

Predicting E-commerce Item Sales with Web Environment Temporal Background

Yihong Zhang¹ and Takahiro Hara¹

¹Osaka University, Japan

Abstract. In this paper, we study the effect of Web environment temporal background in predicting e-commerce item sales, especially those in temporary sales. Temporary sales nowadays are a popular strategy for quickly clearing inventories. For traditional recommender systems, predicting the sales of an item is done based on its past purchase records. For temporary sales items, however, such records are not available. In order to make recommendation for such items, contextual information, such as product descriptions, is usually used. We investigate whether temporal background in the Web environment can be additional useful contextual information in recommender systems. It is assumed that items consistent with the temporal background would have higher demands. We propose a method for representing the temporal background using word embeddings of e-commerce activities and social media data, and evaluate their effect on sales prediction. Through empirical analysis with real-world data, we found that temporal background does have positive effects for sales prediction. The findings in this paper can be conveniently incorporated into future recommender system designs.

Keywords: e-commerce, recommender, social media

1 Introduction

Predicting item sales is an important and challenging problem in e-commerce and marketing. Potentially, knowing the outcome of sales before putting the item on the shelf help sellers better manage inventories. And e-commerce websites can also use this prediction to make more accurate recommendations. This is particularly true for temporary sales, in some situations can also be called flash sales, for which the main purpose of the campaign is to clear a certain amount of inventories [DL18]. Running a successful temporary sales campaign will involve several considerations, including choosing the right product to offer, promoting ahead of time, and using the right word for campaign descriptions. Among them, the timing to start the campaign is of utter importance. It would be much easier to sell the item when the timing is right. For example, it is known that it is much easier to sell air conditioners in early June in Japan as talks about the summer holiday start to appear. Continuous information that reflects on such moments can be considered as the *temporal background*. In this paper, our aim are two folds. First, we would quantify temporal background in Web environment, representing it in a way that can be processed computationally. Then, we would investigate to what extent temporal background can influence item sales through empirical analysis.

We have a similar goal with recommender systems, which in recent years have attracted significant research efforts [SKKR01]. Typically, a recommender system suggests a ranking of available products that the user may purchase in the future. For items in temporary sales that have no previous record of being purchased, however, such recommendation systems based

on past transactions are not useful. This is known as the *cold start problem* [LVLD08]. The cold start problem provides a challenge to recommend new items to users, and the typical solution is to use the contextual information associated with the item or the user that are available before there is any transaction [LLK14]. In this paper, we are not proposing a solution to the cold start problem. Instead, we focus on temporal background as a factor that can potentially help solving the cold start problem. The outcome of our study can reveal to what extent temporal background can help predicting sales of items that have no previous sales records. If there is a clear influence, then future cold start recommender systems can incorporate temporal background as an additional contextual information.

To be discussed in detail in Section 3, we build representations of temporal backgrounds from two data sources, including purchases records of an e-commerce website, and text messages of a social media platform. This specific e-commerce website hosts exclusively temporary sales that are usually available for a period between 7 to 14 days, and thus can be easily influenced by the temporal background. We are provided by this website with all all purchase records that occurred during a period of one year. These purchase records will be used as both the target to be predicted and the data for building the temporal background. The prediction is thus done for the number of item purchases in this e-commerce website, based on the temporal background constructed from the same purchase records and the social media data. The temporal background built from purchases records has a local and closer association with the temporal aspects of products, while that from social media represents a broader environment that reflects the social interest of the moment.

To summarize, our main contribution with this paper are three folds. First, we propose a method to represent temporal background from e-commerce and social media data. Second, we propose a method to predict product sales based on temporal background. Third, using real-world datasets, we verify our approach and reveal the answer about to what extent temporal background can be used to predict item sales. In Section 3 we will introduce our method and in Section 4 we will present experimental results.

2 Related Work

In this paper, we propose the concept of temporal background that can be generated using e-commerce purchase records and social media data. While we consider this to be a new concept, there are a number of previous works already studied the predictive relationships between social media and product sales. Gruhl et al. for example, proposed to use online blogs to predict book sales performance, which was quantified as sale ranks, published by Amazon [GGK+05]. They first studied the correlation between sales rank and blog mentions. From selected top-ranked books, they extracted book titles and author names, as the queries to generate mention frequencies, then the correlations were calculated. Asur and Huberman proposed a similar approach to predict movie revenues from discussions on Twitter [AH10]. They selected a number of movies and extracted related tweets using keywords present in the movie title. Based on the correlation with the Hollywood Stock Exchange index, they studied two statistics in the tweets, URLs and retweets representing promotional material, and rate of tweet mentions. Zhang and Pennacchiotti conducted a study of predicting product sales on eBay from Facebook data [ZP13]. They used a database containing eBay users who connect their accounts to Facebook. Based on the fact that there are similar categories on eBay and Facebook, they found that there is a strong correlation between liking Facebook pages and product purchase of the same category. Their prediction fell in the setting of recommendation systems, instead of the separation of past and future, and they claimed that the use of social media data can solve the cold-start problem for recommendation. Pai and Liu proposed a method to predict vehicle sales from tweets and stock market values [PL18]. They collected tweets mentioning brand names and conducted sentiment analysis to find correlation between tweets and sales.

Lassen et al. conducted a study of predicting iPhone sales from tweets [LMV14]. They collected tweets containing the word iPhone and then primarily conducted sentiment analysis on these tweets. They used a linear model to find correlation between tweet sentiment and iPhone sales, which were divided into quarters. In these previous works, a common drawback is that they relied on keywords that can be associated with the product. While this is feasible for products such as books and movies, in many real-world e-commerce scenarios, such associations may not be present. Our work on the other hand generalizes product and social media data into embeddings, so that the prediction does not require keyword association.

There is also a number of works that predicts continuous sales such as stocks from a textual background. For example, Bollen et al. conducted a sentiment analysis of tweets with regard to changes in Dow Jones Industrial Average index [BMZ11]. They extracted sentiment expressions in tweets using a dictionary that has six mood categories. These tweets are not necessarily related to stock market. However, as they concluded, the collective sentiments could be indicators of stock market changes. Particularly, the mood "calm" seems to be a strong indicator of stock market changes in three or four days. Moat et al. attempted to use Wikipedia activities to predict stock market changes [MCA+13]. They found evidences of increases in the number of page views of articles relating to companies or other financial topics before stock market falls. These works, however, rely on the continuous association between the item in question and the background. In our case of temporary deals, such association may not be available. In this aspect, our method is more general in that it can be applied to the case with or without continuous association between specific item and temporal background.

3 Methodology

We aim to develop a method that predicts future item sales from temporal background. The overview of our method is shown in Fig. 1. Our method consists of two main parts. In the first part, we represent temporal background using embeddings, which are vectors of real numbers that can be processed computationally. In the second part, we make prediction of item sales by comparing the temporal background and the new item description, which are projected to a same embedding space. Evaluating prediction results can thus reveal to us to what extent temporal background can influence item sales. In this section, we will present our method in detail.

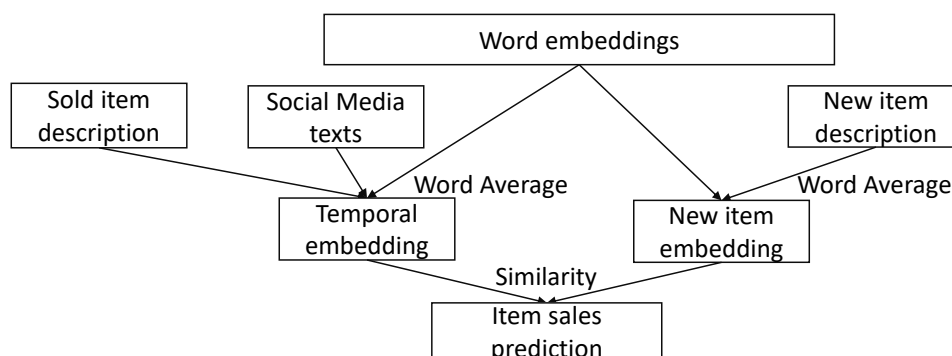


Figure 1. Overview of the method for predicting future item sales from temporal background

3.1 Representing Temporal Background

We found two sources that may contain information about the temporal environment. The first source is past records of e-commerce sales. As we have discussed in the introduction, products in our e-commerce website are temporary deals that are available for a short period of time. As such, the products sold in the past may not connect directly to the products currently on the

market. But the descriptions of products sold in the past may nevertheless contain information about selling aspects given a certain time period. The second source is social media. A social media platform such as Twitter posts millions of messages each day, which contain all kinds of topics that are worth public discussion. Although rarely connect directly to products in an e-commerce site, these messages certainly contain temporal information that reflects interesting aspects occur in the time period. Our first task at hand is to generate a temporal representation of these background data sources.

Given a day, when something interesting or stimulating happens, some topics in social media may become trending, and some products with certain aspects may see sudden rise in sales. When this happens, we say certain temporal aspects are *emerging*. We capture this emergence by observing the change of word frequency. Both the product description and social media messages are text data and can be represented as words, whose frequency we can count. The frequency of a product description word is counted as $\sum sales_i$ where $sales_i$ is the number of sales of product i whose description contains the word. The frequency of the social media words is simply the count of messages that contains the word. Thus we obtain the frequency table of product description and social media words.

We then devise a method for emergence detection based on word frequencies. Similar to some previous works on social media event detection [CALC13], our method involves a foreground and a background. Suppose the period for foreground is fp , and for background is bp , so that word frequencies in these periods are $F_{fp} = \{f_{t-fp}, \dots, f_{t-1}\}$ and $F_{bp} = \{f_{t-fp-bp}, \dots, f_{t-fp-1}\}$. We set inc_{fp} to *True*, if $f_{t-1} > \mu(F_{fp})$, where $\mu(\cdot)$ is the mean function, i.e., the frequency in the last day in the foreground period increases compared to the mean of foreground period, and *False* otherwise. Similarly we set inc_{bp} for the background period. Finally the emergence e_t of the word at time t is set as:

$$e_t = \{ 1, \text{ if } inc_{fp} \text{ OR } (inc_{bp} \text{ AND } \mu(F_{fp}) > \mu(F_{bp}))0, \text{ otherwise}$$

With this formula, we aim to capture two phases of surges of words in social media. First, inc_{fp} captures a new surge. Second, $inc_{bp} \text{ AND } \mu(F_{fp}) > \mu(F_{bp})$ captures the sustenance of a previous surge. Both phases can be considered as a part of an emergence. With this calculation, we obtain for each time unit the emerging words in product sales and social media.

From here, we can follow a naive approach and represent each time unit as bag-of-words, which are the emerging words. However, this approach does not consider the meaning of words, which may cause error, for example, when two words of similar meaning are counted separately. Considering this, we would like to generalize the words into meanings. More specifically, we use distributed representation of words, also called word embeddings, which nowadays are commonly used in text-based analysis. Based on an implementation made available online¹, we learn a set of Japanese word embeddings using Wikipedia. The result word embeddings have 50 dimensions², each represent a certain semantic aspect of a word. To represent a group of words, we take the average vector of embeddings for these words. To represent a time unit which consists of emerging words from two sources, we concatenate the group embedding of the product words and social media words, and as the result, we have a vector of 100 dimensions to represent the time unit. This is our representation of the temporal background.

3.2 Predicting Future Temporal Background from Past Data

We have shown how embeddings representing the temporal background can be generated from e-commerce and social media activities in the time unit. However, as a prediction problem, data for the current time unit is not known before hand, and embeddings of current time unit need to be somehow generated from past data. We can consider embeddings across time as a

¹<https://github.com/philipperemy/japanese-words-to-vectors>

²Although one can train word embeddings with different dimensions, we chose 50 to balance the complexity and generality

multivariate time series, and the task can be considered as a time series forecasting problem, for which many solutions have been proposed [WB04]. Since our focus is on the association between temporal background and item sales, we will only discuss two simple solutions here. The first solution is simply taking the embedding from the previous time unit as the prediction. This method can potentially work because major trends in e-commerce and social media only change gradually. There may be a problem, though, when some abnormal event happens and disrupts the continuity of embeddings across time.

The more popular method for forecasting time series nowadays, however, is through neural networks [LCYL18]. For our task, we build a simple recurrent neural network (RNN), which takes input of h embeddings from previous time units and outputs embedding of current day time unit. Between input and output there are two hidden layers, one contains 48 long short term memory (LSTM) nodes, and the other one contains 30 fully connected nodes. When training this neural network, predictions are iteratively compared with the target embedding, and the mean absolute error (MAE) is used to update neurons through back propagation. The LSTM layer essentially learns how a number of past values lead up to the current value in the time series. The fully connected layer is expected to capture interaction between two data sources, which cannot be captured by simply using the previous day values. We set h to 3 in our experiments, but different values such as 4 or 5 result in similar forecasting performance.

3.3 Evaluating the Effect of Temporal Background in Item Sales Prediction

After obtaining the temporal background representations, we need to find a connection between them and item sales. The popularity of an item will depend on many factors, including inherent quality, brand awareness, discount rate, and so on, and temporal background may be just one among them. Nevertheless, we consider this hypothesis:

Hypothesis (Temporal background consistency) *An item that is more consistent with the temporal background tends to have higher demand.*

As an example scenario, in Japan autumn is strongly associated with appetite. When people actively talk about food in autumn in social media, food products in e-commerce sites are so expected to have higher sales. Although temporal background cannot always be associated with products in this way, for example, when people in social media are talking about a recent political event, we argue that the consistency between product and temporal background can always to some degree influence the product sales.

To measure the consistency between product and temporal background, we apply the cosine similarity. Given a product embedding v_p and the temporal background embedding v_t , which are real value vectors, the consistency between them is calculated as:

$$c(v_p, v_t) = \text{cos_sim}(v_p, v_t) = \frac{\sum_i v_{pi}v_{ti}}{\sqrt{\sum_i v_{pi}^2} \sqrt{\sum_i v_{ti}^2}}$$

After quantifying the consistency between the item and temporal background we can compare the ranking based on it and the actual sales number ranking. There are several measurements we can take. One example is Recall@k, which is calculated as

$$\text{Recall@}k = \frac{TP@k}{TP@k + FN@k}$$

where $TP@k$ is the number of actual top items in the selected k suggestions (True Positives), and $FN@k$ is the number of actual top items not in the selected k suggestions (False Negatives). Recall@k tells the ability the prediction method has to find top items given a certain number of choices.

Another possible metric is *Average Precision* (AP). First we get Precision@k as

$$Precision@k = \frac{TP@k}{TP@k + FP@k}$$

where $FP@k$ is the number of items that are not actual top items, among k suggestions (False Positives). Then AP is calculated as

$$AP = \sum_{k=1}^K (Recall@k - Recall@(k-1)) * Precision@k$$

Essentially, a higher AP tells that the top items are more concentrated in the top suggestions by the method.

4 Experimental Analysis

4.1 Dataset Preparation

We obtain a product sales dataset from a Japanese e-commerce website. The products, called *deals* by the site, are discount coupons that are made available for a limited period of time, usually between 7 and 14 days. Customers who bought these deals can exchange them with real products. The products include several categories of items, including food, cosmetics, home appliances, hobby classes, travel packages, and so on. The dataset provided to us are of a period between October 2016 and August 2017. In total, there are 68,271 products made available and sold at least once during this period, and attracted about 1.6 million purchases. The number of available deals each day is about 1,000 on average. Each deal in the dataset is associated with a textual description written mostly in Japanese.

We obtain a social media dataset by collecting Japanese tweets through Twitter API³. To align with the period of e-commerce dataset, we develop a procedure to search past tweets without monitoring Stream API. In addition to time requirement, it is also desirable that the tweets are talking about Japanese domestic affairs, which reflects the background in which the e-commerce business was operated. Our procedure is thus as the following. First, we collect a list of Japanese politician Twitter accounts⁴. From them we remove a few top politician accounts such as Abe Shinzo as they would attract foreign followers. Next we collect the follower of these politicians, who are expected to be Japanese citizens. Then we select from these citizen accounts whose earliest tweets are dated earlier than October 1st, 2016. This is to ensure that the accounts are active during the entire period of e-commerce dataset. Finally, we collect tweets in the said period from these selected accounts. These tweets become our social media source of temporal background in this experimental analysis. In total this dataset contains about 1.7 million tweets from 11,673 accounts.

We use the natural language processing package *kuromoji*⁵ to process the Japanese text in the e-commerce and social media datasets. The package can effectively perform segmentation and part-of-speech (POS) tagging for Japanese text. After POS tagging, we select only nouns to represent the information in the text. These nouns are converted to temporal background embeddings following the method described in Section 3.

We use day as the time unit. Some of the components in our method such as the RNN model for embedding prediction require training data. We thus split our dataset into a training set and a testing set. The training set consists of 300 days of data, and the testing set consists of 20 days of data.

³<https://developer.twitter.com/en/docs>

⁴Such a list can be found online as political social media accounts are usually public. An example list is provided by the website Meyou with the url <https://meyou.jp/group/category/politician/>

⁵<https://github.com/atilika/kuromoji>

4.2 Direct Prediction of Sales from Product Description

The method proposed in this paper uses two steps to predict the product sales, first embeddings are generated for the product and temporal background, and then sales are predicted by comparing these embeddings. It is also possible, however, to learn a model that directly projects product description to sales, i.e., without the intermediate step of comparing it with temporal background. In this experimental analysis, we implement and test such a method. Using the training set described above, we train the model by setting the response variable as the daily sales number of the product, and the explanatory variable as the 50-dimension word embedding of the product description.

From here there are many possible machine learning techniques that can be applied, for example, linear regression or support vector machines. Since we expect non-linear relationship between dimensions in the word embedding and the sales number, we choose random forest (RF) as our learning model. In previous works, it has been shown that random forest can effectively predict product sales with token-based social media timing signals [ZHS20]. We train a random forest model using the training dataset, and then apply it to the testing dataset, by producing one sales number prediction for each product. The predicted sales numbers are ranked for each day in the test period and evaluated in the same way as we evaluate our method.

4.3 Embedding Prediction Accuracy

As discussed in the methodology section, we use two methods to predict the temporal background embedding of the current day using the embeddings in the past. First we use simply the embeddings of the previous day (P1), then we train a RNN model to forecast current embedding using embeddings of past h days. It would be interesting to see their prediction accuracy for the current embedding, which is used to predict product sales. After investigation, we found that the mean absolute error (MAE) for P1 method is 0.139, and for RNN method is 0.113. The root mean square error (RMSE) are 0.176 and 0.154 for P1 and RNN methods, respectively. Therefore, it is evident that RNN produces a forecast closer to the actual embeddings to be found for the current day.

4.4 Results and Discussions

We test different methods for predicting item sales ranks, and the accuracy measured as Recall@K and mean AP (mAP@K) is shown in Table 1. We test two K values of 50 and 100. The accuracy for the random method is based on theoretical values. For methods based on temporal background, we made two separate predictions, comparing the new item embedding first with the item part of the temporal background embedding, then with the tweet part. We tested three temporal background embeddings, namely, now, P1, and RNN. "Now" is taken as the current day embedding. It is not predicted and cannot be known before hand, but a comparison between it and predicted embeddings can be interesting. These results are averaged over the 20-day testing period. For each day, we pick top 20 items from all the items available for the day according to actual sales amounts, and then make 100 predictions. The theoretical $Recall@k$ for the random method is thus $100/a$ regardless of k , where a is the number of available items of the day.

Table 1. Average accuracy of prediction methods

	random	RF item	item now	tweet now	item P1	tweet P1	item RNN	tweet RNN
Recall@50	0.055	0.053	0.108	0.090	0.065	0.083	0.063	0.070
Recall@100	0.111	0.140	0.185	0.158	0.155	0.158	0.158	0.165
mAP@50	0.001	0.005	0.013	0.007	0.010	0.006	0.009	0.005
mAP@100	0.002	0.008	0.017	0.010	0.014	0.009	0.014	0.009

There are several insights we can draw from the results. First we look at the comparison with the random method. We can see that all prediction methods are better than the random method, which indicating that both item description and temporal background contain positive clues for predicting item sales. We can also see that using temporal background achieves better prediction accuracy than using the item description, indicating stronger predictiveness.

Second, comparing "now" embedding the predicted embeddings, we can see that using current day embedding achieves a higher accuracy. Even though it is not a prediction, we can see from it how correctly predicted temporal background can improve item sales accuracy. This also explains why RNN-predicted embedding is better than using previous day embedding. Since RNN predicted an embedding based on the embeddings in the last few days, it tends to predict a value that lays between the values in the previous day and the current day. As the result, its sales rank prediction accuracy also tends to lay between those using the previous day and the current day embeddings.

Last we compare between predictions using item and tweet embeddings. According to the result, when measuring Recall@k, tweet-based prediction is better than the item-based prediction. But when measuring AP, item-based prediction is better. It means that tweet-based prediction can generally find more top items, but item-based prediction can give higher rank to found items even though they are fewer. Similar tendencies are observed for both K values of 50 and 100.

4.5 Item Analysis

In order to get a closer view of what exactly happens within the prediction process, we analyze some concrete cases. We first pick the first day of test data and collect the emerging social media words most consistent with the temporal embedding of the day. Top 20 words collected and their cosine similarity scores with the temporal background are shown in Table 2⁶.

Table 2. Consistent emerging social media words on test day 1

entirety 0.708	event 0.695	this 0.688	everyone 0.687	answer 0.685
every time 0.671	treatment 0.667	national 0.658	everyday 0.655	target 0.644
hope 0.639	choice 0.638	opposition 0.634	procedure 0.632	expectation 0.632
report 0.629	purpose 0.627	only one 0.626	today 0.625	investigation 0.623

We can roughly guess that the trending social media topic of the day is about some national events and something that involves report and investigation. Next we pick some items in top positions of the rank predicted by RNN tweet method, which is shown to be the best prediction method. More specifically, we pick one true positive and one false positive items by comparing the predicted ranked and actual sales ranks. The true positive item is ranked 17th by prediction and 6th by actual sales. The false positive item is ranked 1st by prediction and 273rd by actual sales. The descriptions and words most consistent with the temporal embedding of the day are shown in Table 3.

Comparing item description words with social media words, we see that both items ranked high because their descriptions contain words related to trouble, reporting and investigation, which are trending semantics in social media. However, item 2 is a false positive mostly because other factors cause low sales for this item. From these examples, we can see how temporal background influences items sales predictions and its limitations.

⁶Words and sentences shown in this section are translated from Japanese by authors.

Table 3. Description and consistent words for selected items**item 1**

Description: [experience interview / July 21st new open commemoration / existing account OK] all hand 100% placenta concentrate undiluted introduction worries raised with age are approached. CURACION introduction 60 minutes. Highly anticipated pure placenta specialty store

interview 0.667	experience 0.641	existing 0.537	approach 0.479	quality 0.414
introduction 0.359	specialty 0.325	street 0.320	age 0.316	anticipated 0.300

item 2

Description: [second anniversary / specialty shop for food trouble/ better know more about prevention and care] 97.2% satisfaction with German foot care, corns and keratin on the soles of the feet. Painless care as a solution to your concern (original prevention care set menu)

solution 0.645	trouble 0.618	concern 0.593	satisfaction 0.585	menu 0.511
pain 0.490	prevention 0.482	set 0.426	original 0.380	German 0.347

5 Conclusion

Our aim with this paper is to discover the effect of Web environment temporal background in predicting e-commerce item sales. In particular, we would like to verify the hypothesis that items more consistent with the temporal background would have higher demands. For this purpose, we propose a method to generate embeddings for temporal backgrounds from e-commerce and social media activities, and make prediction of item sales based on them. By testing the accuracy of the predictions made using our method, and comparing it to the random baseline, we would be able to tell whether temporal background has positive effects on item sales prediction. Experimental analysis done using real-world data does show this positive effect. However, with item-level analysis, we can see some limitations of temporal background-based prediction. Initially this work is developed to support cold-start recommendation systems. Future works can be done on cleaning and filtering social media data so that its content can be more relevant to e-commerce items and potentially produce stronger positive effects.

Acknowledgement

Statement on competing interests: The authors declare that there is no conflict of interest. This research is partially supported by JST CREST Grant Number JPMJCR21F2.

- [AH10] S. Asur and B. A. Huberman, "Predicting the future with social media," in *Web Intelligence and Intelligent Agent Technology (WI-IAT)*, IEEE, vol. 1, 2010, pp. 492–499.
- [BMZ11] J. Bollen, H. Mao, and X. Zeng, "Twitter mood predicts the stock market," *Journal of computational science*, vol. 2, no. 1, pp. 1–8, 2011.
- [CALC13] Y. Chen, H. Amiri, Z. Li, and T.-S. Chua, "Emerging topic detection for organizations from microblogs," in *Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval*, ACM, 2013, pp. 43–52.
- [DL18] F. Dilmé and F. Li, "Revenue management without commitment: Dynamic pricing and periodic flash sales," *The Review of Economic Studies*, 2018.
- [GGK+05] D. Gruhl, R. Guha, R. Kumar, J. Novak, and A. Tomkins, "The predictive power of online chatter," in *Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining*, ACM, 2005, pp. 78–87.

- [LCYL18] G. Lai, W.-C. Chang, Y. Yang, and H. Liu, "Modeling long-and short-term temporal patterns with deep neural networks," in *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, 2018, pp. 95–104.
- [LLK14] G. Ling, M. R. Lyu, and I. King, "Ratings meet reviews, a combined approach to recommend," in *Proceedings of the 8th ACM Conference on Recommender systems*, ACM, 2014, pp. 105–112.
- [LMV14] N. B. Lassen, R. Madsen, and R. Vatrappu, "Predicting iphone sales from iphone tweets," in *Proceeding of the 18th International Enterprise Distributed Object Computing Conference*, IEEE, 2014, pp. 81–90.
- [LVLD08] X. N. Lam, T. Vu, T. D. Le, and A. D. Duong, "Addressing cold-start problem in recommendation systems," in *Proceedings of the 2nd international conference on Ubiquitous information management and communication*, 2008, pp. 208–211.
- [MCA+13] H. S. Moat, C. Curme, A. Avakian, D. Y. Kenett, H. E. Stanley, and T. Preis, "Quantifying wikipedia usage patterns before stock market moves," *Scientific reports*, vol. 3, p. 1801, 2013.
- [PL18] P.-F. Pai and C.-H. Liu, "Predicting vehicle sales by sentiment analysis of twitter data and stock market values," *IEEE Access*, vol. 6, pp. 57 655–57 662, 2018.
- [SKKR01] B. Sarwar, G. Karypis, J. Konstan, and J. Riedl, "Item-based collaborative filtering recommendation algorithms," in *Proceedings of the 10th International Conference on World Wide Web*, ACM, 2001, pp. 285–295.
- [WB04] Z. Wang and D. A. Bessler, "Forecasting performance of multivariate time series models with full and reduced rank: An empirical examination," *International Journal of Forecasting*, vol. 20, no. 4, pp. 683–695, 2004.
- [ZHS20] Y. Zhang, T. Hara, and M. Shirakawa, "Discovering social media timing signals for predicting temporary deal success.," in *Proceedings of the 28th European Conference on Information Systems*, 2020.
- [ZP13] Y. Zhang and M. Pennacchiotti, "Predicting purchase behaviors from social media," in *Proceedings of the 22nd international conference on World Wide Web*, ACM, 2013, pp. 1521–1532.