Journal of Technological Information, Management & Engineering Sciences (2020),1(2), 53-68

Developing an AI-Driven Demand Forecasting Model for Perishable Goods in Fragmented Retail Markets

[1] **Dr. Imran Yousaf**

[2] **Sidra Amin**

How to cite this article:

Yousaf, I., & Amin, S. (2020). Developing an AI-Driven Demand Forecasting Model for Perishable Goods in Fragmented Retail Markets. *Journal of Technological Information, Management & Engineering Sciences, 1(2), 53-68.*

Received	Accepted	Published
17 August 2020	21 September 2020	26 November 2020

Abstract

The problem of demand forecasting for perishables is an operational problem which is difficult because the product has a finite shelf life, the demand is subject to rapid variations and if the demand is not met, the stock is lost and dissatisfied customers. In decentralized retail markets, such as small shops, independent retailers, local wholesalers, and informal traders, these challenges are particularly acute, as they may lack proper records, digital systems, and market information. This paper presents an artificial intelligence (AI) based demand forecasting system comprising a time-series learning approach, machine-learning attributes, a hierarchical information sharing method, and inventory-based decision rules. Historical sales, price and promotion data, calendar effects, weather, local events, replenishment lead times, stock availability, and the characteristics of product shelf-life are integrated into the proposed system. A comparison of interpretable statistical baselines, machine-learning methods, and a hybrid forecasting model explicitly tackles censored demand due to stockout, limited data, new products, and store-level heterogeneity. Quality of the forecast is assessed using scaled error measures, probabilistic calibration and operational measures (waste, service level, lost sales). A research design and hypotheses are provided but no original retail data are provided. The concept of the framework is that the performance of a forecast should be evaluated with regard to ordering decisions and perishability, rather than simply as a statistical goal. A practical implementation should be simple to use, understandable, non-intrusive and provide practical forecasts even if information is scarce at specific stores.

Keywords: *Food perishability, fragmented retail, artificial intelligence, probabilistic forecasting, machine learning, inventory optimization, food waste, supply chain*

[1] Department of Artificial Intelligence, University of Agriculture, Faisalabad, Pakistan

[2] Department of Computer Science, Government College University Faisalabad, Faisalabad, Pakistan

Copyright (c) 2020 Author(s)

This work is licensed under a Creative Commons Attribution 4.0 International License.

1. Introduction

Demand forecasting assists retailers make ordering, reordering and product allocation decisions. These decisions are time-sensitive for those products which are perishable, such as fresh produce, dairy, meat, bakery products, ready-to-eat foods and some pharmaceuticals. Overstocks can spoil or expire and understocks result in stockouts and lost sales. The economic consequences of an error in a forecast also depend on the direction of the error, on the margin of the product, on its shelf life, on the frequency of replenishment and on the response of the customers.

Poorly structured retail markets cause forecasting to be less accurate. Many independent retailers have limited sales records, different product descriptions and incomplete records of transactions, as well as several wholesalers. In addition to integrated point of sale systems, some stores might need to use handwritten records or messaging applications. Changing wholesale prices, holidays, religious festivals, neighborhood events, school calendars, paydays, and weather are also factors that can affect local demand. A model that is trained on a large chain's central data might not take well when applied to such environments.

However, traditional forecasting methods like moving averages, exponential smoothing and seasonal time-series forecasting are still useful where the forecasting pattern is stable and is transparent. Machine-learning techniques are able to include nonlinear relations and external factors, and deep-learning techniques are able to learn temporal patterns in a large number of related series. However, more complicated models do not necessarily take better effect. Poor performance can be caused by sparse data, inconsistent product identifiers, stockouts, price changes, and structural breaks. A practical AI system must therefore be augmented by data-quality controls, appropriate baselines, uncertainty estimates, and inventory-aware decision making.

This paper suggests an AI-enabled architecture that is suitable for the disintermediation of retailers. The system utilizes store-item forecasting, synchronizes the information that is shared between related stores and products, includes external drivers (when available), and transforms forecasts into

[1] Department of Artificial Intelligence, University of Agriculture, Faisalabad, Pakistan

[2] Department of Computer Science, Government College University Faisalabad, Faisalabad, Pakistan

Copyright (c) 2020 Author(s)

This work is licensed under a Creative Commons Attribution 4.0 International License.

replenishment suggestions that take into account shelf life and out-of-stock risk. It focuses on low cost deployment, explainability, and operation assessment. The contribution is an architecture and a methodology that could be tested in practice, rather than a reduction in waste or in forecasting error that has already been accomplished by the proposed model.

1.1 Research Objectives

The goals are: (1) to design a forecasting architecture that is appropriate for sparse and heterogeneous retail data; (2) to use time-series and machine-learning techniques along with the shared information across stores and products; (3) to develop a validation design that avoids temporal leakage and considers demand availability and stock-outs; (4) to assess forecast accuracy in addition to waste, availability and inventory cost; and (5) to suggest practical governance and deployment practices for small retailers.

1.2 Research Questions

The research questions are whether a hybrid AI model can enhance the accuracy of forecasts compared to basic models in retail environments, where the products are fragmented; how information sharing can support shorter sales histories; how a stockout and product expiry should be represented; and what metrics better reflect the operational value of demand forecasting for perishables products in a retail environment.

2. Literature Review

Classical time-series forecasting provides the foundation for retail demand planning. Moving averages remove the short-term noise from data, exponential smoothing assigns more importance to the most recent data points, and seasonal methods can be used to account for other recurring patterns like weekly or yearly cycles of demand. Box–Jenkins modeling of ARMA models provides a systematic way of modeling temporal dependence (Box et al., 2015). These are good baselines to use as they are fairly easy to interpret and can be effective when data is limited. However, they might need more feature engineering to capture

[1] Department of Artificial Intelligence, University of Agriculture, Faisalabad, Pakistan
[2] Department of Computer Science, Government College University Faisalabad, Faisalabad, Pakistan

Copyright (c) 2020 Author(s)
This work is licensed under a Creative Commons Attribution 4.0 International License.

promotions, weather conditions, stock availability and the interactions between them.

Retail demand may be erratic, particularly in the case of slow-moving goods or small retailers. Croston's method decouples the nonzero demand size from the gap between demands, and is applicable to series with intermittent demands (Croston, 1972). Perishable products are different because they may be required on a regular basis and because their demand is very sensitive to changes but as the shelf life decreases the cost of error can increase. Hence, the forecasting method should be chosen depending on the product type and the demand characteristics of the product, and not depending on any algorithm.

Model features, such as machine-learning based random forests and gradient-boosted trees, can incorporate lagged sales, calendar features, price, promotions, weather, store attributes, and more. They can learn nonlinear relationships without having to provide a complete parametric time-series model. They have a few drawbacks such as sensitivity to leakage of features, problems in extrapolating trends and a need for careful validation. When there are sufficient amount of data, shared representations can be learned across multiple time series with neural forecasting techniques. For instance, DeepAR makes probabilistic forecasts for related time series, and it can also be used to generate prediction distributions in addition to just point values (Salinas et al., 2020).

The more you can add up the statistical strength of multiple related series in a global model, the better that model can be at predicting for a data-poor location. Hierarchical forecasting models integrate forecasts at various levels (store, neighborhood, category, and SKU), which helps to maintain consistency between local and aggregate forecasts. Stores in fragmented markets, however, may have different ranges, opening hours, and customer base, and have different purchasing practices. Pooling should be selective, however: information across similar stores/products can be shared, and patterns local to stores can be represented using store-specific features or adjustments.

One difficulty is that sales do not necessarily equal actual demand. If the inventory is depleted, then it will be possible to purchase something that is either lost or

[1] Department of Artificial Intelligence, University of Agriculture, Faisalabad, Pakistan
[2] Department of Computer Science, Government College University Faisalabad, Faisalabad, Pakistan

Copyright (c) 2020 Author(s)
This work is licensed under a Creative Commons Attribution 4.0 International License.

replaced, which will lead to sales figures being a false reflection of demand. Stockouts might not be true low demands, but rather observations that were not recorded. Without using the information on the number of stocks, a model trained directly from these records could predict the low demand for products that have sold out over and over again. The focus of retail forecasting and inventory management research thus lies in the relationship between demand forecasting, replenishment policies and inventory availability.

There are many error measures that can be used to evaluate a forecast. Let's first discuss the easy-to-understand mean absolute error, and then the more heavily penalizing root mean squared error. The MAPE may not be stable when the actual demand is zero or very small. Scaled measures like mean absolute scaled error allow comparisons between different products, with different units and scales (Hyndman & Koehler, 2006). There are several methods for evaluating probabilistic forecasts such as quantile loss or weighted interval score, and there are various methods for evaluating policies for inventories, such as service level, waste, stock out frequency, and total cost.

Problem formulation and market context.

It was determined that the demand for the product is equal to the sales and the availability of it.

Suppose that y(s,i,t) represents sales observed for a store s, item i and time period t. When the amount of goods on hand is small, actual sales could fall short of the underlying customer demand d(s,i,t). Minimalization of relationship is a simplified one: y = min(d, a), where a is the quantity available for sale during the period and no substitution, return or intraperiod replenishment is considered. Availability may also vary during the day in practice; if possible, the data pipeline should capture the time the product is out of stock or at least a flag that the product was in stock on a daily basis. The demand of the sales series alone can be underestimated when the stockouts are frequent.

3.3 Excessive Exports

Remaining shelf life must be taken into consideration when making the inventory decision. Too much and too little can result in expirable, marked down or

[1] Department of Artificial Intelligence, University of Agriculture, Faisalabad, Pakistan

[2] Department of Computer Science, Government College University Faisalabad, Faisalabad, Pakistan

Copyright (c) 2020 Author(s)

This work is licensed under a Creative Commons Attribution 4.0 International License.

disposed products and lost sales, respectively, and thus affect customer trust. The cost of overforecasting is not necessarily equal to the cost of underforecasting. If the product has a short life, an overforecast will lead to waste, and an underforecast can be very expensive at the time of maximum demand, when it takes a long time to place an order. The objective of the service-level should be taken into account in the preferred forecast quantile, along with these asymmetric costs.

- Fragmentation and data heterogeneity.

Frequently retail locations do not have standard product identifiers, and they do not have synchronized transaction systems. Local names, various package sizes, or varying units may be used to describe similar products. Daily totals may be provided by some stores or transaction data by others. The proposed framework adopts a canonical product dictionary, unit normalisation and data quality flags as well as flexible aggregation to develop comparable features. If the missingness is informative, meaning the reason for the missingness is known, then missing data should be treated as such and not as zero sales, for example.

4. Suggested AI-Based Forecasting Approach

This framework has seven components: (1) Data Ingestion, (2) Data Cleaning and Entity Resolution, (3) Feature Engineering, (4) Baseline Forecasting, (5) AI Forecasting and Information Sharing, (6) Probabilistic Calibration, and (7) Inventory Decision Support. The modules can be deployed centrally or via a light-weight cloud service, and have easy-to-use interfaces for stores lacking in complex local infrastructure.

4.1 Data Sources

Historical sales data, product identifier, store identifier, prices, promotions, inventory, replenishment, supplier lead time, expiry and shelf-life data are core data. The temperature, rainfall, public holidays, local events and wholesale market prices can be optionally added as external data. The features should be accompanied by a time stamp when it was made available. For instance, weather forecasts are forecasts, and not the actual weather that was experienced after

[1] Department of Artificial Intelligence, University of Agriculture, Faisalabad, Pakistan
[2] Department of Computer Science, Government College University Faisalabad, Faisalabad, Pakistan

Copyright (c) 2020 Author(s)
This work is licensed under a Creative Commons Attribution 4.0 International License.

the order was placed. This means that future information does not seep into historical evaluation.

4.2 Data Cleaning and Feature Engineering.4.2 Data Cleaning and Feature Engineering.

The pipeline checks for negative values, duplicate transactions, inconsistent units, missing dates, too large price fluctuations and changes in products. It normalizes all the units (kg, litre, piece, no. of packages), and it retains the conversion rules. Lagged sales, rolling means, rolling variability, day of week, month, and proximity to holidays, price, promotion status, time since last delivery, store size, product category, and remaining shelf life are all features. Only information from the forecast origin can be used for the computation of lags and rolling features.

If the absence of a value makes a statement, then it should be accompanied by indicators. Due to the fact that the extreme demand may be a true 'promotion', 'bulk purchase' or ‘true festival', it is important that outliers are investigated before they are removed. Product mapping may be based on a mix of supplier codes, text normalisation, packaging characteristics, and manual checking. A rather semi-automated catalog-matching workflow might be more achievable for small retailers than a perfect data set prior to deployment.

4.3 Baseline Models

First, forecasts are prepared using the seasonal naive, moving-average, and exponential-smoothing methods. A seasonal naive forecast is based on the latest observation at the same time of the season (e.g. the same weekday in the prior week). These baselines are valuable, because they are cheap and give a reference which a complex model has to outperform. Croston-type methods can be incorporated for intermittent demand. For different product categories, and data densities, model choice may differ.

4.4 Machine-Learning Forecasting Model

The key machine-learning contender is a store-item-time example-trained gradient-boosted decision tree. Historical lags, rolling statistics, price, promotions, calendar, weather, inventory, supplier lead time and store and product attributes

[1] Department of Artificial Intelligence, University of Agriculture, Faisalabad, Pakistan
[2] Department of Computer Science, Government College University Faisalabad, Faisalabad, Pakistan

Copyright (c) 2020 Author(s)
This work is licensed under a Creative Commons Attribution 4.0 International License.

Journal of Technological Information, Management & Engineering Sciences (2020),1(2), 53-68

are all features. What is the target, it could be demand of the next day, some days ahead, or the period of ordering practice (lead-time horizon). Across-stores training enables the model to capture common effects while incorporating the store and product information in the form of categorical or encoded information to capture local differences.

We propose a generalized formulation of $\hat{y}(s,i,t+h) = f\theta(X(s,i,t), E(t), C(s,i))$, where X features historical and operational features, E features external variables, C features store and item context, and $f\theta$ is the learned model. The multi-horizon model can generate a number of forecasts for the future. The model should be trained using a loss relevant to the operational objective, for example absolute error for producing robust point forecasts or quantile losses for inventory decisions.

4.5 Cross-Store Information Sharing

The importance of information sharing when there is a short history at individual stores. The framework learns a global store and product model, and applies the store-specific features for prediction modification. Geography, store format, customer traffic, product assortment, and historical demand patterns, are examples of ways to create similarity groups. Neighborhood or category-level signals can be constructed using hierarchical aggregation for new stores and sparse items. Indiscriminate pooling should be avoided in the system as a large urban store is not necessarily a good analogue for a small rural store.

4.6 Probabilistic Forecasting

A point forecast does not give the information about uncertainty. The model proposed is thus an estimation of several quantiles or a predictive distribution. Models can be fit for levels like the median and higher quantiles of demand, and the calibration of these models can be tested using a held-out validation period. It may be suitable to use the upper quantile for products where stockout costs are high, and the lower quantile when expiry costs are high. The selection of quantities should be based on the economics and operating conditions of the product.

4.7 Inventory Decision Layer

[1] Department of Artificial Intelligence, University of Agriculture, Faisalabad, Pakistan

[2] Department of Computer Science, Government College University Faisalabad, Faisalabad, Pakistan

Copyright (c) 2020 Author(s)

This work is licensed under a Creative Commons Attribution 4.0 International License.

Journal of Technological Information, Management & Engineering Sciences (2020),1(2), 53-68

Forecasts are used to drive an inventory policy based on available stock, orders placed, expected deliveries, lead time, shelf life and target service level. An "order quantity" is a simplified version of the formula that combines forecasted demand for the lead time (D̂L), safety stock (SS) based on the uncertainty of demand and the service level requirements, and the inventory position (IP) including on-hand inventory and confirmed orders. The policy for perishable items should also include age-bucket inventory and expiry dates, requiring that all items are not considered equally fresh.

A practical system can suggest a quantity and present a brief description, e.g. new sales trend, next holiday, weather forecast, expected delivery delay. It is recommended that the user be able to override recommendations and be able to record the reason. While overrides do give valuable information, they are not necessarily correct labels, as they may be due to local knowledge, supply constrains or human error. The system should assess if the recommendations lead to a reduction in waste levels, total costs and an acceptable service level.

5. Research Methodology

This paper uses a quantitative research design for the evaluation of the proposed model. An original retailer dataset was not provided, and therefore, measured improvements in forecast are not claimed. A future empirical implementation could rely on a public and suitable public dataset, or a participating retailer's data that is documented and licensed. Research should be preregistered, or otherwise committed to hypotheses, data exclusions, validation scheme, and primary metrics prior to final testing.

5.1 Dataset and Sampling

As many stores as possible, as many product categories as possible, as many sale dates as possible, as many prices as possible, as many promotions as possible, as many stock availability records as possible, as many replenishment records as possible, as many shelf-life attributes as possible. The sampling time should encompass a sufficient number of seasonal cycles to provide an estimate of the recurring effects and some typical and atypical demand periods. Stores should be

[1] Department of Artificial Intelligence, University of Agriculture, Faisalabad, Pakistan
[2] Department of Computer Science, Government College University Faisalabad, Faisalabad, Pakistan

Copyright (c) 2020 Author(s)
This work is licensed under a Creative Commons Attribution 4.0 International License.

Journal of Technological Information, Management & Engineering Sciences (2020),1(2), 53-68

grouped according to size, geography, format and data quality to determine if benefits extend to the largest or more digitally advanced stores.

5.2 Temporal Validation

Randomly dividing the data for the day may introduce patterns from future observations into the training set. Rather, it adopts the method of rolling origin evaluation: train in a small window of the historical data, predict the next small window, shift the origin of the predictions forward, and repeat. All features should be recreated with information from the origin. Hyperparameters are optimized on validation windows and the final test window is never altered until the final model selection is done. If a model is evaluated at multiple horizons, each horizon should be separately reported.

5.3 Handling Stockouts and Missing Data

The study should identify periods of stock-out and differentiate between zero sales during periods of stock-outs and zero sales during periods of unavailable stock. When stock availability information is available, the demand during stock-outs can be reported as censored or estimated using appropriate methods, and uncertainty reported as well. Sensitivity analysis should also involve models trained on all sales data, and models that remove and/or add known stockout periods. If inventory records are incomplete, the limitation should be expressed in terms of a number, instead of counting on sales being an unbiased representation of unconstrained demand.

5.4 Evaluation Metrics

The metrics used to forecast are MAE, RMSE, MASE, and quantile loss. While MAPE is a measure of the error, it should not be used as an indicator of the error alone because it is undefined or unstable when the actual values are zero and small. Another metric of forecast precision is forecast bias, which assesses if the model over- or under-predicts demand. Evaluating probabilistic forecasts for quantile coverage and interval width. Operational Measures are the ones that include: Waste Rate, Stockout Rate, Fill rate, Service level, Gross Margin, Markdowns and Total Inventory Cost.

[1] Department of Artificial Intelligence, University of Agriculture, Faisalabad, Pakistan

[2] Department of Computer Science, Government College University Faisalabad, Faisalabad, Pakistan

Copyright (c) 2020 Author(s)

This work is licensed under a Creative Commons Attribution 4.0 International License.

Journal of Technological Information, Management & Engineering Sciences (2020),1(2), 53-68

Simple waste rates are the number of units that are expired or discarded divided by the number of units received or available for the period; the denominator must be consistently defined. Fill rate is a percentage of the number of units filled from inventory. The above measures should be considered by category as the shelf life, packaging and replenishment frequency vary. Bigger gains in MAE can be obtained even without gains in waste when the inventory policy is not responding appropriately, and vice versa.

5.5 Baseline Comparisons and Ablation.

Proposed model should be compared with a seasonal naive, an exponential smoothing model, and if applicable a Croston-type model, all without external variables, if any, and a machine-learning model with no external variables. Ablation experiments isolate one of the aforementioned factors at a time. These comparisons indicate if added complexity provides extra signal. The study should include runtime and data requirements and accuracy, and in particular, small retailers might not be able to maintain a costly data pipeline.

6. Proposed Algorithmic Workflow

Step 1: Set Forecast Time Horizon, Demand Measure, Product Scope and Inventory goals.

Step 2: Gather sales, inventory, price, promotion, supplier, "shelf-life" and then optional external data, all with a time stamp.

Step 3: Normalize product identifiers and units; check for missingness, extremes, and duplicates.

Step 4: Identify stock out and unavailable time periods to separate actual sales from potential demand.

5. Step 5: Generate lagged, rolling, calendar, price, promotion, weather, store, and product features without temporal leakage.

Step 6: Establish baseline seasonal and statistical for each appropriate product group.

Train a global machine-learning model with local context features, across stores and products, Step 7.

[1] Department of Artificial Intelligence, University of Agriculture, Faisalabad, Pakistan

[2] Department of Computer Science, Government College University Faisalabad, Faisalabad, Pakistan

Copyright (c) 2020 Author(s)

This work is licensed under a Creative Commons Attribution 4.0 International License.

Journal of Technological Information, Management & Engineering Sciences (2020),1(2), 53-68

Yousaf and Amin

Step 8: Create quantile forecasts or predictive intervals and valid them on validation windows.

Step 9: Translate forecast into replenishment quantities based on inventory position, lead time, shelf life and service levels.

Step 10: Test and measure forecasts and inventory results over rolling origin periods.

Step 11: Run subgroup analysis and ablation analysis and quantify uncertainty of difference in performance.

Step 12: Roll out slowly, gather user overrides, track drift, and retrain within a documented environment.

Expected results and analysis plan will be included.Expected results and analysis plan will be included.

The core hypothesis is that such a global ML model that will leverage store context and product context will achieve better performance across many high-volume and promotion-sensitive products when compared to simple baselines. A second hypothesis is that cross-store information sharing will have greater impact for stores and products that have fewer history. A third possibility is that the joint waste and stockout cost will be lower with probabilistic forecasts and shelf-life-aware ordering than with fixed quantities ordering. These hypotheses are not yet experimental results but should be described as hypotheses until the experiment has been done.

Results should be presented by product category, store group, forecast horizon and data availability. A model could work well on dairy and not so well on fresh produce, or do better on high volume stores than on low volume rural stores. Confidence intervals or repeated rolling-origin results should be included in the report and the proportion of store-item series for which the AI model outperforms the baseline should be identified. Average figures can mask these differences.

Modelling Strengths Key limitations

Standard Easy to execute, audit and cheap Just as effective as standard, but limited response to external drivers.

[1] Department of Artificial Intelligence, University of Agriculture, Faisalabad, Pakistan

[2] Department of Computer Science, Government College University Faisalabad, Faisalabad, Pakistan

Copyright (c) 2020 Author(s)

This work is licensed under a Creative Commons Attribution 4.0 International License.

Trend and seasonal models Recurring seasonality, trend, need stable structure, careful specification Exponential smoothing

Gradient-boosted trees Models nonlinear drivers and interactions Needs feature quality and temporal validation

Global AI model with local context Shares information across sparse stores and items May underfit highly local patterns

When uncertainty is linked with ordering and service targets - inventory, lead-time, and shelf-life data are crucial and need to be reliable.

Business Impact & Cost Analysis 7.1

Purchase, stock, expire, markdown, stockout and emergency replenishment should be combined into a business analysis. Cost assumptions need to be provided and sensitivity to these needs to be run. For instance, a store that makes more deliveries can have less safety stock than a store that has longer supplier lead times. The model should not be deemed successful merely because it minimizes the mean error, if it also results in high percentage of waste or stockout in crucial categories.

7.2 Robustness and Generalization

The robustness test should include holidays, extreme weather, supplier delays, price shocks and times of unusually high demand. When possible, models should be evaluated on stores not used during model training, to test transfer to new stores. New products should be evaluated individually with metadata and category level priors. Changes in sales distributions, prices, sales promotion use, forecast errors and inventory outcomes should be monitored.

8. Discussion

Proposed method involves predictive flexibility and operational constraints. Machine Learning can learn the relationship between the demand and external drivers, and global training can share the information across the non-uniform stores. A probabilistic forecast is a statement of uncertainty, and can be used to offer different levels of service. But quality of data and integration of processes is the basis for the value of the system. A refined prediction is not going to be of any

[1] Department of Artificial Intelligence, University of Agriculture, Faisalabad, Pakistan

[2] Department of Computer Science, Government College University Faisalabad, Faisalabad, Pakistan

Copyright (c) 2020 Author(s)

This work is licensed under a Creative Commons Attribution 4.0 International License.

Journal of Technological Information, Management & Engineering Sciences (2020),1(2), 53-68

use if inventory information is wrong, lead-time from suppliers is not known, or it is not possible for employees to follow up on recommendations.

Hence, it is recommended to implement a staged deployment. The first phase should be to standardize product and sales data and come up with the base forecasts. The second phase can add as global forecasts of machine learning for categories for which there is enough information. Additional probabilistic predictions and inventory decisions based on shelf-life can be incorporated during the third phase. This phased implementation will enable retailers to prove value before introducing further technical complexity. It also helps minimise the chance of confusing model errors with operational failures and vice versa.

In fragmented markets, human expertise will still be necessary. There may be other community events, holidays, road closures, supplier shortages, or other events that managers at the store are aware of that aren't included in the data. A good interface should display forecasts and their uncertainty and a few simple, understandable drivers and give documented overrides. Feedback needs to be systematically reviewed to identify missing data, operating practice limitations and inconsistencies in operating practices.

Ethical, Privacy and Governance concerns 8.1

Retail transaction information could include commercially sensitive information regarding supplier terms, local demand and pricing. Permission, purpose limitation, access control, and authentication safeguards to prevent others from seeing one retailer's confidential business performance are needed for data sharing between independent stores. It is helpful if the aggregated learning is useful, but aggregation does not imply privacy. Written contracts and governing law should dictate data retention, security and cross-store modeling.

The system should record training data, definition of features, versions of models, forecast periods, limitations, and business logic employed for order generation. Human overrides need to be auditable and in case of external data or model serving failure, the model must be able to fallback to a simple baseline. Forecast bias, error by category, waste, stockouts and staff adoption should be monitored.

[1] Department of Artificial Intelligence, University of Agriculture, Faisalabad, Pakistan

[2] Department of Computer Science, Government College University Faisalabad, Faisalabad, Pakistan

Copyright (c) 2020 Author(s)

This work is licensed under a Creative Commons Attribution 4.0 International License.

Journal of Technological Information, Management & Engineering Sciences (2020),1(2), 53-68

Retraining a model should be done only if/after data quality checks and validation, not because an alert has been triggered.

8.2 Limitations

This paper is not an original empirical study. The type of expected model performance is dependent on the availability and reliability of transaction, inventory, expiry and outside data. Many markets are broken and stockout records and product identifiers may be incomplete. Ordering rules, staff behaviors, supplier constraints and consumer substitution are interposed between forecast accuracy and waste. Availability of weather and event information and the reliability of historical patterns can be affected by events and/or weather conditions in certain locations, and by sudden economic or supply shifts.

9. Conclusion

In this paper, the authors suggested an AI-based demand forecasting model for perishable products in an uncentralized retail system. It incorporates data standardization, statistical baselines, global machine learning models, cross-store information sharing, probabilistic forecasting and shelf-life-aware inventory decisions. It explicitly considers stockout-censored sales, local heterogeneity, sparse histories and the asymmetric costs of waste and lost sales.

A key recommendation is to assess forecasting within a decision system for inventory. Model validation should be undertaken using rolling origin principles, appropriate baseline conditions, statistical and operational criteria. The potential for cross-store learning without loss of local differences and commercial confidentiality is an opportunity for data-poor retailers. Further empirical studies should be done using real store data to determine how the framework can lead to improvements in waste and availability, as well as identify the components of the data and model that are most valuable to retailers with less digital infrastructure.

References

Box, G. E. P., Jenkins, G. M., Reinsel, G. C., & Ljung, G. M. (2015). Time Series Analysis: Forecasting and Control (5th Edition). Wiley.

[1] Department of Artificial Intelligence, University of Agriculture, Faisalabad, Pakistan

[2] Department of Computer Science, Government College University Faisalabad, Faisalabad, Pakistan

Copyright (c) 2020 Author(s)

This work is licensed under a Creative Commons Attribution 4.0 International License.

Journal of Technological Information, Management & Engineering Sciences (2020),1(2), 53-68

Croston, J. D. (1972). Forecasting & stock control for intermittent demands. Operational Research Quarterly, 23(3), 289–303. https://doi.org/10.1057/jors.1972.50

Hyndman, R. J., & Koehler, A. B. (2006). Another discussion of measures of forecast accuracy. International Journal of Forecasting, 22(4), 679–688. https://doi.org/10.1016/j.ijforecast.2006.03.001

Makridakis, S., & Spiliotis, E. & Assimakopoulos, V. (2020). 100,000 time series and 61 forecasting methods in the M4 Competition. International Journal of Forecasting, 36(1), 54–74. https://doi.org/10.1016/j.ijforecast.2019.04.014

D. Salinas, V. Flunkert, J. Gasthaus, and T. Januschowski. (2020). Probabilistic Forecasting with AutoRegressive Recurrent Networks (DeepAR). International Journal of Forecasting, 36(3), 1181–1191. https://doi.org/10.1016/j.ijforecast.2019.07.001

Syntetos, A. A., & Boylan, J. E. (2005). The precision of the estimates of intermittent demand. International Journal of Forecasting, 21(2), 303–314. https://doi.org/10.1016/j.ijforecast.2004.10.001

Taylor, J. W., & Letham, B. (2018). Forecasting at scale. The American Statistician, 72(1), 37–45. https://doi.org/10.1080/00031305.2017.1380080

Waller, M. A., & Fawcett, S. E. (2013). The revolution of data science, predictive analytics, and big data will impact the design and management of the supply chain. Journal of Business Logistics, 34(2), 77–84. https://doi.org/10.1111/jbl.12010

[1] Department of Artificial Intelligence, University of Agriculture, Faisalabad, Pakistan

[2] Department of Computer Science, Government College University Faisalabad, Faisalabad, Pakistan

Copyright (c) 2020 Author(s)

This work is licensed under a Creative Commons Attribution 4.0 International License.