Demand Forecasting in the Age of AI

PURDUE UNIVERSITY | The Data Mine

Suhani Dadhich (TA), Aarush Sachdeva, Beck Chandler, Cooper Roach, Daniel Welch, Emily Henderson, Hongda Yuan, Jack Swaisgood, Keya Jadhav, Rupali Gadage, Ruyue Jiang

RED GOLD QUALITY

01 Introduction

- **About Red Gold:** Leading U.S. tomato processor relying on accurate forecasting to drive operations
- Red Gold's current model: Silvon
- Red Gold's current forecast accuracy: **75 - 77%**
- **Objective:** Improve forecast accuracy using statistical and AI-driven models beyond the current baseline

02 Dataset

- **Data Scope:** 3.5 years of SKU-level data (Jan '21 - Jun '24)
- Weekly sales volume across all SKUs
- Captures seasonality indicators and historical demand patterns
- **Method:** Dynamic Time Warping (DTW) clustering to group SKUs by time-series shape
- Cluster labels assigned to each SKU for segmentation

03 Modeling Teams

LSTM: Deep learning for nonlinear patterns
ARIMA: Statistical time-series modeling
AI Team: Feature engineering + advanced models like Artificial Neural Networks (ANN), Ollama (large language model) and Prophet (by Meta)

04 Toolbox

Python, Pandas, NumPy, Scikit-learn, TensorFlow/Keras (LSTM), Statsmodels (ARIMA), Prophet (GAM), Dynamic Time Warping (DTW), XGBoost

05 Methodology

Phase 1: Data inspection, literature review, and basic modelling

- **Step 1:** Understood dataset attributes and fiscal vs calendar year
- **Step 2:** Identified popular industrial methods for demand forecasting
- **Step 3:** Deployed methods including basic ANN, Ollama, Prophet, XGBoost, ARIMA and LSTM

Phase 2: Ensemble methods, advanced models, and a routing system

- **Step 1:** Tried feature engineering on the variable Label PSS Case Pack to improve the accuracy of the basic models
- **Step 2:** Combined ABC-XYZ and DTW to improve model routing beyond a single-model approach
- **Step 3:** Improved prediction accuracy with DTW clusterings
- **Step 4:** Evaluated optimal cluster count through statistical evidence and tested clustering methods like K-Means
- **Step 5:** Adopted advanced forecasting models like Croston's, Holt-Winters, ARIMA, LSTM aligned with DTW-based cluster patterns, using grid search for parameter tuning
- **Step 6:** Used volume-based segmentation with ABC-XYZ and DTW for finer-grained modeling

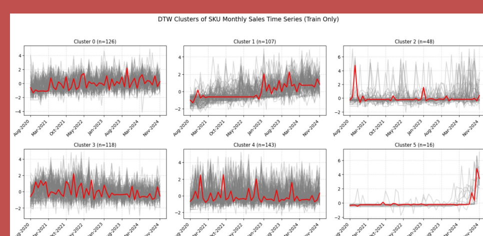

Fig 1. DTW Clustering

06 Key Insights

- DTW clustering significantly improved forecasting performance across models
- ARIMA achieved the highest peak accuracy, reaching up to 89% in specific clusters
- **All models achieved competitive performance:** ARIMA: 77.97%, AI: 76.67%, LSTM: 75.42%
- Model performance varies by cluster; segment-specific modeling is recommended

07 Next Steps

- Incorporate seasonality insights to capture recurring demand patterns
- Enable dynamic updates for changing data
- Build a scalable forecasting system
- Ensure the model adapts effectively to new SKUs

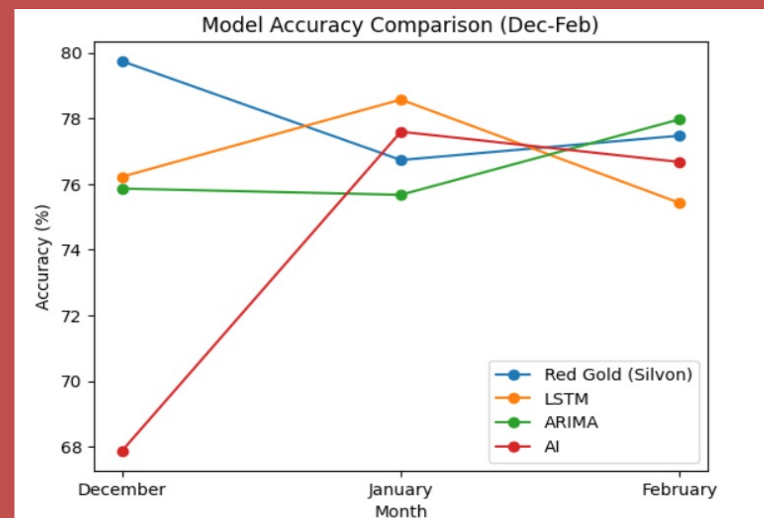

Fig 2. Accuracy Graph

08 Acknowledgements

We would like to extend a special thank you to our Red Gold mentors, Kevin, Logan, Nick, and Brian, for their guidance and support, our Data Mine liaison, Cai, Data Science Team Member, Ashley, for her invaluable contributions, Professor David, Griswold Consulting, our TA, Suhani for all her encouragement throughout the semester, and the entire Data Mine Staff for all their assistance.

The Data Mine Corporate Partners Symposium