



LABOR ANALYTICS FORECASTING

GT x THD ANALYTICS PRACTICUM
FINAL REPORT

Introduction and Problem Statement

Labor forecasting helps The Home Depot schedule the appropriate amount of labor at each store and time of year. Accurate demand forecasts support customer service and sales while reducing unnecessary labor costs. This is particularly important for Outdoor Garden, where demand changes sharply with seasonality, weather, and major promotional events.


Group 5 was assigned to forecast Outdoor Garden units for Fiscal Years 2025 and 2026. Our first objective was to evaluate the accuracy and limitations of the FY2025 forecast. We then used those findings, along with feedback from Home Depot, to develop an improved 52-week forecast for FY2026.

The final approach combines historical demand, weather, fiscal seasonality, promotional events, and store characteristics. We evaluated multiple forecasting methods against historical baselines and used rolling backtests to select the final model.

Data Description

Dataset	Category	Description
Outdoor Garden Units	Volume Drivers	Contains the main target (label) to forecast.
Store Details	Internal Data	Contains detail about The Home Depot stores across regions.
Weather Data	External Data	Contains extensive information about weather attributes during fiscal periods.
Fiscal Calendar	Internal Data	Contains detail on how various time-series (fiscal) attributes are used.
Holiday and Event Tags	Events and Promotions	Indicates whether certain repeating/non-repeating events fall at a fiscal period.
Sister-Store Mappings	Internal Data	Identifies comparable existing stores used to forecast new locations and provides opening dates.

Methodology



Due to presence of real-world complexities such as events and weather, the database tables, metadata and data points themselves revealed a series of complexities. Before diving actively into the practical tasks, it was critical to understand what the data actually means. A variety of steps were taken to make sure the team had a proper understanding of what the data reveals and how it can help us in forecasting. These include the following:

- Working with sponsor documentation and office hours to better understand the labor forecasting problem and available datasets.
- Identifying Outdoor Garden Units as the primary forecasting target tied to labor planning.
- Reviewing data dictionaries and schema documentation to understand dataset structure and relationships.
- Evaluating supporting drivers including weather, fiscal calendar, store attributes, and holiday/event features.
- Mapping dataset relationships using key fields such as store, calendar date and fiscal time periods to support downstream integration

Modeling and Analysis

The final forecasting approach uses a hybrid, event-aware system to predict weekly Outdoor Garden demand at the store level and then allocate those weekly forecasts to individual store-days. Outdoor Garden demand exhibits strong annual seasonality, sharp spring peaks, local weather sensitivity, movable promotional events, and substantial variation across stores. No single model addressed all of these challenges reliably. The final system therefore combines a nonlinear XGBoost point model with a national SARIMAX/seasonal backbone, robust structural demand anchors, cross-fitted calibration, controlled recursive forecasting, and a separately validated daily allocation model.

Data Preparation and Feature Engineering

We combined historical Outdoor Garden units with weather, the fiscal calendar, store details, holiday and event tags, and the sister-store file provided by Home Depot. Before modeling, we standardized store numbers and dates, removed duplicates, checked for missing data, and confirmed that each record matched the correct fiscal week.

The final model incorporated updated sponsor-provided event data. Specific internal file names and event-calendar corrections are omitted from this public copy.

We created features to represent recent demand, prior-year demand, longer-term seasonal patterns, weather, store characteristics, and event timing. The model also included measures of how quickly spring weather was arriving in each market. This helped it account for the fact that the Outdoor Garden season begins at different times across the country.

Sponsor feedback informed treatment of atypical promotional periods and the approach used for newer locations.

Forecasting Models

The final forecast used a hybrid model rather than relying on one forecasting method.

XGBoost served as the main store-level model. It learned relationships between demand and factors like recent sales, prior-year patterns, weather, spring timing, promotional events, and store characteristics. XGBoost was useful because many of these relationships are nonlinear. For example, warm weather may have little effect in January, but a much larger effect during the spring selling season.

Promotional events were modeled separately from normal weekly demand, with broader geographic patterns used when store-level event history was limited.

SARIMAX played a supporting role. It produced a stable national demand forecast based on historical seasonality, weather, and event timing. That forecast was then distributed to stores and used as a reference point for the store-level model. In other words, XGBoost captured the more detailed store-level relationships, while SARIMAX helped keep the overall forecast at a reasonable level.

The final weekly forecast combined [redacted] of the XGBoost prediction with [redacted] of a more conservative

historical baseline. This small baseline component helped reduce the risk of extreme predictions without removing the improvements produced by XGBoost.

The model generated forecasts one week at a time for the full 52 week horizon. A separate day-of-the-week model then divided each store's weekly forecast across its active days. This produced the required daily submission while ensuring that daily values still added back to the weekly forecast.

Model Evaluation and Forecast Results

We evaluated the final model using rolling backtests for FY2023, FY2024, and FY2025. For each test year, the model was trained using only information that would have been available before that year. This allowed us to measure performance under conditions similar to a real future forecast. The final FY2025 backtest produced the following results:

Metric	FY2025 result
Full-year WAPE	REDACTED
Average Weekly Store Mean MAPE	REDACTED
Average Weekly Store Median MAPE	REDACTED
Overall bias	REDACTED
Spring WAPE	REDACTED
Spring Black Friday WAPE	REDACTED
Peak-timing error	REDACTED

Backtesting showed competitive performance against comparison benchmarks and improvement over the strongest historical baseline. Exact benchmark and performance values are omitted from the public copy.

patterns into the FY2026 forecast. Because FY2026 actual demand was not yet available, these graphs should be interpreted as forecast results rather than measures of realized accuracy.

FORECAST VISUALIZATION OMITTED

Confidential demand values removed from public copy.

Figure 1. FY2026 national weekly Outdoor Garden forecast with baseline, model, and forecast-range comparisons.

The national forecast shows the expected seasonal pattern for Outdoor Garden demand. Units rise steadily through February and March before reaching the largest peak in April. A second, smaller increase appears later in the spring, followed by a gradual decline through summer, fall, and winter.

The blue line represents the XGBoost model forecast, while the darker baseline line provides a more conservative historical reference. The final forecast remains close to the XGBoost prediction but

retains a small contribution from the baseline to reduce the risk of extreme results. The Forecast Value and P50 lines overlap because P50 represents the central forecast.

The P10 and P90 lines show the model's forecast range based on historical error patterns. This range is widest during the spring peak, when Outdoor Garden demand is both highest and most difficult to predict. It becomes narrower during the more stable second half of the year.

FORECAST VISUALIZATION OMITTED

Confidential regional demand values removed from public copy.

Figure 2. FY2026 weekly Outdoor Garden forecast for Georgia with baseline, model, and forecast-range comparisons.

The Georgia forecast follows the same broad seasonal pattern as the national forecast, but at a regional scale. Demand increases rapidly during March, reaches its largest peak in April, and then gradually declines through the remainder of the year.

The timing of Georgia's peak is consistent with the national pattern, while its forecast level reflects the state's own stores, weather conditions, and historical demand. This illustrates how the model combines national seasonality with more local information rather than applying the same forecast pattern uniformly to every store.

The greater distance between the P10 and P90 lines during April also shows that uncertainty is concentrated around the spring selling peak. Outside the spring season, the model, baseline, and central forecast move more closely together.

Overall, the backtests show that the final model improved on the historical baselines, while the Forecast Explorer demonstrates how the model applies the expected national seasonal pattern at a more local level. The model also passed the required checks for accuracy, bias, peak timing, long-range stability, new-store treatment, and daily-to-weekly reconciliation. Final FY2026 accuracy will need to be measured once actual demand becomes available.

Deployment

The forecasting workflow was implemented through forecasting application , a Streamlit-based application used to validate data, train models, generate forecasts, and download the final output files. The Tableau dashboard provides a separate presentation layer for exploring the forecast. Users can review demand by store, region, division, or fiscal week and compare historical and forecast patterns. Tableau uses the validated forecast produced by forecasting application rather than training the model itself.

The main project deliverable is the daily forecast file, final daily forecast file . The pipeline also produces weekly forecasts and validation files documenting data quality, forecast integrity, and daily-to-weekly reconciliation.

Key dashboard features include:

- Outdoor Garden forecasts by store and fiscal week
- Historical and forecast demand comparisons
- Regional and divisional summaries
- Weather and event indicators
- Filters for store, region, division, and fiscal week
- Downloadable forecast tables

Scenario analysis was treated as an optional enhancement and was not required for the final submission.

Lifecycle Management

The forecasting solution was designed to support ongoing updates as new sales, weather, and event information become available. Rather than rebuilding the forecasting model each year, the existing forecasting pipeline can be rerun using updated historical data to generate new predictions for future fiscal periods.

To maintain forecast quality, model performance should be monitored after each forecasting cycle by comparing forecasted Outdoor Garden demand with actual sales once they become available. Performance metrics such as Average Weekly Store MAPE, Store-week WAPE, Bias, MAE, and RMSE can be recalculated to measure forecast accuracy and identify opportunities for improvement.

Regular monitoring also enables identification of changes in customer purchasing behavior, weather patterns, or promotional events that may reduce model performance over time. If forecast accuracy begins to decline, the forecasting pipeline can be updated by incorporating new historical data, engineering additional predictive features, adjusting model hyperparameters, or retraining the forecasting models.

Data quality should also be monitored throughout the forecasting process. Incoming sales, weather, fiscal calendar, and holiday/event datasets should be validated for missing values, duplicate records, and consistency before forecasts are generated. Automated validation checks can help identify potential data quality issues early and improve the reliability of forecast outputs. The forecasting pipeline was designed to be modular, allowing individual components such as data preparation, feature engineering, model training, forecast generation, and dashboard updates to be maintained independently. This approach simplifies future enhancements while minimizing disruption to the overall forecasting workflow.

As new forecasts are generated, the Tableau dashboard can be refreshed with the latest prediction data, allowing business users to continuously access updated demand projections through the same interface. This lifecycle management process supports continuous improvement while ensuring the forecasting solution remains accurate, maintainable, and scalable as business conditions evolve.

Conclusion

Group 5 developed a complete FY2026 forecasting system for Outdoor Garden units. The final hybrid approach combined a store-level XGBoost model with a national SARIMAX backbone, historical

In historical backtesting, the final model improved on comparison baselines and identified peak timing accurately. Sponsor feedback and data-quality checks were incorporated into the final forecasting workflow.

The final system produced daily forecasts across the full FY2026 horizon. All required data-quality, forecast-integrity, and reconciliation checks passed before submission.