

Grant Baker

Machine Learning
and Software
Engineer

EDUCATION

2015-2019

B.S. / M.S. Applied
Mathematics, University
of Colorado Boulder
Minor: Computer Science
Emphases: Mathematical
Modeling, Machine
Learning, Optimization

SKILLS

Proficiencies

- Backend Software
- AWS Architecture
- Deep learning
- Data analysis
- Statistical modeling
- Engineering leadership

Languages

- Python
- Scala
- JavaScript

Frameworks

- TensorFlow
- Keras
- XGBoost
- Scikit-learn
- Spark
- Pandas
- Kubernetes

CONTACT

grant.john.baker@gmail.com
303-501-4136
linkedin.com/in/gbaker0

PROFILE

Experienced software engineer and machine learning specialist, adept at integrating machine learning models with software applications. Passionate about programming and mathematics, I thrive on solving complex problems and driving innovation in machine learning.

WORK EXPERIENCE

LEAD MACHINE LEARNING ENGINEER | BRYTE

February 2021 - Current

Leading a team of software engineers to build and maintain cloud applications for the Bryte system, including support for IoT devices, mobile clients, and web applications. Providing modeling and ML expertise in many applicable areas across the product.

- Researched, developed, and implemented the Sleep Concierge, a GPT-4-powered chat assistant which empathetically and actionably helps users interpret their sleep data
- Developed a custom highly-performant Python REST API framework for Lambda applications and rewrote our API with it
- Strengthened cloud security and reliability by implementing best practices including SSO login, audit logs, monitoring and alerting systems, least-privilege permissions, short-term credentials
- Created an optimization framework for providing users with recommendations for selecting the optimal settings for their bed
- Wrote continuously updating models for predicting bed inflation and deflation durations, individualized by hardware
- Leading a team of 4 software engineers through many projects, including the API, internal web tools, user-facing web apps, partner integrations, database migrations, user management, monitoring systems, and a complete overhaul of our software tooling to use CloudFormation infrastructure-as-code

MACHINE LEARNING ENGINEER | ORACLE DATA CLOUD

June 2018 - February 2021

Improved, extended, and maintained a terabyte-scale machine learning and data engineering pipeline powering the foundational Identity Graph product, used across the business.

- Researched and implemented a scalable community detection algorithm in Scala Spark, powering several new layers and opening new business opportunities
- Researched and implemented scalable machine learning training and inference steps with Spark XGBoost for improved edge prediction for each type of link in the Identity Graph
- Built a custom Kubernetes orchestration engine, which coordinated data pipeline steps with Luigi in Python and executed graph construction business logic in Scala Spark