Taylor Bosier

(402) 984-9277 · bosier83@gmail.com · tbosier.github.io

Summary

Lead Data Scientist, 7+ years building production ML, Bayesian inference and optimization models that operational teams run daily — across warehousing, freight and medical devices. Optimization, hierarchical Bayesian inference, causal measurement, and retrieval over large document estates.

Professional Experience

Lead Data Scientist *Aug 2024 – Present*
Stryker (Remote)

- Built a causal inference pipeline in Databricks SQL and Python evaluating rebate ROI with propensity score matching and inverse propensity weighting over customer-level engagement and eligible-product revenue.
- Designed a modular price-optimization platform recommending SKU- and family-level prices from isotonic demand models and dynamic programming over price tiers, published automatically through Databricks pipelines.
- Built an auditable rebate-forecasting platform (Python, SQL, Databricks, React/TypeScript) with leakage-safe backtesting and audit exports; reconciled a material data-mapping discrepancy and cut report load time by 53%.
- Engineered a grounded contract-AI assistant pairing semantic retrieval and vector search with validated read-only SQL, exact citations and PDF text highlighting, making answers traceable to the clauses behind them.
- Built contract-performance intelligence decomposing year-over-year revenue change into price, volume and mix, with contract-year baselines, sold-to analysis and account-level drill-down into where revenue was won or lost.
- Architected a reusable job-execution platform replacing one-off pipeline schedulers, with PostgreSQL run tracking, live cron configuration, authorization, retries, hard timeouts, failure alerts and an operations console.
- Mentored a senior data scientist and full-stack developers; partnered with finance, operations and pricing to put models into daily use.

Data Scientist II *Sept 2021 – Aug 2024*
Knight-Swift Transportation, Phoenix, AZ — *promoted from Data Scientist I, Sept 2023*

- Re-architected and deployed a scheduling optimization model (Gurobi) for driver planning, raising planned drivers at one terminal from 20% to 85% and improving on-time delivery and network design.
- Designed an early-warning detection model surfacing 40+ drivers a week for safety remediation.
- Conducted a study using a Bayesian Poisson-Gamma model to quantify the safety impact of lane-keeping sensors, providing critical evidence for a multi-million dollar fleet hardware decision.
- Rebuilt a driver retention model, taking F1 on the churn class from 0.20 to 0.40; the model fed the retention program that ran alongside a 12% fall in driver turnover.
- Built a frequency-severity model of fleet loss exposure that set the retention and limit structure for the annual insurance placement.

Business Analyst *June 2019 – Sept 2021*
Chain Link Services

- Sited warehouses by k-means clustering over demand geography and forecast inventory demand with Prophet, cutting holding costs and stockouts; restructured pick paths with an apriori model and aisle-traversal heuristics, and set offsite levels with a safety-stock rule.

Technical Skills

Modelling: Python, PyMC, scikit-learn, XGBoost, PyTorch, PySpark · **Optimization & simulation**: Gurobi, Pyomo, PuLP, HiGHS, SimPy, ARENA
Data & cloud: Databricks, Azure, SQL Server, PostgreSQL, DB2, Hadoop, Elasticsearch · **Shipping**: Prefect, Airflow, Git, GitHub, GitLab

Education

New Mexico State University, Las Cruces, NM — M.S. Mathematics (May 2018), B.S. Mathematics (May 2016)