

Online Fashion Commerce: Modelling Customer Promise Date*

Preethi V[†] Nachiappan Sundaram[‡] Ravindra Babu Tallamraju[§]

Abstract

In the e-commerce space, accurate prediction of delivery dates plays a major role in customer experience as well as in optimizing the supply chain operations. Predicting a date later than the actual delivery date might sometimes result in the customer not placing the order (lost sales) while promising a date earlier than the actual delivery date would lead to a bad customer experience and consequent customer churn. In this paper, we present a machine learning-based approach for penalizing incorrect predictions differently using non-conventional loss functions, while working under various uncertainties involved in making successful deliveries such as traffic disruptions, weather conditions, supply chain, and logistics. We examine statistical, deep learning, and conventional machine learning approaches, and we propose an approach that outperformed the pre-existing rule-based models. The proposed model is deployed internally for Fashion e-Commerce and is operational.

Keywords— LightGBM, asymmetric loss, quantile regression, breach control model, fbprophet, lstm

1 Introduction.

The outbound supply chain consists of various stages such as procurement from the vendor, warehouse processing, linehaul shipping, and delivery from the last-mile center. The supply chain and logistics-related findings from PwC[2] show that 39% of the customers want their goods delivered within 2 days, while 50% consider 3-5 days as acceptable. However, it is important to have better predictability over speed[2]. Predicting delivery dates accurately at the time of order placement is one of the biggest challenges in the e-commerce space, given the uncertainties involved in each of the supply-chain legs. Once an order is placed, a promise date for delivery is given to a customer based on the expected time taken at every leg. The uncertainties at various stages include traffic, weather conditions, curfew, etc. The additional factors that can be modeled include delays due to holidays, weekends, pendency from an increased load, office location deliveries on weekends. A shipment delivered

after a promised date leads to a breach resulting in a bad customer experience.

In a large developing country like India, building a strong supply chain network capable enough to reach all remote parts of the country on time is a challenge due to limited infrastructure availability in certain parts of the country, unexpected traffic, road blockages, weather conditions, and variable holiday patterns. PwC study showed that retailers are moving from large distribution centers in the heartland to smaller ones close to population centers[2]. Myntra has also strategically located the last-mile centers based on demand concentration, thereby ensuring that the majority of delivery requirements are met by myntra logistics (ML). However, support from third-party logistics (3pl) is imperative to reach the remote parts of the country and places where the demand is less. This adds to the uncertainties, as we do not have clear visibility into the 3pl partners' network in terms of transit hub¹ and last-mile center locations, manpower, and their load from other vendors. There are also challenges arising from limited data availability for low-demand regions and operational variability during fixed and flexible holidays, and weekends at the pick-up, transit, and last-mile locations.

In the fashion industry, since it is necessary to have a variety of products to meet a varied set of customers with differing needs, the individual stock-keeping-unit (SKU) depth (stock available for each item) tends to be low. The demand needs to be met by a small number of warehouses where goods are stored as well as processed and also by fulfillment centers to cover regional demand. Strategically distributing the SKUs across them is challenging. Thus, there might be cases where the fulfillment cannot be done from the nearest fulfillment center or warehouse. This leads to a longer linehaul and subsequently more uncertainties resulting from traffic congestion, hub wait times, and regional holidays in transit hubs.

Promised delivery times are also affected by hard-to-decipher customer addresses, as the addresses may not follow a predefined structure. Incomplete addresses, spell variations, and avoidable additional text such as

*Supported by GSF grants ABC123, DEF456, and GHI789.

[†]Myntra Designs, India.

[‡]Ericsson, India.

[§]Myntra Designs, India.

¹A hub is the distribution center where shipments are collected, sorted, and reloaded.

directions to reach the address, times of their availability pose challenges in automated address clustering models[1] designed to ease delivery. Thus quick and successful delivery depends on address validity.

Predicting delivery times during high revenue days (HRD) is another big challenge, as the variability in delivery times would be relatively higher and the shipping cutoffs at the transport and dispatch hubs may not be adhered to. The prediction has to be made based on past HRD events, which typically happen only a couple of times in a year. There would be changes happening in the supply chain since past events' occurrence, for instance, opening, merging, and closing of delivery centers, dispatch hubs, process changes, change in revenue targets, etc. Thus, planning data such as volume and capacity plans at various supply chain touchpoints play a major role in predicting delivery time accurately in case of HRD. Delivery time accuracy is thus also influenced by the accuracy of plan numbers.

In this paper, we briefly describe how accurately predicted times could help in examination of unexpected delays and processing rate changes thus enabling optimisation of the supply chain operations. We compare time series, deep learning and machine learning-based approaches for the current problem. We also examine asymmetric loss functions to meet the business constraint of keeping breaches under control as the cost of a breach is higher than the cost of promising a date much later than the actual delivery date. We also demonstrate the effectiveness of this solution for a large-scale Indian e-commerce company.

In the next section, we briefly describe research work related to the current problem. Section 3 contains problem description and challenges. Section 4 discusses various methodologies applied to solve the problem and concludes on the proposed solution. We present the results of various modeling approaches and real-world deployments in Section 5 and, conclusions and further work in Section 6.

2 Related Work.

In this section, we briefly review the existing work on delivery time prediction. Ravindra Babu, et al.[1] describe a classification approach to assign addresses to area-codes, based on the address-text. They describe the impact of unstructured, incorrect, or incomplete addresses on the accuracy of classification. Such addresses can impact the on-time delivery of products, especially in India, where addresses do not follow a regular format like in the Western-world. Jeff Berman[2], discusses customer expectations. Customers have become accustomed to getting products as and when they want them, which creates challenges for

the supply chain. Faster deliveries need a high degree of alignment at the supplier, storage, transport, and manpower levels, which can be expensive. Predictions from models, whether rule-based or data-driven, can be erroneous. The predicted delivery date can fall before or after the actual delivery date. The former error is more detrimental for the company (late deliveries) and should be penalized more as compared to the latter error. Models that we build should capture this asymmetry[3].

In [4], authors talk about predicting lead times for made-to-order industrial-grade equipment. Here, numerous categorical variables affecting times were reduced using the PCA technique. The random forest model gave the best accuracy. In [5], shipping time between factory to port in two continents is predicted using the Random forest model by taking into account port congestion and a major holiday. In this study, it is shown that as the shipment gets closer to the destination, MAE reduces and random forest performs similar to traditional approaches. In almost all the problems, the value lies in the ability to make correct predictions right at the origin time, as customers would not prefer to see updated times at every stage. In [8], the authors predict purchase order fulfillment times using random forest and quantile regression forests. In this paper, dimension reduction technique has been applied to identify the most influential levels of categorical variables similar to [4]. This paper explores tree-based models, time series and deep-learning based approaches to model delivery times end-to-end, while also demonstrating the use of custom loss functions and feedback-based breach control model to handle asymmetry.

3 Problem Description and Challenges.

The objective of the current problem is to accurately predict delivery dates at the time of order placement. Once an order is placed, the item has to traverse through various touchpoints in the supply chain before reaching the customer's location.

The order processing starts from getting the goods from the vendors in case items are not available in the warehouse. The order fulfillment time from the vendor would vary based on the item availability, vendor processing speed, transportation time and related delays, use of co-loader vehicles, and weekend delays due to reduced manpower. It would also depend on the designated times at which the purchase order could be issued and goods can be picked up in a day. Further processing is done at the warehouse.

The order processing starts at the warehouse for items available in the warehouse and there is no vendor processing leg involved for such items. It broadly involves the activities of picking, consolidation, and

packing. An uncertain delay is introduced because of wait times for multi-item packets, cascading effect resulting from occasional choking in one or more previous stages, dependencies due to warehouse closure, delays due to working-shift changes, variable manpower during weekends and flexible holidays, etc.

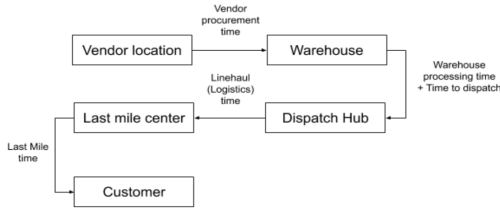


Figure 1: Order processing stages

Subsequently, a shipment traverses through a sequence of hubs where transit time is influenced by traffic congestion, weather conditions, hub wait times until shipping cutoff, regional holidays in transit hubs, etc.

When the packet reaches the last-mile hub, one of the factors causing delay is variable manpower on holidays and weekends. Each holiday duration depends on the nature of the holiday such as religious festivals. The delay in packet delivery is a function of current dependency, available manpower, customer availability, and deliverability of an address. However, it is difficult to localize which packets would be delivered late as it depends on customer availability and delivery re-attempt. The order processing stages are summarized in Figure 1.

4 Proposed Solution.

We propose to solve this problem in three stages. The following is the set of models built for predicting the processing time involved in various supply-chain legs:

Vendor procurement model: The time taken from order placement to shipping is predicted using this model. This is applicable only for items not stored in the warehouse. This time is modeled using LightGBM[6] and explained in Section 4.8.

Warehouse model: This model predicts the time taken from order placement to shipping only for items stored in the warehouse. This time is modeled using time series (fbprophet[9]) and LightGBM as described in Section 4.6.2.

Shipping time model: This model predicts the time taken to reach customer location after shipping referred to as the shipping time. This constitutes the linehaul and last-mile times as indicated in Figure 1. The shipping time for ML during Business As Usual(BAU) days is predicted using tree-based models, time-series, and deep learning-based approaches. Tree-based models, namely,

LightGBM, Random forest, and XGBoost are described in Section 4.2. The time series modeling approach using fbprophet is described in Section 4.6.1. The Deep-learning based approach (LSTM) is explained in Section 4.7. The shipping time for HRD is modeled using LightGBM and explained in Section 4.3. Shipping time for 3pl partners is modeled using LightGBM and explained in Section 4.4.

The various methods used to control breaches such as asymmetric custom loss function, quantile regression, and feedback-based breach control model are discussed in Section 4.5.

At the time of order placement, we propose to give the delivery date to the customer by adding the outputs from the Warehouse model and Shipping time model to the order placement time for items stored in the warehouse. In case of items not stored in the warehouse, the outputs from the Vendor procurement model and Shipping time model will be added to the order placement time.

In this section, we initially present the rule-based model which was operational before the proposed model. Subsequently, we discuss tree-based models for predicting the shipping time of ML and 3pl. The modeling approach for HRD is also described. We then describe ways to control breaches. We discuss time series models to predict shipping time and warehouse time. Subsequently, we present the LSTM model to predict shipping time and the tree-based model to predict vendor procurement time. Then, we briefly discuss applications in supply-chain optimization.

4.1 Rule-based model to predict delivery time.

In the past, a rule-based model was used to predict delivery times. At the level of vendor-warehouse, vendor processing times are configured. Similarly, at a warehouse level, processing times are configured for each of the warehouse legs. In the case of linehaul, the hop-to-hop times are configured and static last-mile center processing time is added. Each of these times is mapped to the nearest shipping cutoffs wherever appropriate. Fixed additional processing times are added for weekends and holidays. Thus, rule-based model doesn't adapt based on recent performance changes and is designed based on heuristics.

4.2 Tree-based models to predict shipping time for Business As Usual(BAU) days.

4.2.1 Model choice. We examined LightGBM[6]. It has the ability to represent categorical features with a large set of levels. The current problem involves categorical features with up to hundreds of levels, which

blow up the feature space when one-hot encoded. The selection of optimal split points in tree-based models would thus consume a large amount of time and we may not be able to meet the latency requirements. LightGBM uses exclusive feature bundling to optimally bundle mutually-exclusive features. This works especially when the feature space is sparse as in the case of one-hot encoded features.

Time consumed for the selection of optimal split points is also directly proportional to the number of training examples. Given the huge number of orders placed in the last few weeks before prediction, it is necessary to reduce the number of training examples without hurting accuracy. Gradient-based one-side sampling[6] achieves this by randomly-sampling instances with smaller gradients while not changing the data distribution (using weights). Thus, the instances with larger gradients are majorly used to determine the split points. Random forest and XGBoost models were also built on the dataset.

4.2.2 Model features. Following are the broad list of features considered to take into account various factors affecting shipping time:

Historical times and pendencies: The pendency in the linehaul and last-mile plays a major role in determining accurate times. The performance would vary for each logistic lane and last-mile center, which is modeled using statistical measures. Pendency is modeled using the recent inflow and outflow of shipments.

Geographical features: The geographical location level features of customer addresses such as city tier, address type (home or office), area to name a few play an important role in determining delivery time. This helps the model to learn the patterns in time by aggregating geographical locations and thus generate near-accurate predictions for locations with limited training data.

Load features: The expected load at various touchpoints plays a key role in determining delays. The on-ground order and ship volumes from the past days play a major role in determining pendencies at all further stages in the supply chain process.

Manpower features: Various manpower types available at the last-mile center are modeled individually to capture productivity and process differences.

Last-mile performance features: Arrival, delivery volumes and performance at the last-mile center on the past days are modeled taking into account cumulative pendency from past days.

Holiday type and seasonality features: The transit hub and last-mile center holidays are captured using these features. Sudden change in the manpower, arrival,

and delivery volumes at the last-mile center might occur due to full or partial closures on holidays/weekends or varied outflow on some weekdays. Hence, the holiday and weekday types are included as features so that the model can effectively learn the relationships between variables.

4.2.3 Model tuning. The examples are weighted using a decaying weighting function as the most recent supply chain operations have a greater impact on current delivery times. As the LightGBM model tends to overfit, we tuned various parameters as indicated in Table 1. The validation set has to be carefully chosen in a time series context as it should be similar to the test set in terms of order day and also meet the recency criteria in this case.

Table 1: Parameter tuning for LightGBM

Metric	Final Value	Default Value	Parameter value list
Boosting iterations	1000	-	[1000, 1500, 2000, 2500]
Tree depth	-1	-1	[4, 5, 8, -1]
Number of leaves	15	31	[15, 31, 63, 127]
Data fraction	0.6	1.0	[0.6, 0.7, 0.8, 1.0]
Feature fraction	0.6	1.0	[0.6, 0.7, 0.8, 1.0]

4.2.4 Holiday and weekend handling times. Weekends and holidays play a major role in determining delivery times. Every last-mile center behaves differently on different holidays based on the holiday type(fixed or flexible), the intensity of the holiday in that region, and the adjacency of the holiday to a weekend. Weekend behavior also differs for every last-mile center. The effect of holiday or weekend is modeled by selecting the most recent similar holidays(based on absenteeism rates) or weekends. The proxy holidays/weekends are then compared with BAU days that lie close to them. The chosen BAU days are subjected to outlier treatment. Statistical measures that are robust against outliers are used to compare and subsequently derive holiday and weekend handling times.

4.3 Tree-based model to predict shipping time for HRD. Delivery times of orders placed on both HRD and post-HRD days are influenced by huge pendencies at various supply chain touchpoints. Accurate delivery time prediction lies in the ability to accurately project pendencies into the future, apart from modeling other

factors such as manpower, capacity, volume, holiday, weekend effects, etc.

4.3.1 Modeling approach. LightGBM model is trained using sale and post-sale data on the proxy and pre-proxy days taken from the recent sale event. The proxy day is chosen based on the number of days elapsed from the start of the sale until the prediction date. The planning data for the upcoming event is fed as input to the model along with the on-ground performance features derived from the past days of the sale. This enables the model to correct itself based on recent data in case of deviation from the plan.

4.3.2 Model features. Apart from the features considered for a BAU day, the following are the additional set of features considered to take into account various factors affecting shipping time during HRD

Planned volume and capacity features: The planned ship volume is modeled by combining sources appropriately to project landings at subsequent supply chain touchpoints. The planned last-mile capacity is indicative of the processing capability or cumulative outflow.

Pendency features: Three major components contribute to the delivery time once shipments arrive at the delivery center – arrival volume on the predicted landing date, pendency accumulated from the days prior to the predicted landing date, and delivery center outflow modeled using capacity plan data. These features are modeled using three different data sources or methods. In addition to modeling using plan data, most recent delivery center performance is captured as the event progresses and arrival volume and pendency at the delivery center is projected on the landing day using recent order and ship volume data, in order to combat any deviation from the plan.

4.4 Tree-based model to predict 3pl shipping time. LightGBM is used to model the shipping time for 3p courier partners. The holiday and weekend behavior at the pickup location as well as at the 3pl last-mile center was handled using the methodology described earlier based on the regional holiday calendar of the courier partners. However, some of the factors influencing shipping time as described above cannot be explicitly captured in the case of 3pl as we are not aware of the last-mile center locations, corresponding pincode mapping and arrival times. Pincode decoding is used to predict times for low-demand regions where limited or no recent data is available.

4.5 Methods to control breaches.

4.5.1 Asymmetric loss functions. Most machine learning models operate by default on loss functions that are minimized by central tendencies, either the mean or the median. The sum of squared residuals has its minimum at the mean, while the sum of absolute residuals is minimized by the median. Here, the goal is to penalize underestimations, that is when the predicted time is less than the actual time. A penalization factor is applied for the negative residuals on the default MSE (mean square error) loss[3]. The minimization of this loss function is indicated in equation(4.1).

$$e_i(\beta) = \begin{cases} \alpha(y_i - f(\beta, X_i)), & \text{if } y_i > f(\beta, X_i) \\ (y_i - f(\beta, X_i)) & \text{otherwise} \end{cases}$$

$$(4.1) \quad \beta^* = \underset{\beta}{\operatorname{argmin}} \left(\frac{1}{n} \sum_i e_i^2(\beta) \right)$$

The problem in using any asymmetric loss function of the above type is that predictions are still aimed at the expected value. We need to get at a prediction that is greater than the true value by a specified probability, that is, the breach cut-off. Multiplying the prediction obtained using MSE loss by a factor does not help us get at the desired result, as the central tendencies and higher percentiles are not linearly related.

4.5.2 Quantile regression. To maintain breaches below a specified percentage, we need to be able to predict the quantile of the probability distribution of future times. One way would be to get the full predictive densities and pick the desired quantile. However, the more powerful algorithms only give point predictions.

Quantile regression[7] is a modeling approach that aims to predict the percentiles by optimizing an objective that has its minimum at the specified quantile. The quantile loss function is the sum of asymmetrically weighted absolute residuals. The minimization of this loss function is shown in equation(4.2) and the loss function is visualized in Figure 2. The solution to the below equation is the estimate of the conditional quantile function. LightGBM supports quantile regression which can be specified using the objective parameter.

$$(4.2) \quad \beta^* = \underset{\beta}{\operatorname{argmin}} \sum_i \rho_\tau(y_i - f(X_i, \beta))$$

4.5.3 Feedback-based Breach control model design. In this approach, we built a feedback mechanism to correct the output from the LightGBM model (with default MSE loss function). We feed the predictions made in the last few days as input to a breach control

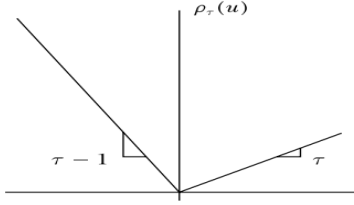


Figure 2: Quantile loss function[7]

model, along with the corresponding weekend, holiday handling times, holiday, seasonal features, and statistical measure-based features used to capture variations in linehaul and last-mile time. The target is constructed as a weighted combination of linehaul standard deviation, the mean and standard deviation of packet inflow and outflow at the last-mile. These weights are tuned based on past model performance to get the breaches below the specified percentage at a delivery date level. Model training is done on this dataset with ideal additional times as the target to predict the optimal time required to get breaches just below the desired level. The output from this model is added to the predictions made by the LightGBM model.

4.6 Time-series modeling approach.

4.6.1 Time series model to predict shipping time. In order to model time-dependencies, we used the time series approach. The Fbprophet model handles trend, seasonality, and holiday effects that occur on irregular schedules as indicated in equation(4.3). It can be seen from the weekly seasonality plot shown in Figure 3² that the shipping times are higher for orders shipped on Thursday, in case of short lanes, and for orders shipped on Monday and Tuesday, in case of longer lanes. This pattern is seen as these shipments would typically land in the last-mile center on a weekend. The *changepoint prior scale* parameter helps determine how flexible the trend should be. This parameter was tuned to avoid overfitting. Cap and floor parameters were set to specify the range within which the predictions can lie. Regional holiday list was inputted. The data were sampled based on hourly frequency. Predictions were made at an hourly level and the upper end of the prediction interval was used as the final result to control breaches. The model fitting is shown in Figure 4². Black dots indicate training data/known target values. The blue line indicates the predictions made on the train and test data. The blue shaded region indicates the interval

²The y-axis numbers are hidden to not reveal internal business numbers.

for the predictions.

$$(4.3) \quad y(t) = g(t) + s(t) + h(t) + \epsilon_t$$

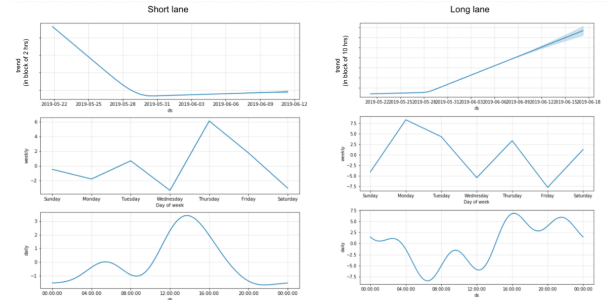


Figure 3: Trend, weekly and daily seasonality plots for short and long lanes

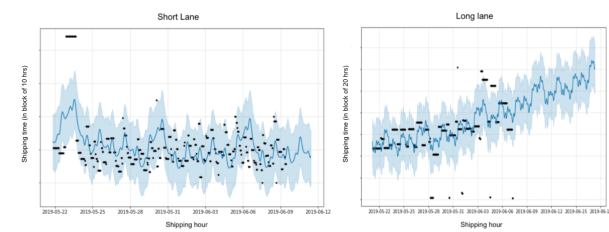


Figure 4: Fbprophet model fitting for a short versus long lane.

4.6.2 Time series model to predict warehouse time. Fbprophet model was used to predict the warehouse processing time. The time taken from order placement to pick is shown in Figure 5². The periodic patterns are seen due to shift timings. Models were built separately for single-item and multi-item orders as multi-item packets have an additional wait time involved to consolidate all items. The seasonal effects by the hour and day of the week are captured by the fbprophet model as the model is built on the past few weeks of data. The time-series data was resampled based on minute-level frequency. Re-training is done every hour to capture the recent changes and predictions are made at a minute level for the next hour.

As the warehouse processing time is small even when compared to a short lane, the model was able to capture recent changes effectively and provide accurate real-time predictions. A LightGBM model was also built using the projected load, multi-item based, seasonal and performance-based features.

4.7 Deep-learning based approach.

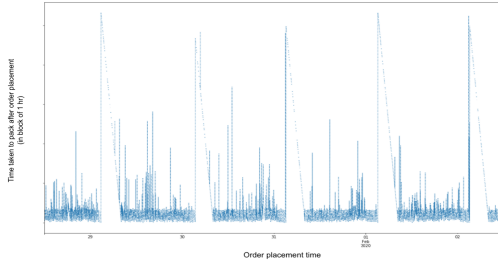


Figure 5: Time series for Warehouse processing leg (Order placement to pick)

4.7.1 LSTM model to predict shipping time.

LSTM is a powerful recurrent neural network that handles sequence dependence among input variables. The network was designed to have 7 LSTM blocks in the hidden layer to capture seasonality. This LSTM network was trained for 300 epochs using a batch size of 100. The lookback value was chosen to cover few weeks after tuning. The other relevant features affecting shipping time were also included. From Figure 6, we can see that the model is able to fit the dataset well.

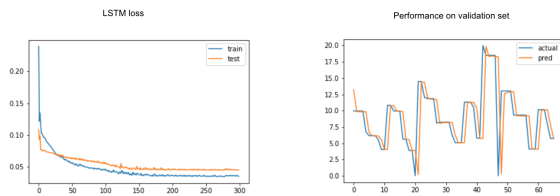


Figure 6: LSTM loss and performance on the validation set

4.8 Tree-based model to predict vendor procurement time.

LightGBM was used to model the time taken to procure the items not available in the warehouse from the vendor and subsequently ship the items. Several features were built including vendor type, features based on designated hours at which purchase order is generated and the times at which pick up happens, location and route-based features, seasonality features, the maximum time allowed to process a purchase order, and pick up in coloader vehicles.

4.9 Supply-chain optimization.

Modeling delivery times accurately not only helps in improving customer experience and reducing customer churn but also in providing business insights to optimize the supply chain operations. When the on-ground processing time shows deviation from the predicted time, it can be used to flag and subsequently correct unexpected pendencies or delays at an early stage. Analysis on last-mile

center processing speed showed that making all delivery attempts within a fixed time duration post arrival would help in reducing breach. This further enables us to provide more accurate times.

5 Experimental Results.

In this section, we discuss the experimental results, where we considered a large dataset for training and evaluation. We examine various models such as variants of LightGBM, XGBoost, Random forest, LSTM, and time series models, and compare the performance with the current baseline of rule-based models. We also present results for HRD and different linehaults. We summarise the results for various supply chain legs such as Warehouse, Shipping(3pl), and Vendor procurement.

5.1 Data. The training data constitutes deliveries made in the last few weeks prior to the prediction date amounting to more than a **million training examples** and the testing is done on the orders placed on the subsequent day.

5.2 Evaluation metrics. Accuracy is measured as the percentage of total packets ordered in a day that get delivered on the same day(0) as the promise date or the previous day(-1) and is denoted as Accuracy(0 to -1). Breaches refer to the percentage of total packets ordered in a day that get delivered after the promised date. These are delivery date-based metrics measured at an order-day level.

5.3 Results. Table 2 shows order-day level metrics averaged over a period of one week in the month of September, thus covering weekend effects. The results cover all possible lanes amounting to around thousands and correspond to only items stored in the warehouse. The models mentioned in the table were used to predict shipping time. The time from order placement to shipping is taken from the rule-based model.

It can be seen from Table 2 that the quantile regression approach achieved the highest accuracies, almost doubling the accuracy achieved using the baseline model, while also reducing the breaches. LightGBM trained with asymmetric loss function resulted in relatively higher breach, next to the rule-based model. Though the penalization factor was tuned, breaches could not be deterministically controlled using this approach. LightGBM in conjunction with the breach control model has achieved relatively a lower breach while taking a slight hit on accuracy. The breach control model is trained based on breaches at a delivery date level to meet recency criteria. For every order date, the factors are tuned to meet breach constraints at all the delivery dates when

Table 2: Comparison of model performance covering all possible lanes on Business As Usual(BAU) days

Model	Accuracy	Breach
LightGBM (Quantile regression)	68.10%	4.68%
LightGBM (Asymmetric loss function)	61.54%	5.50%
LightGBM (with breach control model)	58.30%	3.90%
XGBoost (with breach control model)	58.13%	3.82%
LSTM (with breach control model)	57.25%	6.20%
Random forest (with breach control model)	55.45%	3.85%
Time series model (fbprophet)	40.30%	2.74%
Rule-based model (Baseline)	36.56%	5.63%

those shipments landed. The factor tuning is thus also influenced by performance on other order dates. Therefore, this method gives a more conservative breach result. LightGBM model accuracy is slightly higher than that achieved using XGBoost, while model training with LightGBM was 8x times faster. The Random forest model provided lower accuracies when compared to both LightGBM and XGBoost models.

Table 3: Comparison of model performance covering all possible lanes on HRD and post-HRD days

Model	Accuracy	Breach
LightGBM - HRD	65.15%	7.86%
LightGBM - Post-HRD	74.58%	4.82%
Rule-based model - HRD	44.23%	5.44%
Rule-based model - Post-HRD	54.48%	3.06%

Table 3 shows order-day level metrics averaged over the HRD and post-HRD days in the month of December. Accuracy(0 to -2) is the performance metric measured in the case of HRD as predictability is more challenging compared to BAU days. The results cover all possible lanes and correspond to both items stored in the warehouse as well as those procured from the vendor. The models mentioned in the table were used to predict shipping time. The time from order placement to ship is assumed to be known to indicate the performance of the shipping time model without considering the inaccuracies resulting from the pre-ship leg. Note that since the breaches at the warehouse side are controlled during HRD, the breaches at the time of order placement were still under control despite a slightly higher breach percentage on the shipping leg. It can be seen from Table

3 that the best LightGBM implementation achieved an improvement of more than 20 percentage points over the baseline model in the case of both HRD and post-HRD days.

Table 4: Comparison of model performance for a short versus long lane

Model	Short Lane (Delhi - Gurugram)		Long lane (Delhi - Kolkata)	
	Accuracy	Breach	Accuracy	Breach
Time-series	83.55%	0.52%	50.34%	2.54%
LSTM	83.55%	0.52%	70.65%	7.24%
Light-GBM	83.55%	0.43%	75.26%	5.28%

The comparison of LSTM (in conjunction with breach control model), fbprophet (Time series model), and LightGBM (Quantile regression implementation) models was carried out for a short versus long lane as shown in Table 4. It describes order-day level metrics averaged over a period of one week in the month of September and corresponds to only items stored in the warehouse. LSTM and fbprophet worked on par with LightGBM in the case of shorter lanes. While for longer lanes, the time series model has to predict much farther into the future due to non-availability of recent target values. This can be seen in Figure 4². This leads to reduced confidence in the predictions. Fbprophet picked up recent trend changes and made predictions much ahead into the future based on these changes, thus resulting in overfitting. In the supply chain problem, the recent past has an impact only on the immediate future but the effect fades away as we go into the future. The LightGBM model tends to average out the effects while taking into account the potential effect from other factors on the prediction day. In the case of LSTM as well, predictions much ahead into the future are made based on existing predictions. However, it performed much better than the time series model. The results follow a similar pattern across all lanes as shown in Table 2.

Table 5: Comparison of Model Performance for the various supply chain legs

Supply chain leg	Metrics	Rule-based	Time series	Light-GBM
Warehouse	Accuracy	82.23%	88.19%	86.20%
	Breach	3.79%	6.54%	5.32%
Shipping (3pl)	Accuracy	9.41%	-	55.52%
	Breach	1.24%	-	3.70%
Vendor - customer	Accuracy	8.65%	-	40.84%
	Breach	1.34%	-	5.06%

The results shown in Table 5 are also reported at an order-day level by averaging over a period of one week. In the case of 3pl, the models mentioned in the table were used to predict shipping time. The time from order placement to shipping is taken from the rule-based model, similar to the reporting done in Tables 2 and 4. For the warehouse leg, the accuracy and breaches are reported for the predicted ship date (time from order placement to ship is predicted by the models mentioned in the table). These results correspond to only items stored in the warehouse. The results reported against vendor-customer are accuracies and breaches for the predicted delivery date. These results correspond to only items not stored in the warehouse as the vendor procurement stage is applicable only for such items. The models mentioned in the table are used to predict the time taken from order placement to shipping and the time taken after shipping to reach the customer is taken from the best LightGBM implementation. As the warehouse and vendor procurement models are in the early stages, we have reported the metrics in Tables 2 and 4 by taking the time until shipping from the rule-based model.

It can be seen that the fbprophet model has resulted in an accuracy improvement of 6 percentage points over the rule-based model. The breaches are slightly higher as no breach control mechanism is applied at the warehouse processing leg, it is compensated for when combined with predicted times of other order processing legs. It also performed slightly better than the LightGBM model in this case, as the recent changes in performance could be captured to give accurate predictions. The LightGBM model for 3pl has achieved a 45% jump in accuracy while working under the constraints imposed by 3pl data. Again in the case of vendor processing time where we do not have complete control, accuracy has improved by 31%, with a breach percentage slightly above the cutoff level.

6 Summary.

Promise time computation is a challenging problem in e-commerce given uncertainties at various stages such as variability in vendor processing time, consolidation wait time, dependencies and additional uncertainties at the warehouse, traffic disruptions, weather conditions and hub wait times during linehaul, pendency, manpower variation, and customer availability at last-mile, regional holiday and weekend effects at all of the above legs. We propose a solution that integrates all these aspects. The solution is deployed and operational for Myntra. In the paper, we discuss the challenges in detail and provide empirical evidence through data analysis. We have considered a number of models before zeroing on the final solution such as multiple variants of LightGBM,

XGboost, Random forest, time series model (fbprophet), and deep learning approaches. We have carried out elaborate experiments on more than a million patterns and compared a range of approaches covering these models, while also specifying methods to meet the business constraints using **asymmetric loss functions** and a **feedback-based breach control model**. The LightGBM implementation of quantile regression gave the best results among the chosen models, almost doubling the accuracy achieved with the baseline model. We have proposed a model for handling high revenue days, where predictability is more challenging due to higher variability in times, deviation from the plan, and process changes. We also built a promise time prediction model for 3pl. As future work, we propose to enhance vendor procurement and warehouse models and comprehensively provide end-to-end predictability.

References

- [1] T. Ravindra Babu, Abhranil Chatterjee, Shivram Khandedarker, A. Vamsi Subhash, and Sawan Gupta. Geographical address classification without using geolocation coordinates. In *Proceedings of the 9th Workshop on Geographic Information Retrieval*, GIR '15, New York, NY, USA, 2015. Association for Computing Machinery.
- [2] Jeff Berman. Logistics management. 2017. <https://www.logisticsmgmt.com>.
- [3] Peter F Christoffersen and Francis X Diebold. Further results on forecasting and model selection under asymmetric loss. *Journal of applied econometrics*, 11(5):561–571, 1996.
- [4] Aditya Jariwala Rishabh Mohan Matthew A. Lanham Douglas Halim, Alexander Hartman. Delivery time estimation for industrial equipment. 2018. <https://kranmert.purdue.edu>.
- [5] Antoine Jonquais and Florian Krempf. Predicting shipping time with machine learning. 2019.
- [6] Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. Lightgbm: A highly efficient gradient boosting decision tree. In *Advances in neural information processing systems*, pages 3146–3154, 2017.
- [7] Roger Koenker and Kevin F Hallock. Quantile regression. *Journal of economic perspectives*, 15(4):143–156, 2001.
- [8] Jundi Liu, Steven Hwang, Walter Yund, Linda Ng Boyle, and Ashis G Banerjee. Predicting purchase orders delivery times using regression models with dimension reduction. In *ASME 2018 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*. American Society of Mechanical Engineers Digital Collection, 2018.
- [9] Sean J Taylor and Benjamin Letham. Forecasting at scale. *The American Statistician*, 72(1):37–45, 2018.