries, and rate-limit handling across integrations.
- Built and open-sourced **ttd-databricks**, a PySpark SDK enabling clients to push segment data from Databricks to TTD's Data APIs, supporting incremental Delta Lake batch pipelines with checkpointing, per-row error tracking.
- Engineered a tree-based boolean attribution engine to process pixel fires across client websites, directly supporting \$27M in attributed platform spend.

Software Engineer I

Apr 2024 - Oct 2025

- Built a Snowflake Native App ingesting 10B+ user id-segment pairs, enabling fully self-serve onboarding for 24 data providers.
- Architected a Kafka-based data propagation system, cutting end-to-end processing time 90% and infrastructure costs \$300K/year.
- Designed and deployed a Redis-based distributed rate limiter to protect data ingestion services from exploitation, projected to save 150+ engineering hours per year.

Software Engineer Intern

Jun 2023 - Sep 2023

- Eliminated redundant cross-datacenter propagation by tracking user-location affinity and routing segment data only to relevant datacenters, cutting propagation traffic and network costs by 60%.

Goldman Sachs

Jul 2021 – Jul 2022

Bangalore, India

Software Engineer

- Engineered a Celery-based scheduling service exposing APIs to create and manage Kubernetes CronJobs, autonomously orchestrating database operations across 4,000+ middleware application endpoints.
- Developed GSMDWUtility, a Python library providing reusable REST middlewares (OIDC JWT validation, APM labeling, metadata augmentation) to standardize microservices across departments.
- Mentored an intern in building a monitoring dashboard for bare-metal OS patching, tracking canary-to-fleet rollout status, OS version distribution, and patch success across on-prem servers.

Couture AI

Jan 2021 - Jun 2021

Hyderabad, India

Machine Learning Intern

- Built a video recommendation system leveraging LightGBM and collaborative filtering, with a candidate generation pipeline spanning 20M videos and 6M users.

Goldman Sachs

May 2020 – Jun 2020

Bangalore, India

Software Engineer Intern

- Developed a PostgreSQL audit agent that parsed database logs across 800+ servers, classifying connection events, SQL queries, and errors into structured audit streams for centralized collection.

EDUCATION

University of California, San Diego

Sep 2022 - Mar 2024

Master of Science in Computer Science, GPA: 4.0/4.0

San Diego, USA

BITS Pilani

Aug 2017 - Jul 2021

Bachelor of Engineering in Computer Science, GPA: 8.71/10.0

Goa, India

PROJECTS

Surfstore - A Fault-Tolerant Dropbox Clone

- Built a fault-tolerant distributed file store scaling to 128 nodes with concurrent shared-file access and full CRUD support, using RAFT-replicated MetadataStore and BlockStore services to maintain availability with up to 50% node failures.

Networked Sort - Distributed Sorting Program

- Developed a distributed sort over raw socket programming spanning 16 server nodes, with an even data-partitioning algorithm that sorted 2GB+ datasets in under 15 seconds.

SKILLS

Languages: Python, C#, SQL, C++, Go, Java, C, Shell

Frameworks: .NET 8, FastAPI, Django, SQLAlchemy, MongoEngine, Celery

Infrastructure: Kubernetes, Docker, AWS, Kafka, Redis, RabbitMQ, Databricks

Data & Processing: Apache Spark (PySpark), Delta Lake, Kafka, Databricks, Snowflake

Databases: PostgreSQL, MongoDB, Snowflake, Aerospike

Observability & Tools: Prometheus, Grafana, Honeycomb, Git, Linux

APIs: REST, GraphQL, gRPC, OpenAPI