https://doi.org/10.31705/BPRM.v2(1).2022.8

Retail Demand Forecasting using Light Gradient Boosting Machine Framework

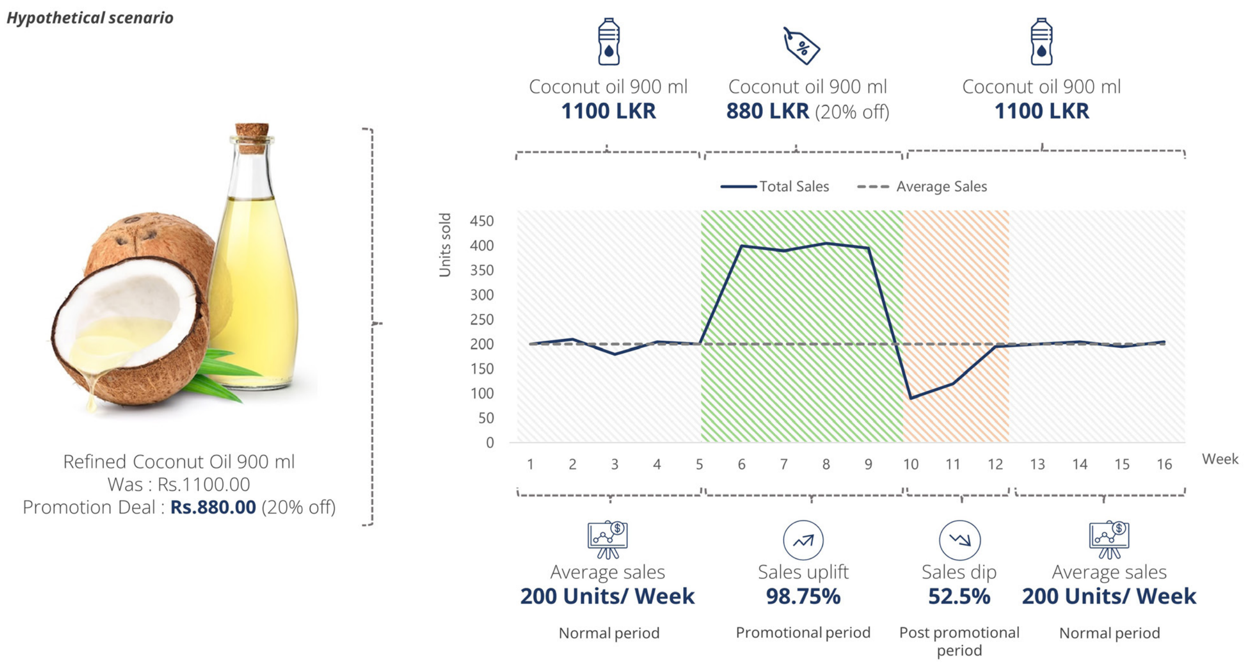

Background

Preparing product-level demand forecasts is crucial to the retail industry. Importantly, reliable inventory and replenishment decisions for retail products depend on accurate demand forecasts. This allows retailers to enable better pricing and timely promotion plans while leading to huge cost reductions [1]. Often, retail promotions create demand irregularities for products. Customers may change their buying behaviour by purchasing more products for future consumption (stockpiling), thereby increasing sales in the promotional period. Then, for a brief time, sales may fall below normal levels before gradually returning to normal levels. The period with a dip in demand is known as the post-promotional period [2]. Thus, a retail promotion has three distinct periods: (1) normal, (2) promotional, and (3) post-promotional, each with its own set of demand fluctuations (see Figure 1).

The common practice in the retail industry is using univariate approaches such as Auto-Regressive Moving Average (ARIMA) and Exponential Smoothing (ES). They often use human judgment to predict demand in promotional periods. This can result in systematic errors leading to ineffective demand forecasts. In contrast, multivariate models can incorporate promotional information along with demand forecasts. There are several sophisticated regression-based retail forecasting methods employed in the industry such as "SCAN*PRO", "CHAN4CAST" and "Promo-Cast™". However, these models are not that popular in the practice due to the high investment and the complexity associated with the models [3].

Machine learning (ML) based methods have risen as a promising alternative for demand forecasting over the past decade (e.g., Support Vector Machines, Regression Trees, Gradient Boosting, Neural Networks, etc.) [4]. Although ML models require high computational capacity, they offer flexibility and high predictive accuracy in the face of a large number of observations. Thus, ML models are a viable option for predicting retail demand, which is a complex process due to the high number of products available across many stores [1].

Over the past few years, Gradient-boosted Regression Trees (GBRT) have become one of the most popular ML methods. Light Gradient Boosting Machine (LightGBM) is one of the most popular implementations among them [4]. In fact, LightGBM was one of the winning methods of the recently concluded "M5 Forecast Competition" [4]. It is also considered a viable substitute for Artificial Neural Networks (ANNs) [1]. Thus, we employed LightGBM to determine whether it is appropriate for retail demand forecasting in the context of retail promotions.

Almost all previous studies focus on the promotional period, regardless of the normal and post-promotional periods. As a result, comparing the forecast performance of ML approaches to univariate methods in each period is difficult. Therefore, the primary contribution of our study is to analyse and compare the two commonly implemented univariate methods (i.e., ES and ARIMA) with LightGBM implementation in the presence of retail promotions.

Model Development

We collected sales data of 156 weeks from a US-based retailer. The dataset comprises sales data including promotional information for 50 products in the top 4 product categories across 75 stores. We used a Moving Average (4-period) approach to predict baseline demand considering demand during normal periods. Then, we used the promotion calendar to separate promotional periods from normal periods. We used Eq. (1) to classify post-promotional periods.

T_k = (Baseline projection at k^{th} period) – (Actual demand k^{th} period) (1)

If T_k is negative immediately after promotion, k_{th} week is named as a post-promotional period ($k \geq 0$). Table 1 provides an overview of the features, with the first set of features being timeseries-specific and the rest being business-specific. Furthermore, we developed an additional feature (i.e., the promotional period as a variable) to see if it improves the LightGBM model.

Type	Feature
Timeseries features	Date, Week 1 lagged sales, Week 2 lagged sales, Week 3 lagged sales
External features	Store ID, Product ID, Product category ID, Product subcategory ID, Product discount size, Display, Feature, Temporary price reduction
Additional feature	Normal period, Promotional period, Post-promotional period

Table 1: Model Features

We implemented 4 candidate models using different feature combinations in our study;

ES - sales data
ARIMA - sales data
LightGBM1 - timeseries, external and additional features
LightGBM2 - timeseries and external features

As univariate methods, we employed ES and ARIMA models using auto.arima() function and ets() function in R forecast package. We used Sckit-learn API to implement LightGBM model.

 https://doi.org/10.31705/BPRM.v2.2021.1

Analysis and Results

In our analysis, we focused on two main areas.

First, we provide a summary analysis of the magnitude and direction of the promotion effect identified by each model using Eq. (2).

$$\text{Forecast Magnitude} = (\text{Final Forecast} - \text{Baseline Forecast}) / \text{Baseline Forecast} \qquad (2)$$

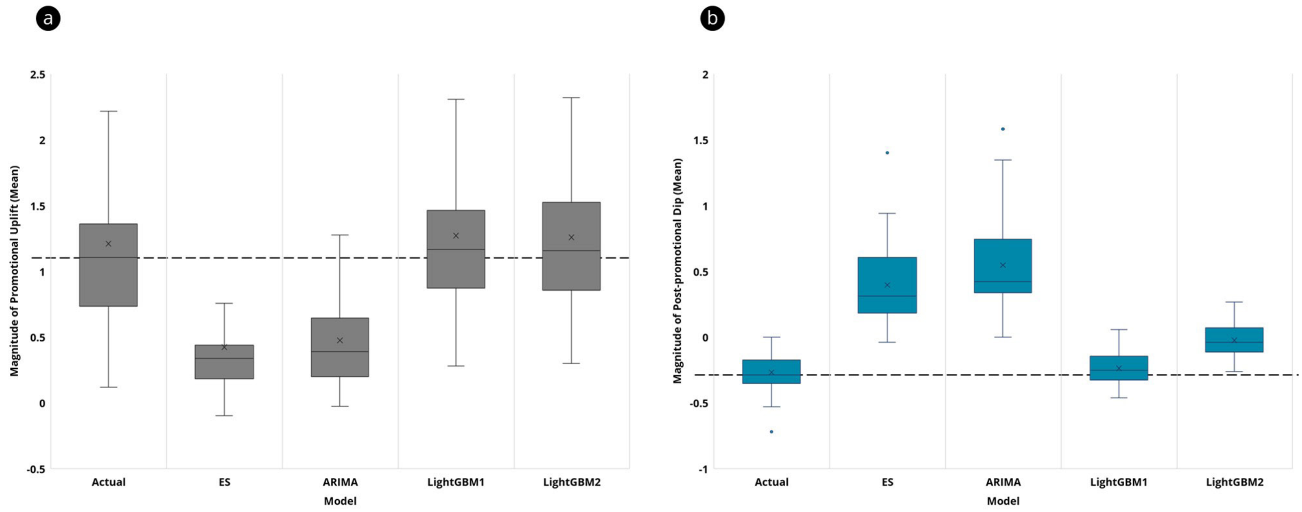

Figure 2: (a) Magnitude and Direction of Promotional Uplift; (b) Magnitude and Direction of Post-promotional Dip

Figures 2 (a) and (b) show the size and direction of the promotional and post-promotional effects identified by each model. It clearly demonstrates that LightGBM models were able to identify correct promotional uplift (ANOVA: $F(4,358) = 17.82$, $p < .000$) and post-promotional dip (ANOVA: $F(4,358) = 129.5$, $p < .000$) compared to univariate methods. However, only LightGBM1 was able to identify the correct size and sign of the post-promotion dip (mean post-promotion dip = -0.23, $p < .000$). Thus, providing the promotional period as a variable has improved the ability to identify the post-promotional dip of the LightGBM model.

Second, we compare the accuracy of forecasts using sMAPE (Symmetric Mean Absolute Percentage Error) as per Eq. (3).

$$\text{sMAPE} = \frac{\sum_{t}^{n} \dfrac{|f_t - a_t|}{(|f_t| + |a_t|)/2}}{n} \text{ X } 100$$

where: A_t : actual demand at t^{th} week, f_t : forecasted demand at t^{th} week and n: number of samples.

Figure 3 depicts the comparison of mean values of sMAPE for all the models. Noticeably, in all periods, both LightGBM models show a significant performance compared to ES and ARIMA models (ANOVA: normal period - $F(4,358) = 16.1$, $p < .000$; promotional period - $F(4,358) = 32.92$, $p < .000$; post-promotional period - $F(4,358) = 50.32$, $p < .000$). In both normal and post-promotional periods, the LightGBM1 model outperforms the LightGBM2 model significantly. However, both LightGBM models performed similarly during the promotional period. Thus, the addition of a promotional period as an additional variable improves the prediction performance of LightGBM models during both the normal and post-promotional periods.

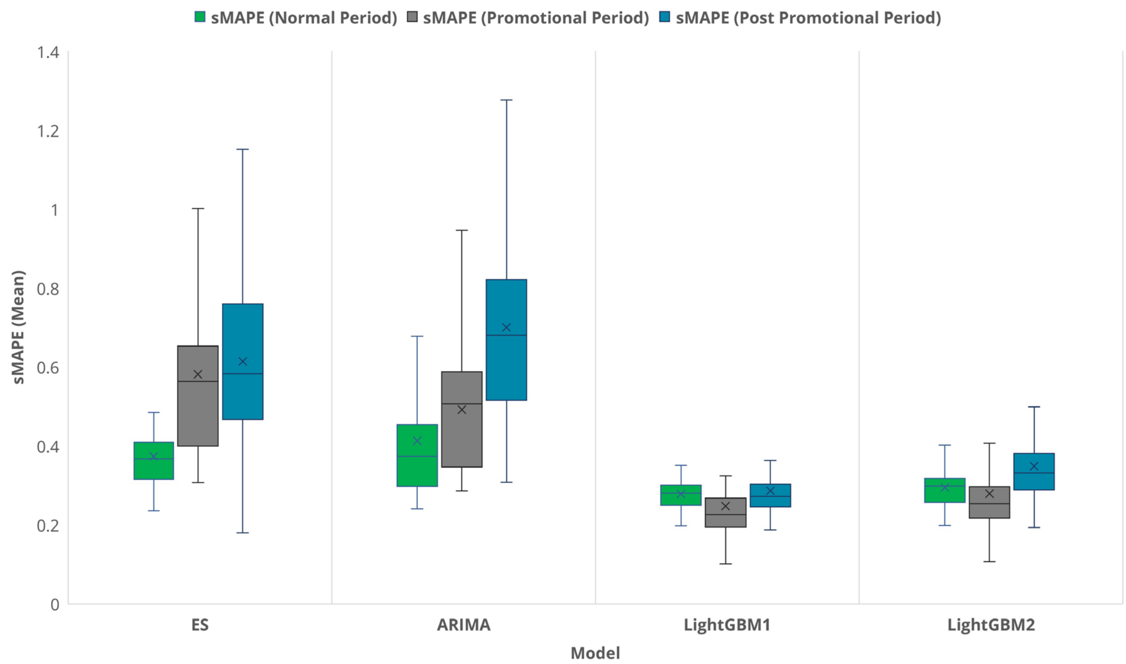

Figure 3: Comparison of Forecast Performance using sMAPE

The Way Forward

Results of our study reveals that the LightGBM models improved the forecasting performance across all three periods compared to the established univariate models. At any given time, retailers often manage thousands of SKUs across several stores, making it redundant to use univariate methods. Our study shows that the LightGBM model can automate the retail sales forecasting process in the presence of sales promotions. Thus, retailers no longer require extra time and effort to make forecast adjustments to univariate methods to cope with promotional periods. Our findings indicate that LightGBM model can automate the forecasting process and provide significant performance even with the standard approach. Hence, we demonstrate the way retailers can successfully apply LightGBM method in forecasting sales.

References:

[1] R. Fildes, S. Ma, and S. Kolassa, "Retail forecasting: Research and practice," Int. J. Forecast., no. xxxx, 2019, doi: 10.1016/j.ijforecast.2019.06.004.

[2] H. C. Hewage, H. N. Perera, and S. De Baets, "Forecast adjustments during post-promotional periods," Eur. J. Oper. Res., vol. 300, no. 2, pp. 461–472, Aug. 2021, doi: 10.1016/J.EJOR.2021.07.057.

[3] H. N. Perera, J. Hurley, B. Fahimnia, and M. Reisi, "The human factor in supply chain forecasting: A systematic review," Eur. J. Oper. Res., vol. 274, no. 2, pp. 574–600, 2019, doi: 10.1016/j.ejor.2018.10.028.

[4] S. Makridakis, R. J. Hyndman, and F. Petropoulos, "Forecasting in social settings: The state of the art," Int. J. Forecast., vol. 36, no. 1, pp. 15–28, 2020, doi: 10.1016/j.ijforecast.2019.05.011.

Article by

Chamara Hewage, Niles Perera

Center for Supply Chain, Operations and Logistics Optimization, University of Moratuwa, Sri Lanka