

Jarod Ferkin

jarodferkin.com | jtferkin@gmail.com | linkedin.com/in/jarod-ferkin-740189201 | github.com/herodthegreater |

Winter Garden, FL · Open to relocation

SUMMARY

Data Science senior at UCF (graduating Spring 2027) who turns messy data into decisions. Ran consumer and market research for a local service business that **reversed its planned growth strategy**, built per-genre classifiers on **114K tracks** (cross-validated ROC-AUC up to 0.88), and designs experiments with pre-registered success criteria. Works in Python, SQL, and R, and explains results in plain language.

EDUCATION

B.S. Data Science **University of Central Florida** Expected Spring 2027
Coursework: Statistical Methods 3, Applied Time Series, Big Data Analytics, Algorithms for ML, Data Management, Data Visualization · A.A. General Studies, Valencia College

ANALYTICS & DATA SCIENCE PROJECTS

Consumer & Market Analytics · first client confidential Client Engagement September 2026 – Present

- Mined **2,500+ client and competitor Google reviews** for topic share: the service the business was really known for was not the one it planned to lead with, while the nearest competitor owned that one, evidence that **reversed a planned strategy**.
- Ran unit economics showing a price war with the category leader was unwinnable (matching its price meant a **58% cut** on the entry item) and shifted strategy to a multi-service membership.
- Modeled the trade area from **US Census ACS** data across 4 ZIP codes and mapped 6 competitors' offers and ads via the Meta Ad Library; spotted a two-peaked rating distribution (19 one-star vs. 4 two-star), a billing-friction signal to fix before any membership campaign.
- Closed a measurement gap (bookings captured **no acquisition source**) with a live **Supermetrics dashboard** reporting cost per booked visit, plus an automated weekly competitor watch; set up a 5-campaign A/B framework with pre-registered kill criteria tied to 60-day conversion.

Text/topic analysis, US Census ACS, competitive analysis, experiment design, Supermetrics, Claude API

Predicting Song Failure on Spotify · Per-Genre Classification December 2025

- Framed a binary "failure" target (popularity < 30) on a **114K-track Spotify dataset**; consolidated 114 sub-genres into 5 parent genres and built a separate model per genre (65K tracks) with stratified splits and variance-threshold feature filtering.
- Tuned and compared models: Random Forest reached a **5-fold cross-validated ROC-AUC of 0.63–0.88** by genre (0.88 for Latin), well above a Naive Bayes baseline (0.58–0.62); acousticness, speechiness, and duration drive failure differently by genre, and weak rock results traced to ambiguous labels, not the model.

Python, scikit-learn, pandas

Pigeon · Model-Compression Experiments Live: herodthegreater.github.io/pigeon July 2026 – Present

- Ran controlled experiments shrinking a YOLOv8n detector: a one-shot 5% channel cut drove bird mAP@50 to 0, while three 7.2% cuts with regrowth kept **67% of native with 31% fewer parameters**; reported FLOPs and measured latency separately after finding they diverge.
- Benchmarked compression of a MobileNetV3 bird classifier on a 500-image eval set: per-tensor INT8 **collapsed accuracy to chance (50%)** while per-channel INT8 held 94.8%; sensitivity-aware mixed precision reached **5.0x smaller at 94.4%**. Results ship on a live site.

Python, PyTorch, torch-pruning, FastAPI, ONNX Runtime Web, experiment design

AI-Native vs. Traditional Data Scraping Engine March 2026

- Built a 9-script pandas pipeline to **clean, transform, and evaluate** hundreds of historical records, benchmarking LLM-based (Claude API) extraction against traditional parsing, including a custom rowspan matrix-expansion algorithm for messy tables.

Stall Card (formerly barnManager) · Barn Operations Platform Live: barnmanager.fly.dev June 2026 – Present

- Designed the **relational data model** (SQLite via PocketBase) for a live operations app built for a working barn: a breeding event log with aggregates **derived from events so they can never drift**, evolved through **20 versioned migrations** as client feedback came in.
- Built a grounded question-answering feature: Claude Haiku answers plain-language questions from the barn's own records only, returns horse links only for horses in that barn, and says so when the records don't answer.

SQLite, PocketBase, TypeScript, React, Claude API

EXPERIENCE

DriveShack June 2024 – April 2026

Manager / Front-of-House Team Lead · Orlando, FL

- Grew revenue **20%+ in one fiscal year** as part of a leadership turnaround team, implementing **data-driven budgeting**.
- Bridged logistics teams, staff, and clients on six-figure events; designed training that ended reliance on external hiring.

SKILLS

Languages & Tools: Python (pandas, scikit-learn, PyTorch), SQL, R · Excel (advanced formulas, pivot tables), Power BI · SQLite, Scrapy, Claude API
Statistics & ML: classification, cross-validation, model evaluation (ROC-AUC, F1), A/B testing & experiment design, time series, segmentation, text/topic analysis, Census/ACS analysis