

Achal Shah

Seattle – WA – 98101

✉ m@achals.com • 📄 github.com/achals

Areas of interest: Scalability and Infrastructure, Distributed Systems, Machine Learning Systems

Experience

Tecton

Software Engineer

Seattle, WA

April 2020–Current

- Early engineer at Tecton, building an enterprise-ready feature store to make world-class machine learning accessible to every company. Our goal is to enable ML teams to build great features, serve them to production quickly and reliably, and do it at scale.
- I worked on various platform challenges, and am currently working on the Feast project.

Uber

Sr. Software Engineer

San Francisco + Seattle

Sept 2015–April 2020

- Michelangelo: Scaled Uber's Machine Learning platform. Worked on sharding the model serving layer, and added deep learning model serving capabilities.
- Reliability Platform: Built and operated Hailstorm, an high-throughput, distributed, load generation platform. Helped Uber make peak days "officially boring".

Amazon.com

Software Development Engineer

Seattle, WA

July 2013– Sept 2015

- Amazon Fulfillment Technologies - Outbound Flow: The team owned services to predict incoming work that warehouses would receive, and the throughput that warehouse associates would achieve, to help optimize planning and operational decisions.
- Amazon Fulfillment Technologies - Sortation Team: The team owned software to direct warehouse associates to sort items into individual shipments with minimal human effort. I helped scale the services to serve associates in new regions, and wrote new tools to interact with mechanical conveyance systems.

Raytheon BBN Technologies

Language Understanding Graduate Intern

Cambridge, MA

May 2012–August 2012

Worked on the FUSE project, which involved predicting scientific trends based on bibliometric data.

- Implemented feature extraction models for the task of citation prediction. Also implemented cluster quality metrics such as cosine similarity, coherence values, distribution of subjects.

Education

University of Pennsylvania

Master of Science and Engineering, Computer and Information Science

Philadelphia

2011–2013

Relevant Coursework: Internet and Web Systems, Database and Information Systems, Mathematical Statistics, Machine Learning

Institute of Technology, Nirma University

Bachelor of Technology, Computer Engineering

Ahmedabad

2007–2011

Academic Papers and Projects

Personality, Gender, and Age in the Language of Social Media

PLOS ONE, September 2013

HA Schwartz et al., DOI: 10.1371/journal.pone.0073791

Wrote the prototype to use LDA on a large collection of tweets to cluster words into topic-based word lists. The topics were then correlated with user personality traits.

Distributed Search Engine CIS 555: Internet and Web Systems

April 2013

Built a scalable search engine from scratch. Individual components involved a Mercator-style crawler that started with a list of seed URLs, a component to create an inverted index from the crawled pages, a hadoop implementation to compute the page rank of pages and a servlet-based webserver to serve the results of web queries.

Predicting Amazon Ratings CIS 520: Machine Learning

August 2011

Implemented and compared different machine learning algorithms along with some feature selection methods on a dataset of Amazon product reviews and their corresponding ratings to predict the ratings for unseen reviews.

Automatic Text Summarization CIS 530: Computation Linguistics

November 2011

Built a system to use Named Entity Recognition, coreference resolution and other lexical features to generate summaries from a collection of documents.

Independent Study: Dependency Eigenwords with Prof. Lyle Ungar

September 2012

Use dependency parse contexts to create eigenwords.

Skills

Fluent in: Java/Kotlin, Python, Go

Familiar with: C/C++, Rust, Javascript, Scala

Execution Frameworks & Technologies: Apache Spark, Ray, Dask, Kubernetes, gRPC, Spring

Databases and Data stores: MySQL, DynamoDB, Redis, Cassandra