Daniela Ortiz

Data Scientist · Forecasting, Experimentation & Causal Inference

Minneapolis, MN · (612) 555-0137 · daniela.ortiz@email.com · linkedin.com/in/danielaortiz · github.com/dortiz-stats

SUMMARY

Data scientist with 7 years in retail and healthcare, working in Python and SQL. Builds demand forecasts and designs A/B tests for a national retailer, and uses quasi-experimental methods when randomization is not an option. Ships models to production with the ML platform team rather than handing off notebooks.

EXPERIENCE

Data Scientist II · Target — Mar 2022 – Present

Minneapolis, MN

- Built a LightGBM demand forecast for perishable categories across 1,900 stores, lowering weekly WMAPE from 34% to 27%.
- Designed and read 40+ A/B tests on site search ranking; added CUPED variance reduction that cut typical test length from 21 to 14 days.
- Estimated the sales effect of a regional price change with difference-in-differences across 212 stores; finance used the estimate to set Q4 pricing.
- Shipped the markdown model with the ML platform team as a Python package, retrained weekly in Airflow, with drift alerts in Datadog.
- Presented quarterly model reviews to merchandising leaders, translating forecast intervals into safety-stock targets for 60 categories.

Associate Data Scientist · Optum — Jul 2019 – Feb 2022

Eden Prairie, MN

- Trained an XGBoost 30-day readmission model on 1.2M claims; care managers used its top-decile list to prioritize outreach calls.
- Cleaned and joined claims, pharmacy and eligibility tables in Snowflake into a member-month feature set reused by 3 teams.
- Wrote SHAP-based reason codes so nurse care managers could see why each member was flagged before calling.
- Audited model calibration by age band and sex before launch and documented the results for the compliance review.

EDUCATION

M.S. Statistics · University of Minnesota — May 2019

B.A. Economics, minor in Mathematics · Macalester College — May 2017

SKILLS

Languages: Python (pandas, scikit-learn, statsmodels), SQL, R

Methods: Time-series forecasting, Gradient boosting, A/B testing and CUPED, Difference-in-differences, Propensity score matching, Model explainability (SHAP)

Platforms: Snowflake, Databricks, Airflow, MLflow, Git, Tableau