retsci

End of studies projects 2027

Tunisia



About RetSci

Retail Science. Simplified.

We are a pioneering retail consultancy, based in the US and operating globally since 2015.
We are retail experts with a passion for making advanced analytics & data science easy and accessible to all retailers & wholesalers. We help you get the most value from your data and technology while building tools and processes to support your teams.

Our Services

- Software Implementation
- AI & Retail Science
- Managed services

internships Topics agenda

1 AI-Assisted Forecasting for Intermittent Retail Demand



Smarter Oracle Forecasting: Parameter Tuning by Product Behavior 2

3 **Automated Forecast Diagnostics:** Building a Root-Cause Analysis Tool

AI-Assisted New Item Forecasting: Finding Look-Alike Products with Images and Attributes 4

AI-Assisted Forecasting for Intermittent Retail Deman

Project Description

Around 80% of retail items at forecast level sell irregularly, with many days of zero sales, yet they are forecast with the same generic methods used for steady sellers. In this project you will design and build a dedicated forecasting approach for these items that predicts when a sale is likely to happen and how large it will be. You will benchmark it against the current Oracle system forecast on real retail data, working with the same AI development tools our team uses (Claude, Cursor, Databricks).

Expected Deliverables

1) Analysis and segmentation of intermittent items in a client dataset (sales frequency, zero-sales patterns, volumes)
2) Definition of a dedicated forecasting process for intermittent items, from data preparation to forecast output
3) Probabilistic forecasting model predicting the timing and size of sales
4) Benchmark against the Oracle system forecast using accuracy metrics suited to intermittent demand
5) Documented, production-ready code in Databricks
6) Final presentation of results and recommendations

Required Skills

1) Statistics and probability (Markov Chains, distributions, estimation)
2) Python (pandas, NumPy) and SQL
3) Forecasting basics and error metrics (MAE, bias, MASE)
4) Data analysis and visualization
5) Clear written and spoken English

Nice-to-have Skills

1) Intermittent demand methods (e.g., Croston, TSB)
2) Databricks / PySpark
3) Retail or supply chain knowledge
4) Probabilistic or zero-inflated models
5) Using AI coding assistants (e.g., Claude Code, Cursor)
6) Git / GitHub

Smarter Oracle Forecasting: Parameter Tuning by Product Behavior

Project Description

Oracle Retail AI Foundation (AIF) offers many forecasting parameters, but today the same generic settings are applied to every product group, from fast-moving basics to slow-selling fashion items. You will build a data-driven method to find the best parameters for each type of product and validate it directly in Oracle's latest demo environment. Your results will feed into real client implementations. You will build it with AI tools (Claude, Cursor, Databricks).

Expected Deliverables

1) Product characterization framework grouping items by demand behavior (velocity, seasonality, volatility, lifecycle)
2) Tuning pipeline in Databricks to test parameter combinations for each product group
3) Parameter matrix mapping each product profile to recommended AIF settings
4) Validation in the Oracle test environment, comparing accuracy against current generic settings
5) Playbook for reusing the matrix on new client implementations
6) Final presentation of results and recommendations

Required Skills

1) Python (pandas) and SQL
2) Time series forecasting fundamentals and error metrics
3) Experiment design and model tuning basics
4) Data analysis and reporting
5) Clear written and spoken English

Nice-to-have Skills

1) Oracle Retail or other enterprise forecasting software
2) Databricks / PySpark and Delta tables
3) Clustering and segmentation methods
4) Optimization tools (e.g., Optuna, Bayesian search)
5) Retail or CPG knowledge (fashion, assortment, promotions)
6) Using AI coding assistants (e.g., Claude Code, Cursor)

Automated Forecast Diagnostics: Building a Root-Cause Analysis Tool

Project Description

When a forecast goes wrong, diagnosing the underlying cause often requires hours of tedious manual investigation. In this project, you will design and implement a post-forecast decision-tree diagnostic engine that automatically identifies, categorizes, and reports forecasting anomalies.

The engine will flag critical issues such as seasonality anomalies, inflated demand, eligibility errors, and outlier-handling gaps. The resulting diagnostic reports will highlight primary issue types, their frequency, and the product segments most heavily impacted—enabling teams to systematically improve forecasting accuracy and proactively address performance deterioration.

You will build upon our existing codebase and data pipeline, with the ultimate goal of delivering a production-grade tool ready for integration into an Oracle-native marketplace product. Development will be accelerated using modern AI developer tooling (e.g., Claude Code, Cursor, and Databricks).

Expected Deliverables

1)Decision tree design (problem, symptom, root cause, detection logic, and thresholds)
2)Decision tree execution codebase
3)Testing and validation
4)Code documentation and maintenance guide
5)Final presentation of results and recommendations

Automated Forecast Diagnostics: Building a Root-Cause Analysis Tool

Required Skills

1) Python and SQL
2) Data validation and quality checking
3) Forecasting basics (seasonality, outliers, error metrics)
4) Writing clean, maintainable code
5) Clear written and spoken English

Nice-to-have Skills

1) Databricks / PySpark, PL/SQL
2) Anomaly detection or rule-based systems
3) Git basics
4) Dashboards and data visualization
5) Retail or CPG knowledge
6) Using AI coding assistants (e.g., Claude Code, Cursor)

AI-Assisted New Item Forecasting: Finding Look-Alike Products with Images and Attributes

Project Description

Forecasting new products is one of the most frequent challenges in retail: with no sales history, planners manually search for similar past items and adjust their demand. In this project you will build a proof of concept that automates this search. It will combine existing product attributes with image analysis to identify the most similar historical items and use them to generate a starting forecast for new products. You will build it with AI tools (Claude, Cursor, Databricks).

Expected Deliverables

1) Assessment of how well existing product attributes (class, subclass, color, price, etc.) already identify similar items
2) Image-based similarity approach to refine and narrow down the pool of candidate items
3) Look-alike matching pipeline that ranks the most similar historical items for each new product
4) New item forecast built from matched items, compared against the current approach
5) Scalability assessment for running the approach on full client assortments
6) Final presentation of results and recommendations for next steps

Required Skills

1) Python (pandas) and SQL
2) Machine learning fundamentals (similarity, clustering)
3) Basics of image processing or computer vision
4) Data analysis and forecasting basics
5) Clear written and spoken English

Nice-to-have Skills

1) Pretrained image embedding models (e.g., CLIP, ResNet)
2) Databricks / PySpark
3) Retail or fashion merchandising knowledge
4) Vector search or nearest-neighbor methods
5) Using AI coding assistants (e.g., Claude Code, Cursor)
6) Git / GitHub

We look forward to getting to know you!

