

Alexander Lee

westford14@gmail.com | www.linkedin.com/in/alex-y-lee | (06) 234-38786 | Arnhem, NL

SUMMARY

Senior Software Engineer with over 10 years of experience across infrastructure, data orchestration, machine learning, CI/CD implementation, developer tooling and API development. Has worked in various job functions ranging from Data Scientist to DevOps engineer and is eager to apply these skills in organizations that benefit from collaboration, reliability, and creative problem solving. Excels in owning the lifecycle of a product from roadmapping to execution utilizing cloud deployment and modern tools and languages.

SKILLS

- **Databases and Cloud Services:** AWS, GCP, Azure, Elasticsearch, Redis, MongoDB, Postgres
- **Languages:** Python, Go, TypeScript, R, SQL, GraphQL, Stata, Bash
- **Tools and Technologies:** Kubernetes, Istio, Prometheus, Kubeflow, Terraform, Atlantis, Helm, CircleCI, Github Actions, Drone, Airflow, Prefect, Dagster, Confluent, Spark, Bazel

EXPERIENCE

APrime Technology

Senior Software Engineer II

May 2022 - Feb. 2025

New York, NY

Contributed to various projects for clients as a contractor for Prime Technology, detailed below.

Client: Usage-Based Billing Platform

- Improved Kubernetes management with Argo Rollouts, custom CircleCI orbs, and Helm
- Ran a successful data orchestration platform migration from AWS Step Functions to Dagster; this included initial scoping and POC creation, stakeholder buy-in, and roadmapping planning
- Modernized migration tooling by implementing a custom Go library that can only be invoked through CI/CD
- Demoed a Bazel building platform for all microservices to unify gRPC spec creation and deployment practices

Client: Employee Benefits Enablement Platform

- Worked with multiple cross-functional teams to perform a zero downtime VPC migration
- Deployed a cross-region, cross-account offsite data backups with realtime data replication via AWS that adhered to all SOX compliance and HIPAA regulations
- Revamped the company's entire Terraform stack allowing for more modular builds and implemented better CI/CD practices with Atlantis

Client: FinTech Index Fund

- Created an end to end data and prediction pipeline in Prefect for a FinTech company using a shrinkage based time varying Bayesian algorithm
- Translated all code from R to Python with parallelization to reduce overall runtime by nearly a third

Other Client Projects:

- Migrated an education platform's cloud infrastructure from AWS to GCP; spearheaded all scoping, runbook creation, downtime mitigation, and services matching
- Implemented multiple API changes to a GraphQL, TypeScript stack for a Real Estate empowerment platform
- Used Prefect to improve and harden a prescription drug reimbursement platform's data pipelines while maintaining SOX and HIPPA compliance

New York Times

Senior Software Engineer, Machine Learning Infrastructure

July 2021 - May 2022

New York, NY

- Revamped the CI/CD processes in Drone utilizing path-based builds to speed up deployments by four times
- Used Istio in the ML Platform clusters to enable canary and shadow deployments and allow for requests based horizontal pod autoscaling
- Increased unit and integration test coverage and rolled out context based logging for the Golang model service
- Led effort to move backend and frontend services to multiple regions via Terraform

EXPERIENCE

Bark

Lead Machine Learning Engineer April 2020 – June 2021

Jan. 2019 – June 2021

New York, NY

- Standardized model training with Kubeflow, GKE, and CircleCI for a team of data scientists
- Created a custom feature store and experiment tracker using BigQuery and GCS
- Shifted prediction serving from a batch prediction to a realtime framework using TensorFlow Serving
- Employed Graph Neural Networks to improve our recommendation engine performance
- Improved our survey sampling strategy to capture a true randomize sample of our users for better NPS understanding
- Utilized a hierarchical GLM to capture unattributed churn due to COVID-19
- Handled a database migration from MongoDB to Amazon Redshift

Machine Learning Engineer Nov. 2019 – April 2020

- Implemented a Linear Programming solver to optimize warehouse assignments for the add-to-box experience
- Created a robust test suite for the add-to-box codebase
- Continued leading technical recruitment for all data science related roles
- Trained and deployed a suite of BERT-based NLP models to predict toy attributes from user comments and feedback

Senior Data Scientist: Jan. 2019 – Nov. 2019

- Moved our recommendation engine from external consultants internally, resulting in improved performance
- Utilized survival models to predict retention of subscribers and impact business decisions
- Deployed a platform for data scientists to securely share their insights with the rest of the company

Trendalytics

July 2018 – Jan. 2019

New York, NY

Software Engineer, Machine Learning

- Utilized XGBoost to improve our search relevancy NLP results

Institute for Health Metrics and Evaluation

July 2016 – June 2018

Seattle, WA

Machine Learning Data Specialist: April 2018 – June 2018

- Led a team of 3 that deployed an image recognition Convolutional Neural Network into production
- Gave a lightning talk and clinical track poster presentation at the 2018 Machine Learning for Health Care 2018 conference at Stanford University.

Data Analyst: July 2016 – April 2018

- Produced figures and tables for multiple publications published in high impact scientific journals such as PLOS ONE, the Lancet, and the New England Journal of Medicine

National Human Genome Research Institute, NIH

June 2014 – June 2016

Washington, DC

Scientific Program Analyst

- Presented a poster on graph algorithms in bioinformatics at the 2015 NHGRI Symposium

AWARDS

- Innovator of the Year 2018 – **Institute for Health Metrics and Evaluation**
- First Prize, Summer Coding Challenge 2017 – **Institute for Health Metrics and Evaluation**
- Distinguished Achievement Award – **National Human Genome Research Institute, NIH**

EDUCATION

Wageningen University and Research

August, 2024 - May 2026

Wageningen, NL

MS Candidate, Data Science for Food and Health

- Thesis: A unified approach to adjust for correlated measurement error in multiple exposures using external validation data, with implementation in R

Haverford College

June, 2014

Haverford, PA

BS in Chemistry, Concentration in Biochemistry, Minor in History of Art

- Led the college's student-run first-year orientation program, managing a team of nearly 200