Alex Rudy

hello@alexrudy.net
alexrudy.net
github.com/alexrudy

Overview

- Strategy and execution focused, with experience from data warehouses through state of the art machine learning pipelines.
- Strong track record of deploying high-impact compliant machine learning models for a variety of business applications.
- Experience with a wide variety of modeling techniques: tree-based ensemble models, neural networks, hierarchical bayesian models.
- Proven system architect: high-read production Postgres tables, data warehouses and ETL pipelines ML pipelines.
- Developed and deployed extensively in Python, SQL, C and Rust. Experience with C++, Go, JavaScript, Ruby, Fortran & R.

Experience

Staff Machine Learning Engineer, Platform — August 2022 – Present

Discord – San Francisco, CA – Technical leadership and building the feature store to power Machine Learning at Discord

Principal Machine Learning Engineer — March 2021 – June 2022

CloudTrucks – San Francisco, CA – Data Science, Algorithm Development, Infrastructure and Product Engineering

- Architect for a query caching and crawling system, solving numerous pain points with our existing Django on-demand load search, API, and RPA system. I re-designed the core truck load data models used at CloudTrucks to power our entire app.
- Built and deployed a route optimization tool backed by an adaptive ruin and recreate algorithm which powers the Schedule Optimizer, a tentpole feature for CloudTrucks, driving signups for our Flex and Virtual Carrier products. Built new algorithms and compatible API in Rust, leveraging pyo3, tokio, and serde to provide backwards compatible interfaces to existing python code.
- Responsible for Docker, Python, Django and CircleCI infrastructure, including Celery task queues, dependency management, testing infrastructure, and continuous deployment of our container images. Partner with our primary infrastructure engineer to maintain our GCP infrastructure, networking, proxying, host management, metrics, and monitoring – using terraform, prometheus and grafana.
- Technical mentor and educator – building the CloudTrucks technical onboarding program, teaching intermediate and advanced python skills, and ensuring that CloudTrucks has a high quality developer experience from day one using docker-compose.

Head of Data Science / Principal Data Scientist — August 2019 – March 2021

Bitly – San Francisco, CA – Lead data science for the Bitly IQ startup within Bitly.

- Created NLP skip-gram models to provide semantic tagging for billions of crawled pages across consumer industries.
- Built a cloud-native machine learning pipeline for building models using TensorFlow on Bitly's click and web history, taking data from BigQuery and google cloud storage, pre-processing with data flow pipelines, training on Google's TPU infrastructure, and deployed using docker containers, all orchestrated using GCP's AirFlow equivalent Cloud Composer.
- Developed a BERT-descendant natural language model to identify suspicious and malicious URLs on the Bitly platform.

Senior Data Scientist — March 2018 – August 2019

Even – Oakland, CA

- Developed a novel ML algorithm to apply active learning and hierarchical clustering to personal financial transaction data. The algorithm allows the Even app to budget for predictable expenses, powering 250,000 personal financial plans.
- Led the design and architecture of a new data warehouse to expose business data to Even employees. Built and evangelized the use of transformed data with business definitions, democratizing data analysis across the company. Leveraged a custom fork of DBT and collaborated to build an ELT pipeline to move data from Amazon RDS to S3, write parquet files, and transform using Amazon Athena.

Data Scientist — July 2017 – March 2018

LendUp Card Services – San Francisco, CA – became Mission Lane in 2018

- Built an underwriting model for applicants with bankruptcies – 80% of credit card approvals and billions of dollars in growth.

National Science Foundation Research Fellow — September 2012 – July 2017

University of California, Santa Cruz, University of California Observatories, and Lawrence Livermore National Laboratory

Education

PhD, Astrophysics, University of California Santa Cruz, CA — July 2017
Advisor: Claire Max. Thesis Topics: Predictive Control for Adaptive Optics, Gas Dynamics of Nearby Galaxies

BA, *cum laude*, Physics, Pomona College, Claremont, CA — May 2011