



How Much Can A Retailer Sell? Sales Forecasting on Tmall

Chaochao Chen^(✉), Ziqi Liu, Jun Zhou, Xiaolong Li, Yuan Qi, Yujing Jiao,
and Xingyu Zhong

Ant Financial Services Group, Hangzhou 310099, China
{chaochao.ccc, ziqiliu, jun.zhoujun, xl.li, yuan.qi, yujing.jyj,
xingyu.zxy}@antfin.com

Abstract. Time-series forecasting is an important task in both academic and industry, which can be applied to solve many real forecasting problems like stock, water-supply, and sales predictions. In this paper, we study the case of retailers' sales forecasting on Tmall—the world's leading online B2C platform. By analyzing the data, we have two main observations, i.e., *sales seasonality* after we group different groups of retailers and a *Tweedie distribution* after we transform the sales (target to forecast). Based on our observations, we design two mechanisms for sales forecasting, i.e., seasonality extraction and distribution transformation. First, we adopt Fourier decomposition to automatically extract the seasonalities for different categories of retailers, which can further be used as additional features for any established regression algorithms. Second, we propose to optimize the Tweedie loss of sales after logarithmic transformations. We apply these two mechanisms to classic regression models, i.e., neural network and Gradient Boosting Decision Tree, and the experimental results on Tmall dataset show that both mechanisms can significantly improve the forecasting results.

Keywords: Sales forecasting · Tweedie distribution ·
Distribution transform · Seasonality extraction

1 Introduction

Time-series forecasting is an important task in both academic [4] and industry [16], which can be applied to solve many real forecasting problems including stock, water-supply, and sales predictions. In this paper, we study the forecasts of retailers' future sales at Tmall.com¹, one of the world's leading online business-to-customer (B2C) platform operated by Alibaba Group. The problem is essentially important because accurate estimation of future sales for each retailer can help evaluate and assess the potential values of small businesses, and help discover potentials for further investment.

¹ <https://en.wikipedia.org/wiki/Tmall>.

C. Chen and Z. Liu—Equal contribution.

However, accurately estimating each retailer’s future sales on Tmall could be way challenging in several reasons. First of all, naturally different goods or products are sold with strong seasonal properties. For example, most of fans are sold in summer, while most of heaters are sold in winter, i.e. we can observe strong seasonal properties on different groups of retailers. Secondly, the distribution of sales among all retailers demonstrates an over-tail power law distribution, i.e., the retailers’ sales spread a lot. Naively ignore such issues could make the performance much worse.

In this paper, we analyze and summarize the characteristics of the sales data at Tmall.com, and propose two mechanisms to improve forecasting performance. On one hand, we propose to extract seasonalities from groups of retailers. Specifically, we characterize the seasonal evolutions of sales by first clustering retailers into groups and then study the seasonal series as decompositions from a series of Fourier basis functions. The results of our approach can be utilized as features, and simply added to the feature space of any established machine learning toolkits, e.g., linear regression, neural network, and tree-based model. On the other hand, we place distribution transformations on the original retailers’ sales. Specifically, we observe that the distribution over retailers’ sales follows a Tweedie distribution after we transform retailers’ sales by logarithmic. Thus, we propose to optimize Tweedie loss for regression on the logarithmic transformed sales data instead of other losses on the original ones. Empirically, we show that the proposed two mechanisms can significantly improve the performance of predicting retailers’ future sales by applying them into both neural networks and tree based ensemble models.

We summarize our main contributions as follows:

- By analyzing the Tmall data, we obtain two observations, i.e., *sales seasonality* after we group different categories of retailers and a *Tweedie distribution* after we transform the original sales.
- Based on our observations, we design two general mechanisms, i.e., seasonality extraction and distribution transform, for sales forecasting. Both mechanisms can be applied to most existing regression models.
- We apply the proposed two mechanisms into two popular existing regression models, i.e., neural network and Gradient Boosting Decision Tree (GBDT), and the experimental results on Tmall dataset demonstrate that both mechanisms can significantly improve the forecasting results.

2 Data Analysis and Problem Definition

In this section, we first describe the sales data and features on Tmall. We then analyze the seasonality and distribution of sales data. Finally, we give the sales forecasting problem a formal definition.

2.1 Sales Data and Feature Description

Tmall.com is nowadays one of the largest business-to-customer (B2C) E-commerce platform. It has more than 180,000 retailers. Among of those retailers,

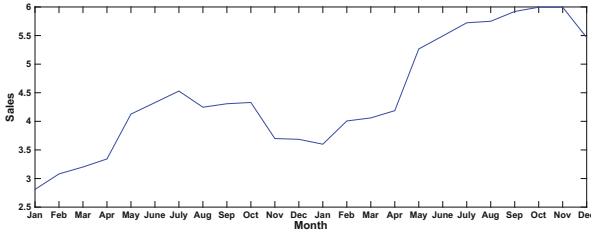


Fig. 1. GMV on Tmall. The horizontal axis denotes the months between Jan. 2015 to Dec. 2016, and the vertical axis denotes the GMV on Tmall where we omit the scale.

there could be giant retailers like Apple.com, Prada, and together with small businesses. The platform is selling hundreds of thousands products in diverse categories, e.g., ‘furniture’, ‘snack’, and ‘entertainment’.

Besides category information, the other features of retailers on Tmall can be mainly divided into three types: (1) The basic features that are able to reflect the marketing and selling capability of each retailer. For example, the amount of advertisement investment, the number of buyers, the rating/review given by the buyers, and so on. (2) The high-level features that are generated from historical sales and basic feature data. Suppose a retailer i generates a series of sales data, e.g., $y_{i,t-2}, y_{i,t-1}, y_{i,t}, y_{i,t+1}$. We are currently at time t and want to forecast $y_{i,t+1}$. Then $y_{i,t}$ can be taken as a feature which indicates the sale amount of previous period, $y_{i,t} - y_{i,t-1}$ is a feature that indicates the increasing speed of the sales, and $(y_{i,t} - y_{i,t-1}) - (y_{i,t-1} - y_{i,t-2})$ is a feature that denotes the accelerated speed of the sales. Similarly, we can generate other high-level features, e.g., the number of buyers, using the basic features available. (3) The seasonality features that are generated by using other machine learning techniques, which aim to capture the seasonal property of different retailers. We will present how to generate these features in Sect. 3.1.

2.2 Seasonality Analysis

The retailers’ sales tend to have different seasonality due to the seasonal items they sell. Take vegetables for example, tomatoes and cucumbers are usually sold more in summer, while celery cabbage is likely sold more in winter. Although the GMV demonstrate seasonal properties, as is shown in Fig. 1, the analysis of the seasonality related to the gross merchandise volume is relative meaningless for the prediction on each retailer. In contrast, the seasonality analysis on each single retailer makes the analysis cannot generalize well in the future. Instead, we further investigate the seasonalities in different groups of retailers.

By analyzing the sales data on Tmall, we observe different seasonal patterns on different categories of retailers. Figure 2 shows the sales of four different categories of retailers, i.e., ‘Women’s Wearing’, ‘Men’s Wearing’, ‘Snack’, and ‘Meat’, where we use two year’s sales data from January 2015 to December 2016. We can observe that, ‘Women’s Wearing’ and ‘Men’s Wearing’ show quite similar

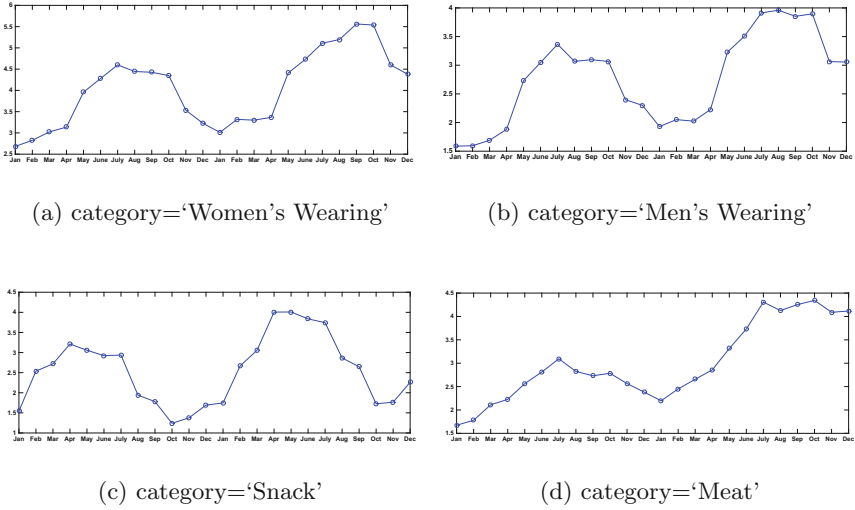


Fig. 2. Sales seasonality of different categories on Tmall. The horizontal axis denotes the months between Jan. 2015 to Dec. 2016, and the vertical axis denotes the total sale among of retailers in each category where we omit the scale.

seasonal patterns, i.e., they both reach peak in summer (July or August) and decline to nadir in winter (January). On the contrary, 'Snack' and 'Meat' show different seasonal patterns. In summary, the seasonalities under different categories could differ quite a lot. Thus if we can somehow partition the retailers into appropriate groups, the shared seasonality among retailers in one group could be statistically useful for characterizing each retailer in the group. Given a group of retailers, how can we characterize the seasonality for the group remains to be solved. We will discuss our approaches in Sect. 3.1.

2.3 Sale Amount Analysis

The sales of each retailer over time-series could be much challenging. To illustrate this, we show the histograms over sales in Fig. 3 (left). It shows that the sales could be much diverse across over all the retailers. In practice, this is very hard to formalize as a trivial regression problem because the errors on those sales from giant retailers could dominate the loss, e.g. least squared loss.

Instead, after we do a logarithmic transformation on the sales of each retailer, we found that the histogram appears to be a clear Tweedie distribution, i.e. Fig. 3 (right), which will be further described in details in Sect. 3.2. As a result, such transformation on the dependent variables makes our forecasting much easier. Note that, there are always some retailers' sale around zero. This is because some shops on Tmall will close or forced to be closed by Tmall due to some reason from time to time, and correspondingly, some shop will be newly opened or reopen. Consequently, some retailers' sales are around zero.

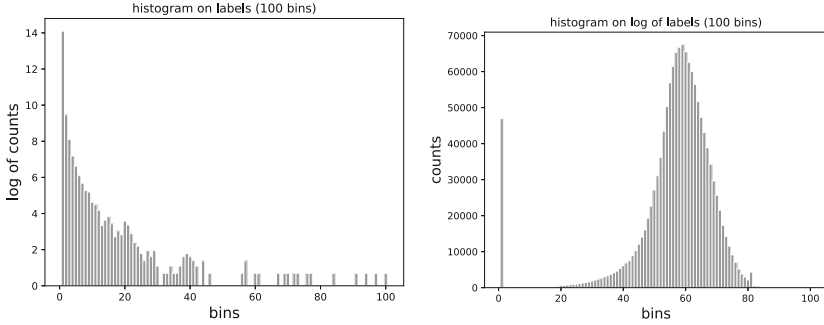


Fig. 3. Sales on Tmall obey Tweedie distribution after logarithmic transformation. We order the sales among retailers in an increasing order, partition the sales into 100 bins with equal frequency binning (x-axis), and show the retailer counts in each bin (y-axis).

2.4 Problem Definition

Assuming any retailer in Tmall.com as i , at month t , We formalize the sales forecasts problem as a regression problem, i.e. given the features of each retailer $x_{i,<t} \in \mathbb{R}^d$, where d is the feature dimensionality and $< t$ denotes the months before t , and the known sales of each retailer $y_{i,<t+1}$, we want to learn a function: $f : x_{i,t} \mapsto y_{i,t+1}$, where $x_{i,t}$ denotes the features of retailers i at month t and $y_{i,t+1}$ denotes its sales at month $t + 1$. That is, give the features of any retailer at month t , we want to predict their corresponding sales at month $t + 1$.

3 Model Design and Implementation

In this section, we will present our designed two mechanisms, i.e., seasonality extraction and distribution transform, for sales forecasting.

3.1 Seasonality over Groups of Retailers

As we reported in Sect. 2.2, the seasonalities under different categories could differ quite a lot, therefore, the remaining problem is that how should we partition the retailers into appropriate groups. Instead of manually partition retailers, we adopt clustering methods for time-series data [14] to do so. Specifically, we group the retailers by using the basic and high-level features described in Sect. 2.1, so that retailers that have similar features are grouped together.

After we partition retailers into groups, we adopt discrete Fourier transform to automatically extract the seasonality for retailers in different groups. Formally, assuming a group of sellers with expected amount of sales annotated as $\tilde{y}(t)$ at time t , thus results into a series of expected sales as observations, i.e.

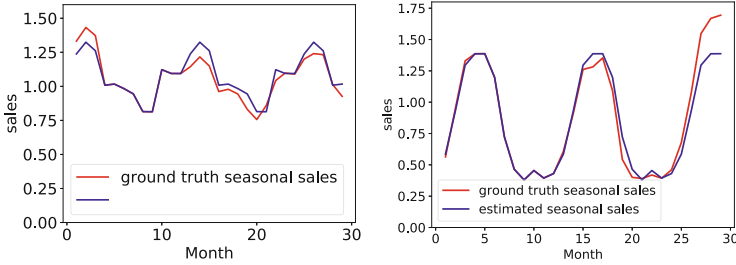


Fig. 4. Seasonality extraction results for two groups of retailers. The left one is mainly the group of retailers who sell purses, and the right one is mainly the group of retailers who sell accessories. In both figures, we use the first 15 months data to learn the parameters in Eq. (1), and further use them to predict the seasonality of all the months’ data (note that we also omit the scale of the sale amount).

$\{\tilde{y}(0), \dots, \tilde{y}(t), \dots, \tilde{y}(T)\}$. Each periodic function $\tilde{y}(\cdot)$ can be expanded by the Fourier series, which is a linear combination of infinite sines and cosines,

$$\tilde{y}(t) = a_0 + \sum_{n=1}^{\infty} a_n \cos(nt) \sum_{n=1}^{\infty} b_n \sin(nt), \tag{1}$$

where $\{a_i, b_i | i \in \{0, \dots, n, \dots, \infty\}\}$ are parameters to be optimized. As a result, the function $\tilde{y}(\cdot)$ can be represented by a Fourier basis.

We now show the results of extracted seasonalities on different groups of retailers. We randomly select two groups of retailers and show their seasonalities and estimates for sales in Fig. 4, where we find the two groups of retailers mainly sell purses and accessories, respectively. In Fig. 4, we use the first 15 months’ data to learn the parameters in Eq. (1), and further use them to predict the seasonality of all the months’ data. It is obvious that our extracted seasonality is very close to the real one in both groups. Similarly, we can extract seasonality for other features, e.g., the number of buyers and the among of advertisement investment, by using the same method.

In practice, we use two types of features extracted from such seasonal patterns: (1) the seasonal values of the target we want the extract, e.g., sales and the number of buyers, in a window of 12 months centered around the month t . (2) the variation, i.e. the difference among those seasonal values. Hopefully, such seasonality or trend measures for each group of sellers can be fed into any classifiers, so as to characterize the seasonal patterns for each seller. We will empirically study the effectiveness of these seasonality features in experiments.

3.2 Tweedie Loss for Regression

As we described in Sect. 2.3, based on our observation, the sales on Tmall will clearly obey Tweedie distribution after a logarithmic transformation. From Fig. 3 (right), we see that the sales after logarithmic transformation is a combination of

Poisson distribution and Gamma distribution, which is a special case of Tweedie distribution, i.e., a compound Poisson-Gamma distribution. That is, we assume that (1) the status of retailers, i.e., closed or not, are independent identically distributed and they obey Poisson distribution; (2) the sales of retailers are also independent identically distributed and they obey Gamma distribution. The Tweedie distribution was first proposed in [22], and then officially named by Bent Jorgensen in [9], which belongs to the class of exponential dispersion.

Tweedie distribution has been popularly used in insurance scenarios [23]. We now formally describe Tweedie distribution in sale forecasting scenario. Suppose Let N be a Poisson random variable denoted by $Pois(\lambda)$, and let $Z_i, i = 0, 1, 2, \dots, N$ be independent identically distributed gamma random variables denoted by $Gamma(\alpha, \gamma)$ with mean $\alpha\gamma$ and variance $\alpha\gamma^2$. We also assume that N is independent of Z_i . Define a random variable Z by

$$Z = \begin{cases} 0, & \text{if } N = 0, \\ Z_1 + Z_2 + \dots + Z_N, & \text{if } N > 0. \end{cases} \tag{2}$$

We can see from Eq. (2) that Z is the Poisson sum of independent Gamma random variables, which is also called compound Poisson-Gamma distribution. In sales forecasting scenarios, Z can be viewed as the total number of retailers, N as the opened retailers, and Z_i as the sale amount of retailers i . Note that the distribution of Z has a probability mass at zero, i.e., $Pr(Z = 0) = exp(-\lambda)$. The existing research has proven that, if we reparametrize $(\lambda, \alpha, \gamma)$ by

$$\lambda = \frac{1}{\phi} \frac{\mu^{2-p}}{2-p}, \alpha = \frac{2-p}{1-p}, \quad \gamma = \phi(\rho-1)\mu^{\rho-1},$$

Eq.(2) then becomes the form of a Tweedie model $Tw(\mu, \phi, \rho)$ with $1 < \rho < 2$ and $\mu > 0$. Here, the boundary cases $\rho \rightarrow 1$ and $\rho \rightarrow 2$ correspond to the Poisson and the gamma distributions, respectively. The compound Poisson-Gamma distribution with $1 < \rho < 2$ can be seen as a bridge between the Poisson and the Gamma distributions.

The log-likelihood of this Tweedie model for the sale y of a retailer is

$$L(y|\mu, \phi, \rho) = \frac{1}{\phi} \left(y \frac{1}{\phi} \frac{\mu^{1-p}}{1-p} - \frac{1}{\phi} \frac{\mu^{2-p}}{2-p} \right) + \log(a(y, \phi, \rho)), \tag{3}$$

where the normalizing function $a(\cdot)$ can be written as

$$a(y, \phi, \rho) = \begin{cases} \frac{1}{y} \sum_t W_t(y, \phi, \rho), & \text{for } y > 0, \\ 1, & \text{for } y = 0, \end{cases}$$

and $\sum_t W_t$ is an example of Wright’s generalized Bessel function [22].

After that, given the parameter ρ for Tweedie model, the other parameters can be efficiently solved by using maximum log-likelihood approach [23]. The Tweedie model can be naturally combined with most existing regression models, e.g., NN and GBDT. That is, we can train a Tweedie loss NN model or GBDT

model instead of the models with other losses, e.g., square loss [23]. Obviously, the results of Tweedie loss regression are much better than those of other loss regression, e.g., square loss, as will be shown in experiments. This is because Tweedie loss fits the real sales distribution after logarithmic transformation of sales, as is shown in Fig. 3 (right).

4 Empirical Study

In this section, we first describe the dataset and the experimental settings. Then we report the experimental result by comparing with various state-of-the-art sales forecasting techniques. We finally analyze the effect of Tweedie distribution parameter (ρ) on model performance.

4.1 Dataset

Features. As we described in Sect. 2, the features of retailers mainly contain three types, i.e., the basic features, high-level features that are generated from historical sales and basic feature data, and the seasonality features that are generated by using other machine learning techniques. This includes 189 features in total, where there are 79 basic features, 102 high-level features, and 8 seasonality features as we discussed in Sect. 3.1.

Samples. We choose the samples (retailers) during Jan. 2015 and Dec. 2016 on Tmall. Note that we only focus on forecasting the relative small retailers whose monthly sale amount is under a certain range (300,000). Because, in practice, the sales of big retailers are very stable, and it is meaningless to forecast their sales. After that, we have 783,340 samples. We use the samples in 2015 as training data, the samples from Jan. 2016 to June 2016 as validation, and the samples from July 2016 to Dec. 2016 as test data.

4.2 Experimental Settings

Evaluation Metric. Most existing research use error-based metric, e.g., Mean Average Error (MAE) and Root Mean Square Error (RSME), to evaluate the performance in time-series forecasting [1, 6]. However, these metrics are way sensitive to those retailers whose sales are large. As we can see in Fig. 1, the sales on Tmall spread a lot. In practice, the forecasting precision of the retailers with small sales counts the same as the ones with big sales. Therefore, we propose to use Relative Precision (RP) for sales forecasting on Tmall, which is defined as

$$RP@p = \frac{\sum_{i=1}^N \mathbf{1}\left(\frac{|y_i - \hat{y}_i|}{y_i} < p\right)}{N}, \quad (4)$$

where N is the total number of retailers, y_i as the real sale and \hat{y}_i as the forecasted sale, $p \in [0, 1]$, and $\mathbf{1}(\cdot)$ is the indicator function that equals to 1 if the expression in it is true and 0 otherwise.

Table 1. Comparison result on test data.

Model	NN	NN-S	NN-T	NN-ST	GBDT	GBDT-S	GBDT-T	GBDT-ST
RP@0.1	0.1693	0.1723	0.3236	0.3338	0.1719	0.1859	0.3159	0.3263
RP@0.2	0.1933	0.1987	0.3484	0.3534	0.2095	0.2242	0.3394	0.3520
RP@0.3	0.2603	0.2657	0.3950	0.3956	0.2681	0.2816	0.3821	0.3966

As we can see from Eq. (4), RP is actually the percentage of the retailers whose forecasting error is in a certain range p . Intuitively, the smaller p is, the smaller RP will be. Because one has higher demanding for the forecasting performance when p is smaller.

Comparison Methods. Our proposed mechanisms, i.e., seasonality extraction and distribution transform, has the ability to generalize to most existing regression algorithms. To prove this, we apply the mechanisms into two popular regression models, i.e., Neural Network (NN) and Gradient Boosting Decision Tree (GBDT). We summarize all the methods, including ours, as follow:

- **NN** has been used to do time-series forecasting and proven effective where we use square loss [1, 20].
- **NN-S** uses extra our proposed seasonal feature in Sect. 3.1 for NN, and its comparison with NN will prove the effectiveness of seasonality extraction.
- **NN-T** uses our proposed Tweedie-loss in Sect. 3.2 for NN, and its comparison with NN will prove the effectiveness of our proposed Tweedie-loss regression after sale distribution transform.
- **NN-ST** extra uses our proposed seasonal feature in Sect. 3.1 for NN-T, which is the application of our proposed two mechanisms in NN.
- **GBDT** is developed for additive expansions based on any fitting criterion, which belongs to a general gradient-descent ‘boosting’ paradigm and suits for regression tasks with many types of loss functions, e.g., least-square loss, Huber loss, and Tweedie loss [7]. Specifically, we use the GBDT algorithms implemented on Kungpeng [26]—a distributed learning system that is popularly used in Alibaba and Ant Financial, where we also use square loss.
- **GBDT-S** uses extra seasonal feature for GBDT, similar as NN-S.
- **GBDT-T** uses Tweedie-loss for GBDT, similar as NN-T.
- **GBDT-ST** uses extra seasonal feature for GBDT-T, similar as NN-ST.

Parameter Setting. For NN, we use a three-layer network, with Rectified Linear Unit (ReLU) as active functions, and optimized with Adam [12] (learning rate as 0.1). For GBDT, we set tree number as 120, learning rate as 0.3, and regularizer of ℓ_2 norm as 0.5. We will study the effect of parameter ρ of Tweedie regression in Sect. 4.4.

4.3 Comparison Results

We summarize the comparison results in Table 1, and have the following comments.

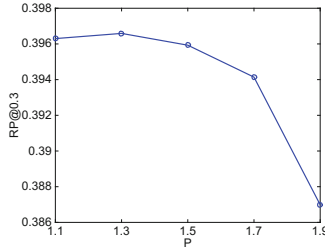


Fig. 5. Effect of Tweedie distribution parameter (ρ) on GBDT-ST on validate data.

(1) The forecasting performance of NN and GBDT are close, and the performance of GBDT is slightly higher than NN. This is because GBDT can naturally consider the complicate relationship, e.g., cross feature, between features. (2) Our proposed seasonality extraction mechanism can clearly improve the forecasting performance of both NN and GBDT. For example, GBDT-S improves the forecasting performance of GBDT by 8.14% in terms of RP@0.1, and GBDT-ST further improves the forecasting performance of GBDT-T by 3.29%. (3) Our proposed distribution transform mechanism can significantly improve the forecasting performance of both NN and GBDT. For example, NN-T improves the forecasting performance of NN by 91.14% in terms of RP@0.1, and NN-ST improves the forecasting performance of NN-S by 93.73% in terms of RP@0.1 (4) In summary, our proposed two mechanisms consistently improve the forecasting performances of both NN and GBDT models. Specifically, NN-ST improves the forecasting performance of NN by 97.14%, 82.82%, 51.98% in terms of RP@0.1, RP@0.2, and RP@0.3 respectively. And, GBDT-ST improves the forecasting performance of GBDT by 89.82%, 68.10%, 47.93% in terms of RP@0.1, RP@0.2, and RP@0.3 respectively. The results not only demonstrate the effectiveness of our proposed mechanisms, but also indicate the generalizability of them.

4.4 Effect of Tweedie Distribution Parameter (ρ)

As described in Sect. 3.2, the Tweedie distribution parameter (ρ) bridges the Poisson and the Gamma distributions, and the boundary cases $\rho \rightarrow 1$ and $\rho \rightarrow 2$ correspond to the Poisson and the Gamma distributions, respectively. The effect of Tweedie distribution parameter (ρ) on GBDT-ST is shown in Fig. 5, where we use the validate data. From it, we find that GBDT-ST achieves the best performance when $\rho = 1.3$. This indicates that the real sales data on Tmall fit the Tweedie distribution when $\rho = 1.3$.

5 Related Works

In this section, we will review literatures on time-series forecasting, which are mainly in two types, i.e., linear model and non-linear model.

5.1 Linear Model

The most popular linear models for time-series forecasting are linear regression and Autoregressive Integrated Moving Average model (ARIMA) [8]. Due to their efficiency and stability, they have been applied to many forecasting problems, e.g., wind speed [10], traffic [21], air pollution index [13], electricity price [3], and Inflation [19]. However, since it is difficult for them to consider complicated relations between features, e.g., cross feature, their performance are limited.

5.2 Non-linear Model

Non-linear models are also adopted for time-series forecasting. The most popular ones are Support Vector Machine (SVM), neural network, and tree-based ensemble models. For example, SVM are applied to financial forecasting [11] and wind speed forecasting [15]. Neural network are also used in financial marketing forecasting [2] and electric load forecasting [18]. Recently, Gradient Boosting Decision Tree (GBDT) are also adopted to forecast traffic flow [24].

Moreover, model ensemble is also popular for time-series forecasting. For example, ARIMA and SVM were combined to forecast stock price [17]. Hybrid ARIMA and NN models were also used for time-series forecasting [5, 25].

In this paper, we do not focus on the choices of regression models. Instead, based on our observation, we focus on extracting seasonality information and transforming label for better forecasting performance. Our proposed seasonality extraction and label distribution transform can be applied into most forecasting models, including NN and GBDT.

6 Conclusions

In this paper, we studied the case of retailers' sales forecasting on Tmall—the world's leading online B2C platform. We first observed sales seasonality after we group different categories of retailers and Tweedie distribution after we transform the sales. We then designed two mechanisms, i.e., seasonality extraction and distribution transform, for sales forecasting. For seasonality extraction, we first adopted clustering method to group the retailers so that each group of retailers have similar features, and then applied Fourier transform to automatically extract the seasonality for retailers in different groups. For distribution transform mechanism, we used Tweedie loss for regression instead of other losses that do not fit the real sale distribution. Our proposed two mechanisms can be used as add-ons to classic regression models, and the experimental results showed that both mechanisms can significantly improve the forecasting results.

References

1. Ahmed, N.K., Atiya, A.F., Gayar, N.E., El-Shishiny, H.: An empirical comparison of machine learning models for time series forecasting. *Econometr. Rev.* **29**(5–6), 594–621 (2010)
2. Azoff, E.M.: *Neural Network Time Series Forecasting of Financial Markets*. Wiley, Hoboken (1994)
3. Bianco, V., Manca, O., Nardini, S.: Electricity consumption forecasting in Italy using linear regression models. *Energy* **34**(9), 1413–1421 (2009)
4. Box, G.E., Jenkins, G.M., Reinsel, G.C., Ljung, G.M.: *Time Series Analysis: Forecasting and Control*. Wiley, Hoboken (2015)
5. Cadenas, E., Rivera, W.: Wind speed forecasting in three different regions of Mexico, using a hybrid ARIMA-ANN model. *Renew. Energy* **35**(12), 2732–2738 (2010)
6. Carbonneau, R., Laframboise, K., Vahidov, R.: Application of machine learning techniques for supply chain demand forecasting. *Eur. J. Oper. Res.* **184**(3), 1140–1154 (2008)
7. Friedman, J.H.: Greedy function approximation: a gradient boosting machine. *Ann. Stat.* 1189–1232 (2001)
8. Hannan, E.J.: *Multiple Time Series*, vol. 38. Wiley, Hoboken (2009)
9. Jorgensen, B.: Exponential dispersion models. *J. Roy. Stat. Soc. Ser. B (Methodol.)* **49**, 127–162 (1987)
10. Kavasseri, R.G., Seetharaman, K.: Day-ahead wind speed forecasting using f-ARIMA models. *Renew. Energy* **34**(5), 1388–1393 (2009)
11. Kim, K.J.: Financial time series forecasting using support vector machines. *Neurocomputing* **55**(1–2), 307–319 (2003)
12. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. arXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014)
13. Lee, M.H., Rahman, N.H.A., Latif, M.T., Nor, M.E., Kamisan, N.A.B., et al.: Seasonal ARIMA for forecasting air pollution index: a case study. *Am. J. Appl. Sci.* **9**(4), 570–578 (2012)
14. Liao, T.W.: Clustering of time series data—a survey. *Pattern Recogn.* **38**(11), 1857–1874 (2005)
15. Liu, D., Niu, D., Wang, H., Fan, L.: Short-term wind speed forecasting using wavelet transform and support vector machines optimized by genetic algorithm. *Renew. Energy* **62**, 592–597 (2014)
16. Makridakis, S., Hibon, M.: The M3-competition: results, conclusions and implications. *Int. J. Forecast.* **16**(4), 451–476 (2000)
17. Pai, P.F., Lin, C.S.: A hybrid ARIMA and support vector machines model in stock price forecasting. *Omega* **33**(6), 497–505 (2005)
18. Park, D.C., El-Sharkawi, M., Marks, R., Atlas, L., Damborg, M.: Electric load forecasting using an artificial neural network. *IEEE Trans. Pow. Syst.* **6**(2), 442–449 (1991)
19. Pufnik, A., Kunovac, D., et al.: Short-term forecasting of inflation in Croatia with seasonal ARIMA processes. Technical report (2006)
20. Qi, M., Zhang, G.P.: Trend time-series modeling and forecasting with neural networks. *IEEE Trans. Neural Netw.* **19**(5), 808–816 (2008)
21. Sun, H., Liu, H., Xiao, H., He, R., Ran, B.: Use of local linear regression model for short-term traffic forecasting. *Transp. Res. Rec.: J. Transp. Res. Board* **1836**, 143–150 (2003)

22. Tweedie, M.: An index which distinguishes between some important exponential families. In: *Statistics: Applications and New Directions: Proceedings of Indian Statistical Institute Golden Jubilee International Conference*, pp. 579–604 (1984)
23. Yang, Y., Qian, W., Zou, H.: Insurance premium prediction via gradient tree-boosted Tweedie compound Poisson models. *J. Bus. Econ. Stat.* **36**, 1–15 (2018)
24. Yinga, X., Jungangb, C.: Traffic flow forecasting method based on gradient boosting decision tree (2017)
25. Zhang, G.P.: Time series forecasting using a hybrid arima and neural network model. *Neurocomputing* **50**, 159–175 (2003)
26. Zhou, J., et al.: KunPeng: parameter server based distributed learning systems and its applications in Alibaba and ant financial. In: *SIGKDD*, pp. 1693–1702 (2017)