

MARK BRUNER

Cincinnati, OH • contact@markbruner.io • linkedin.com/in/mark-bruner • markbruner.io

SENIOR DATA SCIENTIST | INNOVATION-DRIVEN ML, DECISION SCIENCE & MLOPS

SUMMARY

Innovation-minded data science leader (6+ years) who turns ambiguous business problems into measurable outcomes by modernizing decision strategies (rules → ML), building scalable data and ML pipelines, and driving adoption through clear measurement and storytelling. Known for building creative, scalable solutions that improve speed, reliability, and ROI—e.g., reduced recurring data errors from 2,000+ to <100, improved data reliability 95%+, saved 10–15 hours/week via automation, cut delivery cycle time 4 months → 2 months (50%), and delivered 25% improvement in forecasting accuracy.

INNOVATION & IMPACT HIGHLIGHTS

- Modernized an in-production rules/SQL decision approach toward ML: improved Best Prospects precision by 5% in <4 weeks and created a credible roadmap for leadership buy-in and rollout.
- Built Python-driven data quality frameworks that slashed recurring issues from 2,000+ to <100 and improved data reliability by 95%+ across integrated enterprise systems.
- Designed automation and intake workflows (Python, Azure DevOps, Power Automate) saving 10–15 hours per week for the analytics team and 1–2 hours per request for stakeholders.
- Developed a linear-programming scheduler to optimize report queues, cutting total build time from 4 months to 2 months (50% faster).
- Created an automated performance bias audit (hypothesis testing + NLP/sentiment analysis) reducing audit time from 80 hours to <2 while increasing rigor and repeatability.
- Built predictive models including 25% improvement in forecasting accuracy; and a deep learning model predicting post-surgical return risk within 48 hours with 80% test accuracy.

CORE SKILLS

Machine Learning: XGBoost, scikit-learn, classification/propensity, NLP (text features/sentiment), clustering/segmentation, linear programming

Decision Science: experimentation, KPI design, optimization (linear programming), model evaluation & monitoring

Data & Engineering: Python, SQL, Git, CI/CD, testing/validation frameworks, data quality & governance, data pipelines

Visualization: Power BI, Tableau; stakeholder-ready dashboards and executive storytelling

Platforms/Tools: Azure DevOps; Oracle Autonomous Data Warehouse (ADW); Google Cloud Platform

PROFESSIONAL EXPERIENCE

CINTAS — Data Scientist, IT Data Science & Engineering (Best Prospects) | Nov 2025—Present

- Owned the model modernization narrative from SQL scoring to an XGBoost proof-of-concept (PoC); communicated lift/precision tradeoffs and operational impact to stakeholders and leadership.
- Productionalized the current SQL-based Best Prospects model using Cloud Build and GitLab CI/CD for repeatable releases, controlled changes, and reliable deployment.
- Engineered novel features to represent territory saturation (e.g., customer counts within 1/3/5/10 miles of a prospect) and improved prospect ranking precision by 5% in <4 weeks.
- Delivered weekly stakeholder updates on performance, risks, and roadmap; built alignment across teams for a staged ML rollout.

FIRSTENERGY — Supervisor, HR Reporting & Analytics | Oct 2022–Oct 2025

- Built and scaled an analytics function, establishing governance, Agile delivery, and stakeholder partnership practices across HR reporting and analytics.
- Engineered data pipelines integrating Oracle ADW with 8 enterprise systems, improving data reliability by 95%+ and enabling analytics-ready datasets for advanced modeling.
- Automated analytics intake, QA, and delivery workflows (Python, Azure DevOps, Power Automate) saving 10–15 hours/week for the team; improved stakeholder turnaround by 1–2 hours per request.
- Implemented a Python-based data quality framework that decreased inconsistencies from 2,000+ to <100, increasing trust in dashboards and downstream analytics.
- Built a linear-programming scheduler that optimized report queues, cutting total build time from 4 months to 2 months (50% faster).
- Owned the end-to-end delivery of an executive STIP dashboard (Power BI/Tableau), translating compensation rules into scenario-based insights that drove a 5% base salary payout for 10,000+ employees.
- Partnered with executives to build an interactive geospatial map (Python + BI + API integrations) informing Return-to-Office strategy impacting 11,000+ employees.

FIRSTENERGY — HR Data Analyst III | Apr 2021–Oct 2022

- Created an automated performance audit tool using Python, hypothesis testing, and sentiment analysis to detect potential bias in evaluations of protected groups—reducing audit time from 80 hours to <2.
- Developed predictive analytics models using Python and SQL, achieving a 25% improvement in forecasting accuracy and influencing executive decision-making.
- Led a rapid-response Python remediation for STIP and STIP overtime calculations for 12,000 employees, shipping within 48 hours and enabling on-time payout processing.

BRIGHTHOSPITAL — Data Analyst (Contract) | Jan 2021–Apr 2021

- Built a deep learning model predicting post-surgical patient return risk within 48 hours using unstructured surgeon notes + history; achieved 80% test accuracy and that could have supported operating room scheduling decisions.
- Developed Tableau dashboards to surface OR scheduling inefficiencies and enable leadership to monitor turnaround and capacity.

THE NAVIGATORS — Fundraising Analyst & Team Lead | Jun 2006–Jan 2021

- Conducted geographic and donor-income analysis to identify high-potential fundraising regions; contributed to \$1M+ in fundraising revenue during tenure.
- Led and mentored teams of 8+ and built scalable programs, growing community groups from 3 to 100+ participants.

PERSONAL PROJECT HIGHLIGHT

PLOTWISE — Geospatial Analytics Product (Personal Project)

- Built an end-to-end geospatial analytics platform: automated data ingestion (Python/Selenium), structured storage, and interactive mapping dashboards for property insights and clustering-based market signals.
- Implemented feature engineering and reproducible pipelines (versioned code, automated runs, validation checks) to support scalable, production-style analytics delivery.

EDUCATION

M.S., Business Analytics — Kent State University

B.A., Mathematics — University of Cincinnati

B.S., Education — University of Cincinnati