

Social Media Brand Engagement as a Proxy for E-commerce Activities: A Case Study of Sina Weibo and JD

Weiqliang Lin*, Pedro Saleiro†, Natasa Milic-Frayling‡, Eugene Ch’ng§

*International Doctoral Innovation Center, University of Nottingham, Ningbo, China

†Department of Computer Science, University of Chicago, Chicago, USA

‡School of Computer Science, University of Nottingham, Nottingham, UK

§NVIDIA Joint-Lab on Mixed Reality, University of Nottingham, Ningbo, China

{wyatt.lin, eugene.chng}@nottingham.edu.cn, Natasa.Milic-Frayling@nottingham.ac.uk, saleiro@uchicago.edu

Abstract—E-commerce platforms facilitate sales of products while product vendors engage in Social Media Activities (SMA) to drive E-commerce Platform Activities (EPA) of consumers, enticing them to search, browse and buy products. The frequency and timing of SMA are expected to affect levels of EPA, increasing the number of brand related queries, clickthrough, and purchase orders. This paper applies cross-sectional data analysis to explore such beliefs and demonstrates weak-to-moderate correlations between daily SMA and EPA volumes. Further correlation analysis, using 30-day rolling windows, shows a high variability in correlation of SMA-EPA pairs and calls into question the predictive potential of SMA in relation to EPA. Considering the moderate correlation of selected SMA and EPA pairs (e.g., Post-Orders), we investigate whether SMA features can predict changes in the EPA levels, instead of precise EPA daily volumes. We define such levels in terms of EPA distribution quantiles (2, 3, and 5 levels) over training data. We formulate the EPA quantile predictions as a multi-class categorization problem. The experiments with Random Forest and Logistic Regression show a varied success, performing better than random for the top quantiles of purchase orders and for the lowest quantile of search and clickthrough activities. Similar results are obtained when predicting multi-day cumulative EPA levels (1, 3, and 7 days). Our results have considerable practical implications but, most importantly, urge the common beliefs to be re-examined, seeking a stronger evidence of SMA effects on EPA.

Index Terms—Social media activities, e-commerce platform activities, cross-sectional data analysis, time series, multi-class categorization, quantile level prediction.

I. INTRODUCTION

E-commerce platforms enable promotion and sales of products at large scales and resort to careful monitoring and optimization of product sales cycle in order to ensure quality services to both sellers and buyers. Thus, a substantial effort has been put into studying clickstreams and predicting product purchases [1]–[5]. Clickstream data typically includes searching and browsing for specific brands or products and clicks on product information for feature and price comparison. Patterns of such user activities have been successfully used to predict users’ intent to purchase once engaged with the platform [3]–[7]. Indeed, searching, browsing, and comparing products are aligned with the Purchase Decision Models (PDM) used in marketing and sales management [8] and therefore, good pre-

dictors of consumer purchases. However, vendors are equally interested in creating awareness of their brands and products to drive traffic to their Web sites and e-commerce platforms. This

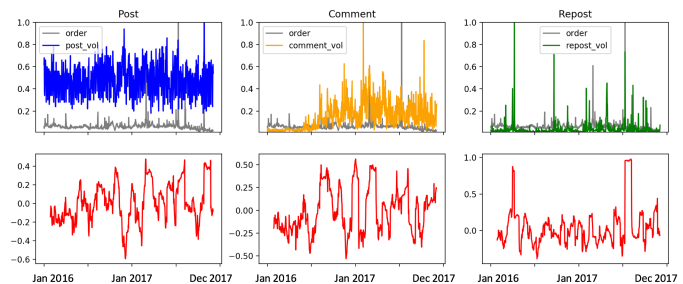


Fig. 1: Rolling correlation between SMA on day t_p and EPA on day t_{p+1} , with a 30-day rolling window over two years.

has led to increased commitment to Social Media Monitoring (SMM) by vendors and, consequently, to the development of tools and services for marketers to manage their Social Media Activities (SMA). The increased investment in SMM and SMA is driven by an underlying belief that social media is a key factor in business success, particularly for product promotion and sales. However, considering the global adoption of e-commerce platforms, it is important to examine that assumption and investigate to what degree vendors’ SMA affects consumer engagement with their products, i.e., the consumers’ E-commerce Platform Activities (EPA). In fact, we expect that the effects of vendors’ SMA are conflated with and dominated by advertising campaigns organized by the e-commerce platform itself and global retail events, such as Black Friday, 11-11, and June 18 in China. At the same time, EPA analyses and predictions of SMA effects are out of reach of individual vendors and marketers as they require comprehensive data collections and analysis and sophisticated machine learning methods.

Our research involves a cross-sectional data analysis of multiple time-series from SMA and EPA and a new approach to predicting EPAs, formulated as a multi-class categorization problem with classes defined by discretized time-series spec-

tra. More precisely, we specify distinct levels of EPA activities that correspond to quantiles of the EPA data distribution over a given time period. We use a data set provided by JD e-commerce platform (JD.com) that comprises search, click-through and product orders for a sample of 33 vendors from five product categories. For each vendor in the sample we collect social media posts, reposts, and commentaries published on the Sina Weibo (Weibo.com) social media platform. The data covers a period of two years. We use time-alignment of vendors' social media posts, reposts, and comments with EPA that correspond to different decision stages in the consumers' purchases: search, clickthrough, and orders. By relating the vendors' SMA and EPA we postulate:

Should the data analysis reveal high correlation between a vendor's daily SMA and its customers' daily EPA, we might be able to create reliable predictors of EPA based on SMA features.

Our correlation analyses of specific SMA and EPA pairs, i.e., Pearson correlation of Post with Search, Clickthrough with Order, and Repost with Comment data for different lag factors (1-15 day shifts), point to a highly variable correlation over time, as illustrated in Figure 1. Across the sample of 33 vendors, the correlation coefficients of SME-EPA pairs vary between 0.08 and 0.56, indicating that predicting daily volumes of EPA purely based on SMA would not be reliable. Nevertheless, the question still remains:

Could SMA features be used to predict changes in the levels of EPA, i.e., highs and lows in search, clickthrough and orders, especially if SMA has a sustained effect over several days, e.g., 3 to 7 days following a specific social media activity?

In order to explore this, we create predictive models for specific levels of cumulative search, clickthrough, and orders added up over 1, 3, and 7 days, where levels are defined by the quantiles of the corresponding distributions over the training time period. We apply Random Forest and Logistic Regression classifiers for categories corresponding to 2, 3, and 5 quantiles, i.e., (i) above the median, (ii) within three quantiles: (0-33%, 33-66%, 66-100%), and (iii) within five quantiles (0-20%, 20-40%, 40-60%, 60-80%, and 80-100%).

Our experiments show that SMA can be used to predict top quantiles of product orders and low quantiles of search and clickthrough across vendors' categories. The performance for the remaining ranges is on par with random predictions. The results are important for optimizing the e-commerce platform operations for peaks in EPA and for evaluating marketing campaigns through SMA. From the research perspective, our work is the first to provide a detailed analysis of SMA-EPA pairs and a method that reveals selective roles that SMA types play in predicting levels of specific EPA types.

II. RELATED WORK

Most of the existing work in understanding purchase behavior and user intent aims at estimating conversion rates for individual products, individual customers, or both [1], [3]–[7], [9]. Computational studies of consumers behavior have adopted ad-hoc approaches and brute-force feature engineering, aiming to uncover factors that explain user behavior, particularly the

intent to buy [1], [2], [5]. In contrast, studies within business and market research have led to conceptual frameworks that are useful for both optimizing and interpreting experiment results, particularly the effects of features used in predictive models [4], [6].

Considering e-commerce platforms, research has focused on behavioral patterns, product characteristics [3], [7] and feedback in product reviews [10], and use them as features for predicting the online product sale. Such predictive features often originate from resources internal to the shopping platform [3], [11], or those collected externally such as social media content [12] or query logs from online search engines [13].

With regards to social media, past studies explored how commercial intent, expressed in tweets, correlates with external events [14], [15]. Homophily was found to have a significant influence on the purchase intent, i.e., users are more likely to purchase products similar to those purchased by their friends due to word-of-mouth effects [16]. This phenomenon is important for e-commerce and has been studied in detail over the past decade [17], [18]. It has also been recognized that other external factors, such as geographic proximity, may explain similarity in purchase patterns among friends on social media [19], [20]. Furthermore, text analysis has been applied to detect purchase intent from social media content [21] and Facebook profiles [22] and provide recommendations based on micro-blogging activities [23].

We complement this work by considering social media monitoring practices and macro-level analyses that do not involve individual user profiles and content analysis but focus on the levels and timing of brand-specific social engagements, i.e., volumes of vendors' post and the followers commentaries and repost. Understanding temporal patterns of user activities has generated a wide interest over the past decade, including the identification of broad classes of temporal patterns based on activity peaks [24]–[27]. Usually, those classes are defined based on specific volumes ranges and duration of activities before and after the peak. Crane and Sornette [24] define endogenous and exogenous origins of peaks based on triggers generated by internal aspects of the social network. They found that popularity is mostly driven by exogenous factors instead of endemic spreading. Yang and Leskovec [27] propose a new measure of time series similarity and clustering.

In our research we apply cross-sectional analysis of data from two distinct platforms rather than studying specific aspects of commercial intent within social media or e-commerce platforms themselves. With the aim to explore the two sets of temporal data, we apply correlation analysis and formulate discretized prediction problems that can leverage SMA features to predict certain levels of e-commerce activities despite a highly volatile correlation of the corresponding data series.

III. PRELIMINARIES AND DATA ANALYSIS

E-commerce platforms are instrumented to capture information about shopping activities and gather detailed statistics about consumers' interactions with facilities that support the purchase workflow. Analyzing such data is instrumental for

optimizing the platform operations. Our research is conducted in collaboration with JD¹ with access to anonymized data about product sales of vendors in different product categories. Furthermore, we include information about vendors’ social engagements on Sina Weibo² platform that hosts Web sites of many JD vendors. The vendors share post and engage the community through comment and discussions. We align Sina Weibo data and JD logs to understand how the vendors’ Social Media Activities (SMA) on Weibo relate to the E-commerce Platform Activities (EPA) of JD users, focusing on patterns of interaction within both services. We perform a macro-level analysis considering time-series of actions; analyses of individual users’ behavior or content shared in social media are not part of this study.

TABLE I: Total social media activities of the sampled vendors within each product category for the two year period (Jan 2016-Dec 2017).

Category	Number of Vendors	Post	Comment	Repost
Phone&EI	4	37,312	5,143,200	2,629,674
Sports	5	6,927	295,952	230,733
Food	9	17,666	1,072,441	496,664
Clothes	6	9,456	227,258	155,334
Home	9	53,857	1,902,164	681,071




Fig. 2: Total Post, Comment and Repost statistics normalized by the corresponding maxima over the two year period (Jan 2016-Dec 2017).

A. Data

Our EPA data set comprises JD records of (1) Search for brands, (2) Clickthrough, and (3) Orders. Search queries are included only if explicitly mention a vendor’s name. The data was collected for a sample of 50 *most popular* JD vendors across 5 product categories: Phone&Electronics, Sports, Food, Clothes and Home on JD’s platform. The vendor’s popularity ranking is based on the JD’s sales performance metrics that is treated as business confidential. Among 50 companies we identified 33 that have had a Weibo account for at least two

¹JD.com is one of the largest e-commerce platforms in China.

²Weibo.com is the largest social media platform in China, with social media features that combine those of Facebook (www.facebook.com) and Twitter (www.twitter.com).

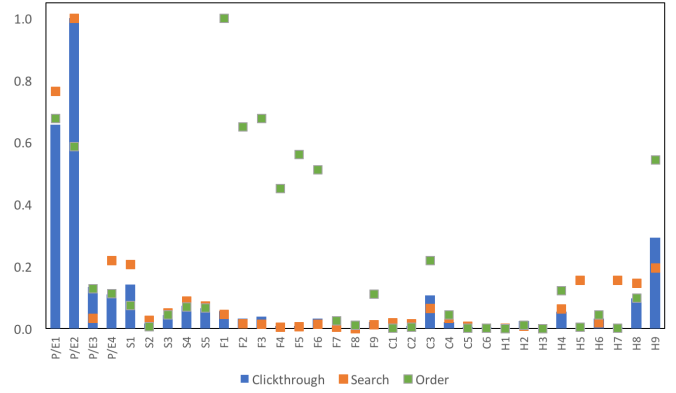


Fig. 3: Total Search, Clickthrough and Order volumes of individual vendors normalized by the daily maxima over the two year period (Jan 2016-Dec 2017).

years, from Jan 2016 until Dec 2017, and published at least 10 posts over that period. The final sample comprises: 4 vendors in Phone&Electronics (P/E), 5 in Sports, 9 in Foods, 6 in Clothes, and 9 in Home. For each vendor we collected: (1) Post, (2) Repost, and (3) Comment statistics.

The EPA and SMA data is segmented per calendar date. Table I shows the SMA distribution per vendor category and Figure 2 plots the total SMA volumes of Post, Repost and Comment activities for individual vendors on a logarithmic scale. As expected, the Repost and Comment statistics are two orders of magnitude higher than Post, illustrating a strong social media effect in response to the vendors’ posting activities. Considering the SMA statistics across five categories (Figure 2), the vendors in the Phone&Electronics category appear to gain most traction through repost and commenting activities. Sports category ranks lowest on posts while the Clothes category is the lowest when considering the sum of all three SMA statistics.

Figure 3 presents the total Search, Comment and Order statistics for individual vendors, normalized by the maximum daily volume over the period of two years. It shows that online purchases incur Search and Clickthrough volumes with 1-2 orders of magnitude higher than Order volumes as they play important parts in the purchase decision process.

B. Correlation Analysis

We calculate the pairwise Pearson Correlation between SMA (Post, Repost, Comment) and EPA (Search, Clickthrough, Order) for the two year time-series corresponding to the individual vendors. Table II shows the highest correlation coefficient for each SMA-EPA pair among all 33 vendors. It confirms low-to-medium correlations of SMA-EPA across vendors, with Post-Search having the highest maximum coefficient of 0.56. The maximum of 0.11 for Repost-Clickthrough and 0.08 for Repost-Orders indicate low correlations between these SMA and EPA types across all the vendors.

Heat maps in Figure 4 present the pairwise SMA-EPA correlation coefficients for each of the 33 vendors across the full two year period by considering the next-day EPA statistics for




Fig. 4: Pearson Correlation between social media activities (Post, Repost, and Comment) and specific e-commerce platform activities (Search, Clickthrough and Order) for 33 vendors using their two year time series. Heat maps show low to medium correlations between SMA and EPA pairs.

TABLE II: The highest correlation coefficients across SMA-EPA pairs for the sample of 33 vendors.

	Post	Comment	Repost
Clickthrough	0.39	0.23	0.11
Search	0.56	0.20	0.23
Order	0.25	0.21	0.08

a given SMA day statistics. Comparing the relative importance of SMA types, it appears that the Post statistics are more highly correlated with Search and Clickthrough, driving the brand awareness and product interest. Considering Comment statistics across vendors, they are more highly correlated with Clickthrough and Orders than with Search, suggesting that comments on posts and reposts may help with purchase decisions. Our analysis is the first to offer empirical evidence that different types of social media engagements relate to different aspects of online shopping activities.

We further consider SMA-EPA correlations for different time frames and time lags.

1) *Rolling Correlation Analysis*: We calculate the Pearson Correlation Coefficient for SME-EPA pairs using the rolling 30-day window over the period of two years (instead of the single two-year time-series) and show that the correlation coefficients for daily statistics varies over time. This is illustrated in Figure 1 (Section I), showing time variations of the correlation coefficients for the vendor with the highest volume of comments. Without stable correlations it would be hard, for example, to use Comment features to predict Orders.

2) *SMA-EPA Lag Analysis*: In our scenario, SMA and EPA occur on different platforms, i.e., Sina Weibo and JD, respectively. Thus, even if SMA has an effect on EPA, one can expect delays, depending on the speed of social media propagation. For that reason, we investigate rolling SMA-EPA correlations with different day lags, allowing for 1-15 day delay. By repeating the 30-day rolling window calculations with 1-15 day shift in time series, we observe a weekly pattern in the rolling correlation, with a spike every 7 days. However, even for the vendor with the highest volume of Comment

activities, the absolute correlation values are always below 0.5. Generally, the volumes of Post activities have the highest correlation with the volumes of Search or Clickthrough, while Comment data has the highest correlation with Order.

C. Problem Formulation

Concluding from the presented analysis, the task of predicting daily volumes of e-commerce activities from levels of social media activities is not well supported by the correlation analysis of daily statistics. However, in many practical scenarios it is sufficient to predict a trend, i.e., a change in the volume range, anticipating highs and lows that may affect product supplies and platform logistics. Furthermore, instead of daily activities, it is sufficient to predict purchase outcomes over a given period of time, e.g., a cumulative sales for 3-7 ahead. Thus, in Section IV we formulate the EPA prediction problem by considering (a) discretized distributions of EPA volumes for each vendor into 2, 3, and 5 quantiles and (b) quantile predictions of EPA totals for the next day, next 3 days and next 7 days.

IV. PREDICTING E-COMMERCE ACTIVITIES

Let us now consider a set of vendors, represented by their brands $B = \{b_1, b_2, \dots, b_i, \dots\}$. Without a loss of generality, we can assume that each vendor is represented by a single brand that has a specific daily stream of social media signals $S_t = \{s_1, s_2, \dots, s_j, \dots\}$, consisting of official Post activities by the vendor and the Repost and Comment activities by the users of the social media platform. Similarly, each brand has a daily stream of e-commerce platform activities (EPA) $E_t = \{e_1, e_2, \dots, e_l, \dots\}$ corresponding to Search, Clickthrough and Order actions on a given day t . Some of the social media users may also be customers of the e-commerce platform but we do not attempt to align the user activities across the two platforms.

We consider the predictive power of the social media signals regarding specific e-commerce platform activities, i.e., $f_v(b_i, t, S_t, E_t)$, a discrete volume function of a brand b_i ,

a timestamp t , a social media stream S_t and a type of e-commerce activity stream E_t . We are, in fact, interested in aggregated volumes of specific EPA types and consider a time frame $T = [t_p, t_{p+h}]$, where the date t_p is the day of prediction and t_{p+h} is the prediction horizon $h = 1, 3, 7$ (1-day, 3-day, 7-day ahead). Therefore, we aim to learn a proxy function f_v that at t_p defines a cumulative function:

$$F_V(b_i, T, S, E) = \sum_{t=t_p}^{t=t_{p+h}} f_v(b_i, t, S_t, E_t)$$

integrating, i.e., summing up f_v over T .

A. Multi-class Predictions

Instead of predicting specific values of $F_V(b_i, S, T, E_t)$ we predict the distribution of F_V in terms of quantile levels that F_V may attain on a given day or a given period of time. Evaluating the performance of quantile predictors enables us to assess whether different social media signals are predictive of specific EPA levels. We cast that as a multi-class categorization problem using supervised learning where the number of quantiles q determines the number of classes. Given a number of quantiles q and a training set with values $F_V(b_i, S, T, E_t)$ from the training time period, we determine the data points $F_V(b_i, S, T, E_t)$ within each quantile, i.e.,

$$Pr(F_V(b_i, S, T, E_t) \leq f_k) = k/q, \quad \forall k \in [1, q - 1]$$

and the corresponding ranges of F_V values. The minimum/maximum values determine the thresholds for assigning a class label to an instance of F_V .

Training of the multi-class predictor is based on partitioning the training set into distribution quantiles with corresponding class labels and value ranges. We use the quantile ranges to assign class labels to F_V values of the test set. If the test distribution is broader than the one of the training set, F_V values below the minimum or above the maximum are assigned to the lowest/highest quantile, respectively.

Given a time of prediction t_p , we extract features from the social media stream S up to t_p and predict F_V for the prediction horizon t_{p+h} . We consider activities on a daily basis and use t_p as 23:59:59 on each day.

B. Feature Construction

Our feature set is based on the three types of SMA signals: Post, Repost and Comment activities. For each signal type we generate features by calculating statistics from SMA streams for the time spans of a pre-specified length. In particular, we consider K-day periods with $K \in \{1, 3, 5, 7\}$. Table III shows the statistics that we calculate for each signal type, i.e., Post, Repost, Comment, over the specific stream length K . The Theil-Sen estimator [28] is a non-parametric trend detector that corresponds to the median over all possible combinations of slopes over K daily measures of the SMA activity. In total, we construct 22 features for each of the 3 SMA types; the total of 66 features that characterize a vendor's SMA.

TABLE III: Description of statistics used as features for each SMA type over K days before t_p . For $K = 1$ we consider only $Feat_{sum}$.

Feature	Description
$Feat_{sum}$	Sum of activity volume over K-day period
$Feat_{mean}$	Mean activity volume over K-day period
$Feat_{max}$	Maximum activity volume over K-day period
$Feat_{min}$	Minimum activity volume over K-day period
$Feat_{var}$	Variance of activity volume over K-day period
$Feat_{stdev}$	Standard deviation of activity volume over K-day period
$Feat_{prev}$	Activity volume on previous day
$Feat_{theil}$	Theil-Sen estimator

V. EXPERIMENTS

We use Logistic Regression and Random Forest [29]–[31] to learn multi-class classification models for 2, 3, and 5 quantiles (2-q, 3-q, 5-q) and vary the prediction horizon by modeling next day (1-day), 3-day and 7-day cumulative volumes of Search, Clickthrough and Orders. We evaluate multi-class prediction results by calculating standard precision, recall and F1 statistics. However, we focus our discussion on the precision since the aim is to assess the effectiveness of SMA features in predicting volume levels of specific EPA types. Thus, identifying *true-positive* instances within a specific quantile is given a priority over avoiding *false-negative* instances. In fact, correct predictions of *high* (top 20%-30%) and *low* (bottom 20%-30%) quantiles are of particular interest since their detection can improve SMA campaigns and optimize e-commerce operations for high levels of activities that on some days, such as global shopping events, can increase 100-fold.






Fig. 5: Sliding window of a 12-month training and 1-month test data with a shift of one calendar month, to cover the test period from Jan 2017 until Dec 2017.

1) *Temporal Cross-validation*: All our experiments are performed using 12-fold time-series cross-validation, as shown in Figure 5, with data sets specific to each vendor. We train and test our models using a two year JD data set (Section III-A). Our starting training set covers the 12 month period from 1 January 2016 to 31 December 2016. The test set is the following month. In each fold, we slide the 12-month training and 1-month test period by one calendar month. Thus, for each experiment we use one year of historical data and predict EPA in every calendar month in 2017. For each vendor we report the average precision statistics across 12-folds. We collate and average precision statistics across 33 vendors and across vendor categories.

2) *Activity Predictions*: For a given quantile scale (2-q, 3-q, 5-q) we predict volumes of activities for the next day t_{p+1} (1-



(a) Precision statistics for 3-q next-day predictions for Orders



(b) Precision statistics for 5-q next-day predictions for Orders

Fig. 6: Precision of Random Forest predictors for 3-q and 5-q categories for the next-day Orders across 33 vendors. The top quantile precisions are higher than random predictions (33% for 3-q and 20% for 5-q).

TABLE IV: Precision statistics for Random Forest 5-q classifiers for the next-day volumes of each EPA type: Order, Clickthrough and Search. Results are aggregated per five categories.

Vendor		0-20%			20%-40%			40%-60%			60%-80%			80%-100%		
		AVG	MAX	MIN	AVG	MAX	MIN	AVG	MAX	MIN	AVG	MAX	MIN	AVG	MAX	MIN
Order	P&E	0.159	0.338	0.013	0.136	0.218	0.054	0.198	0.292	0.070	0.298	0.570	0.149	0.341	0.393	0.271
	Sports	0.129	0.340	0.000	0.161	0.370	0.045	0.154	0.263	0.021	0.247	0.425	0.149	0.456	0.653	0.154
	Food	0.070	0.231	0.000	0.090	0.302	0.000	0.144	0.271	0.030	0.313	0.462	0.181	0.410	0.708	0.213
	Clothes	0.199	0.571	0.000	0.129	0.217	0.011	0.186	0.385	0.065	0.225	0.464	0.018	0.308	0.528	0.076
	Home	0.167	0.596	0.000	0.157	0.317	0.000	0.263	0.517	0.092	0.222	0.333	0.091	0.312	0.481	0.079
Clickthrough	P&E	0.371	0.548	0.192	0.200	0.311	0.099	0.147	0.241	0.025	0.159	0.353	0.000	0.265	0.565	0.055
	Sports	0.462	0.631	0.275	0.264	0.393	0.094	0.144	0.328	0.038	0.135	0.348	0.024	0.148	0.286	0.037
	Food	0.351	0.860	0.083	0.223	0.488	0.056	0.177	0.345	0.030	0.136	0.290	0.000	0.162	0.385	0.036
	Clothes	0.312	0.571	0.013	0.165	0.260	0.030	0.199	0.426	0.012	0.145	0.308	0.000	0.192	0.500	0.023
	Home	0.324	0.645	0.100	0.225	0.347	0.116	0.188	0.385	0.025	0.149	0.259	0.039	0.174	0.365	0.000
Search	P&E	0.521	0.679	0.336	0.231	0.353	0.065	0.109	0.293	0.029	0.077	0.176	0.000	0.255	0.481	0.070
	Sports	0.483	0.763	0.276	0.248	0.350	0.103	0.130	0.222	0.000	0.081	0.173	0.000	0.126	0.250	0.066
	Food	0.308	0.784	0.089	0.193	0.370	0.027	0.183	0.324	0.027	0.175	0.373	0.052	0.185	0.317	0.048
	Clothes	0.329	0.628	0.194	0.251	0.353	0.114	0.232	0.313	0.131	0.127	0.238	0.020	0.119	0.288	0.043
	Home	0.351	0.573	0.179	0.232	0.323	0.041	0.200	0.322	0.083	0.200	0.311	0.098	0.194	0.316	0.086

day) and the multi-day cumulative activities over 3-day and 7-day periods. For each prediction type and each individual vendor, we use the corresponding quantile scale determined over the training data. Our experiments thus involve 891 predictions: for each of 33 vendors, 3 EPA types (Search, Clickthrough, Orders) with three quantile scales (2-q, 3-q and 5-q) and 3 time periods (1-day, 3-day, 7-day). For this discussion we select experiments that shed light on:

- How successful the predictors are in identifying quantiles for individual EPA types across the vendor sample?
- How well the predictors perform for the cumulative 3-day and 7-day activities across the vendor sample?
- Which features significantly contribute to the performance of predictors for the individual EPA type?

A. Next-day EPA Predictions for 3 and 5 Quantiles

Experiments with the next day predictions of EPA for 3-q and 5-q show that the Random Forest (RF) predictors perform higher than random for Orders in the top quantiles (top 33% and 20%, respectively). Figure 6 presents RF precision for Order quantiles for individual vendors. Table IV summarizes RF precision statistics for all three EPA types and the five vendor categories. We highlight the RF results that are better than random predictions (> 0.20 for 5-q). We see that, in addition to the top quantile predictions for Orders, the precision statistics are better than random for Search and Clickthrough in the lowest quantile (bottom 33% for 3-q and

bottom 20% for 5-q). All other quantile predictions for Order, Clickthrough, and Search are on par or lower than random.

We conducted the same experiments with Logistic Regression (LR) and found that RF outperforms LR predictors. In fact, for our sample of product vendors, the t-test performed for LR and RF predictions yield $p < 0.05$ for all the quantile labels (not just the top quantiles).

B. Predictions of Multi-day Cumulative EPA


Considering a selective success of SMA based predictors for quantiles of the next-day EPA, we train predictors for cumulative EPA, i.e., EPA volumes over 3 and 7 days. We expect that multi-day cumulative statistics are less volatile and therefore the quantile levels may be more stable and predictable over time.

Our experiments show that RF predictors for 3-day and 7-day cumulative EPA are consistently performing at the similar order of magnitude for a given quantile level. That is illustrated in Table V, showing the precision statistics of RF predictors for 1-day, 3-day and 7-day cumulative volumes of Order, Clickthrough and Search for 3 quantiles (3-q). Similar observations are made for 5 quantiles and for the LR classifier.

We conclude that our high performing predictors for cumulative multi-day EPA volumes can be flexibly applied to different scenarios that benefit from observing cumulative EPA across multiple days. This includes the precision of top quantile volumes for Orders. Figure 7 shows 3-q predictions

TABLE V: Three quantile prediction results for five categories of products across 33 vendors based on Random Forest, reporting average precision by categories of 1-day (1D), 3-day (3D) and 7-day (7D) cumulative Orders, Clickthrough and Search for each quantile.

Vendor		0-33%			33%-66%			66%-100%		
		1D	3D	7D	1D	3D	7D	1D	3D	7D
Order	P/E	0.232	0.225	0.237	0.327	0.329	0.303	0.522	0.536	0.551
	Sports	0.241	0.212	0.172	0.237	0.240	0.298	0.584	0.598	0.607
	Food	0.141	0.139	0.133	0.288	0.230	0.251	0.663	0.656	0.675
	Clothes	0.287	0.286	0.280	0.297	0.272	0.310	0.397	0.431	0.478
	Home	0.261	0.266	0.244	0.375	0.383	0.351	0.453	0.467	0.464
	AVG	0.226	0.222	0.209	0.310	0.293	0.302	0.528	0.540	0.556
Clickthrough	P/E	0.488	0.480	0.471	0.269	0.303	0.232	0.385	0.362	0.372
	Sports	0.637	0.620	0.623	0.226	0.263	0.214	0.232	0.223	0.266
	Food	0.464	0.471	0.476	0.282	0.285	0.286	0.263	0.266	0.285
	Clothes	0.418	0.399	0.423	0.301	0.311	0.293	0.331	0.328	0.329
	Home	0.475	0.487	0.484	0.327	0.307	0.289	0.285	0.278	0.256
	AVG	0.488	0.486	0.490	0.288	0.295	0.271	0.291	0.286	0.293
Search	P/E	0.597	0.621	0.655	0.182	0.194	0.150	0.277	0.259	0.289
	Sports	0.621	0.613	0.616	0.266	0.248	0.263	0.166	0.157	0.152
	Food	0.442	0.413	0.440	0.316	0.319	0.292	0.299	0.291	0.276
	Clothes	0.455	0.479	0.476	0.375	0.370	0.355	0.225	0.205	0.172
	Home	0.453	0.474	0.524	0.360	0.313	0.298	0.330	0.327	0.326
	AVG	0.493	0.497	0.522	0.315	0.301	0.283	0.271	0.261	0.254



(a) Precision of bottom quantile predictions for Orders (b) Precision of middle quantile predictions for Orders (c) Precision of top quantile prediction for Orders

Fig. 7: Precision of Random Forest predictions into three quantiles for 1-day, 3-day and 7-day cumulative Orders.

of 1-day, 3-day and 7-day cumulative Orders for individual vendors.

C. Feature Significance

Both Random Forest and Logistic Regression allow us to assess the significance of individual feature types in terms of their contributions to the prediction decision. For illustration, we present the analysis of features for the RF classifier trained for 3-q predictions of the next-day EPA volumes. For each classifier we calculate the average relative rank of a feature. Table VI shows the top 10 contributing features of 1-day EPA predictors for Orders, Clickthrough and Search, respectively, across all the vendors and 3 quantiles.

We observe that the relative feature importance is, to a degree, in agreement with the observations from the correlation analysis in Section V-2. Search levels are predicted by Comment and Repost activities considered over different lengths of time: 1, 3, 5, and 7 days. Clickthrough activities seem to be aligned with SMA over 7 and 3 day periods, primarily with consumer comments. Orders, however, are clearly related to Comment volumes over a longer periods, i.e., 5 and 7 days. Overall, we recognize the importance of

features generated from Comment activities, as they contribute to the predictors of all EPA types.

TABLE VI: Ten top-ranked features of the 3-q Random Forest predictors for the next-day EPA.

Rank	Order	Clickthrough	Search
1	$7DComment_{min}$	$7DComment_{min}$	$7DComment_{min}$
2	$7DComment_{sum}$	$3DComment_{mean}$	$5DRepost_{min}$
3	$7DComment_{mean}$	$3DComment_{sum}$	PreviousDComment
4	$7DComment_{max}$	PreviousDComment	$7DComment_{std}$
5	$5DComment_{mean}$	$3DComment_{max}$	$3DComment_{mean}$
6	$5DComment_{max}$	$7DComment_{sum}$	$7DRepost_{std}$
7	$5DComment_{sum}$	$7DComment_{mean}$	$5DRepost_{sum}$
8	$7DComment_{std}$	$7DPost_{sum}$	$3DComment_{sum}$
9	$7DComment_{var}$	$3DComment_{min}$	$7DComment_{var}$
10	$3DComment_{mean}$	$5DComment_{var}$	$5DComment_{var}$

VI. CONCLUDING REMARKS

This paper presents a detailed empirical study of the relationship between SMA of product vendors and EPA of consumers interested in the vendors' products. The study is the first to characterize the correlation of specific SMA engagements, i.e., posts, reposts and comments, and EPA types that correspond to specific stages in the consumers' purchase decisions, i.e., search for brands and products, clickthrough

product information, and product orders. Our analyses uncovered low-to-moderate correlations between the volumes of SMA-EPA pairs, suggesting that predicting daily volumes of EPA based on SMA volumes alone is not well supported. However, moderate correlations of Post-Search, Post-Clickthrough and Comment-Order across vendors, suggest that one may be able to train predictors of EPA distributions and their changes rather than the precise daily volumes. We thus introduce a new approach of characterizing SMA-EPA relationship in terms of EPA quantiles. We formulate EPA predictions as a multi-class categorization problem into 2, 3 and 5 quantiles. Our Random Forest classifiers outperform both the random predictors and the Logistic Regression classifiers for top quantiles of Order and bottom quantiles of Search and Clickthrough. Our study provides unique insights into the varied correlations of SMA-EPA pairs and a mixed success of SMA-based predictors in determining EPA levels. A general view that social media engagements undoubtedly drive consumer traction on e-commerce platforms is not substantiated by the Sina Weibo and JD case study and requires further analysis. In our future work we will expand explorations of this issue with a broader set of product categories and SMA analyses that consider properties of the content exchanged through the SMA.

VII. ACKNOWLEDGEMENTS

We wish to express our gratitude to JD.com for providing access to invaluable data and Bin Xu and Hao Dong for their assistance. We also acknowledge the financial support from the International Doctoral Innovation Centre, Ningbo Education Bureau, Ningbo Science and Technology Bureau, and the University of Nottingham. This work was also supported by the UK Engineering and Physical Sciences Research Council [grant number EP/L015463/1].

REFERENCES

- [1] J. Yeo, S. Kim, E. Koh, S.-w. Hwang, and N. Lipka, "Predicting online purchase conversion for retargeting," in *Proceedings of the 10th ACM International Conference on Web Search and Data Mining*, pp. 591–600, ACM, 2017.
- [2] C. Lo, D. Frankowski, and J. Leskovec, "Understanding behaviors that lead to purchasing: A case study of pinterest," in *Proceedings of the 22th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 531–540, 2016.
- [3] W. W. Moe, "Buying, searching, or browsing: Differentiating between online shoppers using in-store navigational clickstream," *Journal of Consumer Psychology*, vol. 13, no. 1-2, pp. 29–39, 2003.
- [4] C. Sismeiro and R. E. Bucklin, "Modeling purchase behavior at an e-commerce web site: A task-completion approach," *Journal of Marketing Research*, vol. 41, no. 3, pp. 306–323, 2004.
- [5] C. Park, D. Kim, J. Oh, and H. Yu, "Predicting user purchase in e-commerce by comprehensive feature engineering and decision boundary focused under-sampling," in *Proceedings of the 2015 International ACM Recommender Systems Challenge*, p. 8, ACM, 2015.
- [6] D. Van den Poel and W. Buckinx, "Predicting online-purchasing behaviour," *European Journal of Operational Research*, vol. 166, no. 2, pp. 557–575, 2005.
- [7] D. J. Bertsimas, A. J. Mersereau, and N. R. Patel, "Dynamic classification of online customers," in *Proceedings of the 2003 SIAM International Conference on Data Mining*, pp. 107–118, SIAM, 2003.
- [8] K. L. Kotler, Philip Keller, *Marketing Management, volume 14*. Prentice Hall Englewood Cliffs, 2015.
- [9] L. Xu, J. A. Duan, and A. Whinston, "Path to purchase: A mutually exciting point process model for online advertising and conversion," *Management Science*, vol. 60, no. 6, pp. 1392–1412, 2014.
- [10] X. Yu, Y. Liu, X. Huang, and A. An, "Mining online reviews for predicting sales performance: A case study in the movie domain," *IEEE Transactions on Knowledge and Data Engineering*, vol. 24, no. 4, pp. 720–734, 2012.
- [11] E. Young Kim and Y.-K. Kim, "Predicting online purchase intentions for clothing products," *European Journal of Marketing*, vol. 38, no. 7, pp. 883–897, 2004.
- [12] D. Hawkins, D. Mothersbaugh, and R. Best, "Consumer behaviour: Building marketing strategy mcgraw hill," 2012.
- [13] G. Kulkarni, P. Kannan, and W. Moe, "Using online search data to forecast new product sales," *Decision Support Systems*, vol. 52, no. 3, pp. 604–611, 2012.
- [14] J. Wang, W. X. Zhao, H. Wei, H. Yan, and X. Li, "Mining new business opportunities: Identifying trend related products by leveraging commercial intents from microblogs," in *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pp. 1337–1347, 2013.
- [15] B. Hollerit, M. Kröll, and M. Strohmaier, "Towards linking buyers and sellers: detecting commercial intent on twitter," in *Proceedings of the 22nd International Conference on World Wide Web*, pp. 629–632, ACM, 2013.
- [16] F. Kooti, K. Lerman, L. M. Aiello, M. Grbovic, N. Djuric, and V. Radosavljevic, "Portrait of an online shopper: Understanding and predicting consumer behavior," in *Proceedings of the 9th ACM International Conference on Web Search and Data Mining*, pp. 205–214, ACM, 2016.
- [17] J. Leskovec, L. A. Adamic, and B. A. Huberman, "The dynamics of viral marketing," *ACM Transactions on The Web (TWEB)*, vol. 1, no. 1, p. 5, 2007.
- [18] A. Anagnostopoulos, R. Kumar, and M. Mahdian, "Influence and correlation in social networks," in *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 7–15, ACM, 2008.
- [19] D. Crandall, D. Cosley, D. Huttenlocher, J. Kleinberg, and S. Suri, "Feedback effects between similarity and social influence in online communities," in *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 160–168, ACM, 2008.
- [20] C. R. Shalizi and A. C. Thomas, "Homophily and contagion are generically confounded in observational social network studies," *Sociological Methods & Research*, vol. 40, no. 2, pp. 211–239, 2011.
- [21] V. Gupta, D. Varshney, H. Jhamtani, D. Kedia, and S. Karwa, "Identifying purchase intent from social posts," in *ICWSM*, 2014.
- [22] Y. Zhang and M. Pennacchiotti, "Predicting purchase behaviors from social media," in *Proceedings of the 22nd International Conference on World Wide Web*, pp. 1521–1532, ACM, 2013.
- [23] X. W. Zhao, Y. Guo, Y. He, H. Jiang, Y. Wu, and X. Li, "We know what you want to buy: a demographic-based system for product recommendation on microblogs," in *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 1935–1944, ACM, 2014.
- [24] R. Crane and D. Sornette, "Robust dynamic classes revealed by measuring the response function of a social system," *Proceedings of the National Academy of Sciences*, vol. 105, no. 41, pp. 15649–15653, 2008.
- [25] J. Lehmann, B. Gonçalves, J. J. Ramasco, and C. Cattuto, "Dynamical classes of collective attention in twitter," in *Proceedings of the 21st International Conference on World Wide Web*, pp. 251–260, ACM, 2012.
- [26] D. M. Romero, B. Meeder, and J. Kleinberg, "Differences in the mechanics of information diffusion across topics: idioms, political hashtags, and complex contagion on twitter," in *Proceedings of the 20th International Conference on World Wide Web*, pp. 695–704, ACM, 2011.
- [27] J. Yang and J. Leskovec, "Patterns of temporal variation in online media," in *Proceedings of the 4th ACM International Conference on Web Search and Data Mining*, pp. 177–186, ACM, 2011.
- [28] M. G. Akritas, S. A. Murphy, and M. P. Lavalley, "The theil-sen estimator with doubly censored data and applications to astronomy," *Journal of the American Statistical Association*, vol. 90, no. 429, pp. 170–177, 1995.
- [29] D. W. Hosmer Jr, S. Lemeshow, and R. X. Sturdivant, *Applied logistic regression*, vol. 398. John Wiley & Sons, 2013.
- [30] H.-F. Yu, F.-L. Huang, and C.-J. Lin, "Dual coordinate descent methods for logistic regression and maximum entropy models," *Machine Learning*, vol. 85, no. 1, pp. 41–75, 2011.
- [31] A. Liaw, M. Wiener, *et al.*, "Classification and regression by random-forest," *R News*, vol. 2, no. 3, pp. 18–22, 2002.