Selecting Valuable Customers for Merchants in E-commerce Platforms

Yijun Wang[†], Le Wu[‡], Zongda Wu[§], Enhong Chen[†*] and Qi Liu[†]
[†]University of Science and Technology of China [‡]HeFei University of Technology [§]Wenzhou University
wyjun@mail.ustc.edu.cn lewu@hfut.edu.cn zongda1983@163.com {cheneh, qiliuql}@ustc.edu.cn

***Abstract*—An e-commerce website provides a platform for *merchants* to sell *products* to *customers*. While most existing research focuses on providing customers with personalized product suggestions by recommender systems, in this paper, we consider the role of merchants and introduce a parallel problem, i.e., how to select the most valuable customers for a merchant? Accurately answering this question can not only help merchants to gain more profits, but also benefit the ecosystem of e-commence platforms. To deal with this problem, we propose a general approach by taking into consideration the *interest* and *profit* of each customer to the merchant, i.e., select the customers who are not only interested in the merchant to ensure the visit of the merchant, but also capable of making good profits. Specifically, we first generate candidate customers for a given merchant by using traditional recommendation techniques. Then we select a set of the valuable customers from candidate customers, which has the balanced maximization between the interest and the profit metrics. Given the NP-hardness of the balanced maximization formulation, we further introduce efficient techniques to solve this maximization problem by exploiting the inherent submodularity property. Finally, extensive experimental results on a real-world dataset demonstrate the effectiveness of our proposed approach.**

I. Introduction

With the rapid development of the Internet, e-commerce as a new business model is becoming increasingly popular, i.e., online purchasing through e-commerce sites has become one of the most important commercial forms between customers and merchants [1], [2]. Thus, it is necessary for an e-commerce site to provide quality services for its two kinds of important users, customers and merchants, so as to increase the revenues of the e-commerce site.

From a customer perspective, it is important for an e-commerce site to select a relatively small set of products that satisfy the customer's requirements. This problem can be overcome by a recommender system, which focuses on designing a proper utility function to predict the degree of a customer interested in a product according to the customer's profiles (e.g., the consumption histories) [3]. From a merchant perspective, it is necessary to select a relatively small set of more valuable customers for a merchant, so as to guide the marketing efforts of the merchant, and as a result to increase the profits of merchants and benefit the e-commerce ecosystem. Here, it is natural to think that the valuable customers are those who would like to visit the merchant, meanwhile, they are capable of making good profits. Therefore, the value of a customer can be measured by the following two aspects: (1) the **interest** in a merchant, i.e., the more interested the customer is in the merchant, the more valuable the customer; and (2) the **profit** for a merchant, i.e., the higher profits the products (that a customer is interested in) have, the more valuable the customer.

Although personalized recommendation has demonstrated its effectiveness in improving customers' experience [4], directly applying personalized recommendation techniques to select customers for a merchant would result in poor performance, e.g., in general, personalized recommendation only takes the interest metric of customers into consideration, without considering the profit metric. Therefore, how to select valuable customers for a merchant remains pretty much open. Intuitively, one straightforward way for customer selection is to first select interested customers, and then further select profitable customers from them. However, such a strategy will incur several challenges. First, as we know, in online shopping process, it is more likely for a customer to browse interested products at first rather than a merchant, i.e., it is difficult to directly obtain the interest of a customer to a merchant from the customer's profiles. Thus, how to accurately compute the interest of a customer to a merchant based on the customer's profiles becomes the first challenge. Second, different products from the same merchant may have different profits, which requires that the selected customers would likely to purchase the products of higher profits. Thus, how to ensure that the selected customers are as profitable as possible for the merchant is another main challenge.

To address the above challenges, in this paper, we propose a general approach for customer selection. Specifically, we first generate a set of candidate customers for a merchant by using a recommendation algorithm based on the customers' implicit feedback. Second, from these candidates, we select a smaller set of the valuable customers who are not only interested in the merchant, but also capable of generating good profits. In the above process, the selection for the most valuable customers is a NP-hard problem, but its objective function meets the property of non-negative monotone submodularity. Therefore, we present an efficient algorithm to solve the problem. Finally, extensive experimental results on a real-world dataset demonstrate the effectiveness of our approach. Please note that, though we use detailed algorithms in each step, our approach presents a general framework for customer selection, where each step is open to other algorithms (i.e., it can be replaced by other similar algorithms).

$*$ Contact Author.

2374-8486/16 $31.00 © 2016 IEEE
DOI 10.1109/ICDM.2016.77

II. Related Work

Collaborative Filtering based Recommender Systems. In an e-commerce platform, *C*ollaborative *F*iltering (CF) is a popular approach to build a recommender system for improving customer satisfaction. Approaches for collaborative filtering fall into two categories: neighborhood-based methods [5], [6] and matrix factorization models [7]. These traditional CF models rely on the explicit customer-product ratings. However, customers' implicit feedback is more common than explicit feedback in our daily lives [8]. Thus, some recent methods proposed CF techniques based on users' implicit feedbacks [9], [10]. In this paper, our proposed problem of selecting valuable customers to merchants resembles a recommendation task, thus we borrow the key ideas of CF problem. However, our work distinguishes from CF as the merchants care more about the potential profits of the selected customers, which has rarely been considered by CF methods.

Customer Identification for Marketing. Customer identification is an important part of customer relationship management for building long term, profitable relationships with specific customers [11]. This phase involves targeting the population who are most likely to become customers or most profitable to the company through analysis of customers' underlying characteristics [12]. Typically, some researchers focus on segmenting customers into different groups, thus providing targeted marketing for each group. In fact, most of the techniques in data mining could be directly applied to the customer segmentation task [11], [13].

In fact, social influence based targeting has become a hot topic for identifying valuable customers in a social network. We borrow the ideas of social influence targeting for providing marketing efforts for merchants. Recently, some data mining researchers propose to devise marketing efforts from various angles. For example, Ma et al. proposed to identify hesitant and interested customers for a given company or product [14]. Tang et al. proposed to consider the magnitude of influence and the diversity of the influenced crowd simultaneously [15]. Liu et al. proposed an integrated marketing approach which combines targeted marketing and viral marketing [16]. Our work distinguishes from them as we take the interest and profit quality into consideration when selecting customers for a given merchant, which is rarely concerned in previous works.

III. Problem Statement

In this section, we present the preliminaries and problem formulation. Table I lists key notations used in this paper.

Preliminaries. An e-commerce site consists of a set of customers U, a set of merchants M and a set of products V. We use u, v, m to denote a customer, a product and a merchant, respectively. A merchant m is associated with a product set V_m. We represent the profit rate and price of a product v as R_v and T_v. The interaction between all the customers and products forms a consumption matrix, denoted by $C_{|U|\times|V|}$.

Formally, given the historical consumption matrix $C_{|U|\times|V|}$, the profit rate and price of each product, we formulate the

TABLE I
MATHEMATICAL NOTATIONS

Notations	Description
U V M	the set of customers/ products/ merchants
u v m	a customer/ product/ merchant
V_m	the product set of merchant m, $V_m \subseteq V$
U_m	candidate customer set of merchant m
P_u	consumed products of customer u, $L_u \subseteq V$
C_v	candidate customer set of product v
L_u	products which customer u is in their candidate customer sets
$\hat{x}_{uv}$	the preference score of customer u to product v
$\hat{r}_{um}$	the preference score of customer u to merchant m
R_v	the profit rate of product v
T_v	the price of product v

problem we study in this paper as the k-**M**ost **V**aluable **C**ustomers (k-MVC) problem.

Problem Formulation. *Given the customers' past consumption behaviors $C_{|U|\times|V|}$, the profit rate R_v and price T_v of each product v, given a merchant m and its associated product set V_m, the k-MVC problem aims at automatically selecting the k most valuable customers S_m from U, which has the best balance between the interest and profit goals.*

Since the two goals may conflict with each other, we propose to set our optimization function as a linear combination of them. For a given merchant m, we formulate the valuable customer selection as follows:

$$\max_{S_m} F(S_m) = \alpha \frac{I(S_m)}{\overline{I}} + (1-\alpha)\frac{E(S_m)}{\overline{E}}, \quad s.t. \quad S_m \subseteq U_m, |S_m| = k, \tag{1}$$

where $I(S_m)$ represents the interest score of the customer set S_m, and $E(S_m)$ is the expected profit from S_m. $\overline{I}$ and $\overline{E}$ are normalization factors, which is equal to the maximum values of $I(S_m)$ and $E(S_m)$, respectively. In addition, we balance the interest and profit goals using a parameter $\alpha \in [0, 1]$.

IV. The Proposed Framework

In this section, we describe the proposed framework. As shown in Figure 1, for each product in a merchant, we first generate many candidate customers that may be interested in the product by a recommendation algorithm. Second, the union of candidate customers of all the products is considered as the candidate customer set of the merchant. Finally, we select a set of target customers from the candidate customers by balancing the interest scores and the profit scores of customers.

A. Identifying Interested Customers

To identify interested customers for a product, the key is to measure the customer interest accurately. Since the customer feedback is implicit (either 1 or 0, i.e., buy or not), we use the widely used Bayesian personalized pairwise ranking method (BPR) [9] to compute customers' interest scores. We obtain a ranking list for each product by learning a pairwise ranking function $p(>_v \ |\Theta)$ that generates a partial order between each pair of customers which can be obtained based on the following assumption.

Assumption 1 *Given a product, a customer who has shown positive actions (e.g., purchase it, add it to the shopping cart or show likeness to it) has a greater interest in this product*

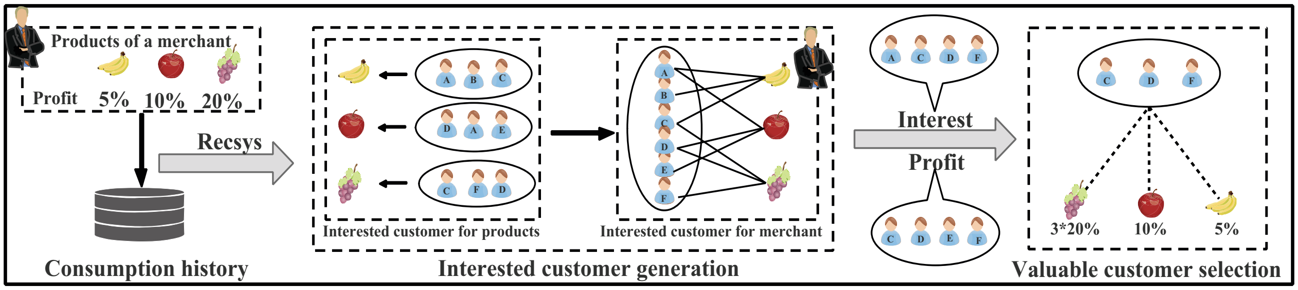

Fig. 1. Flowchart of the Proposed Framework

than the customers who have only clicked the product, without other acts.

To formalize the above assumption, we use U_v^p to denote a set of customers who have taken positive actions (i.e., purchase, collect or cart) on product v, and U_v^c to denote a set of customers who have clicked v. Then, we create training data $D_S := V \times U \times U$ by:

$$D_S = \{(v, i, j) | i \in U_v^p \wedge j \in U_v^c\}.$$

We formulate the prediction of the interest from a customer i to the product v as:

$$\hat{x}_{iv} = Q_v^T P_i, \tag{2}$$

where P_i and Q_v are the factorized low-rank feature vectors for the customer i and the product v, respectively.

Then, we can define the probability that customer i has a greater interest on product v than the customer j as $p(i >_v j|P, Q) = \sigma(\hat{x}_{ijv})$, where $\sigma(x)$ is the logistic sigmoid and $\hat{x}_{ijv} = \hat{x}_{iv} - \hat{x}_{jv}$. Further, if we use $>_v$ to denote all the ranking orders, the posterior probability that we need to maximize can be presented as follows.

$$p(P, Q| >_v) \propto p(i >_v j|P, Q) \cdot p(P, Q). \tag{3}$$

Then, the low-rank feature vectors can be learned by minimizing the following objective function:

$$\begin{aligned} L = \quad & -\textstyle\sum_{(v,i,j)\in D_S} \ln \sigma(\hat{x}_{iv} - \hat{x}_{jv}) + \lambda_i \sum_{i\in U} ||P_i||^2 \\ & +\lambda_j \textstyle\sum_{j\in U} ||P_j||^2 + \lambda_v \sum_{v\in V} ||Q_v||^2. \end{aligned} \tag{4}$$

After obtaining P and Q, we can apply them to estimate the interest scores of the customers to products. Without loss of generality, we select a set of customers with the highest values as the candidate customer set C_v for a product v. Then, the union of all the candidate customer sets is considered as the candidate customer set U_m of the merchant m.

B. Valuable Customer Selection

In this subsection, we describe how to select valuable customers by jointly modeling the customers' interest and profit metrics.

1) Interest and Profit Computation: To obtain the solution of Equation 1, we need to first obtain the interest and profit measures. Given a merchant m, let U_m be the candidate customer set that is generated in the previous subsection, and V_m the product set associated with the merchant m. For each $v \in V_m$, $\hat{x}_{uv}$ is the predicted interest score of the customer u to the product v, where the customer u belongs to the candidate customer set U_m. Then, we can define the interest score of the customer u to merchant m as $\hat{r}_{um} = \sum_{v\in V_m} \hat{x}_{uv}$. Thus, the interest score of the customer set S_m to merchant m could be further defined as $I(S_m) = \sum_{u\in S_m} \hat{r}_{um}$.

Next, we show how to compute $E(S_m)$. Specifically, for each product with the selected customer set S_m, we consider the following three aspects:

How many times the product is covered by the interests of the customers. Specifically, for a customer u and a product v, if $u \in C_v$, then v belongs to the interested product set L_u of u. The intuition is that, for a product v in a merchant m, suppose $L_{u_1}, L_{u_2}, ..., L_{u_k}$ are interested product sets of the selected customers $u_1, u_2, ..., u_k$ of m, then the more product sets contain the product v, the more likely v is to be purchased by the selected customers S_m.

The predicted interest scores of the customers to the product. The predicted interest score $\hat{x}_{uv}$ is used to measure how much a customer u is interested in a product v. Thus, it is obvious that selecting the customers with higher interest scores for a product will increase the probability of the customers to purchased the customers.

The negative effect from the increasing number of the customers. For a product v, due to the restricted factors from the product storage, the customer budget, etc., the probability of a successful purchase would decrease as the size of the interested customer set increases. That is, the more customers would like to purchase a product, the more difficult the product is obtained due to the restrictions from external factors.

Based on the above three factors, we present the profit metric $E(S_m)$ as follows:

$$E(S_m) = \sum_{v\in V_m} f(\sum_{u\in S_m} \hat{x}_{uv}[u \in C_v])R_v T_v, \tag{5}$$

where $f : \mathbb{R} \to \mathbb{R}$ can be any increasing concave function satisfying $f(0) = 0$ (to ensure $E(\varnothing) = 0$).

From the profit metric, we observe that for each product v, if a customer $u \in S_m$ appears in the interested customer set C_v of v, then we compute the interest score $\hat{x}_{uv}$ of u, which combines the coverage times and predicted interest scores mentioned above. In addition, the introduction of an increasing concave function $f(x)$ is used to reflect the negative effect on purchasing with the increasing size of the customer set. Actually, similar kinds of formulation are also used for diversified social influence maximization application [15] and diversified recommendation generation [17].

We can further formulate the k-MVC problem as follows:

$$\max_{S_m} F(S_m) = \alpha \frac{\sum_{u \in S_m} \hat{r}_{um}}{\overline{I}} \quad (6)$$
$$+(1-\alpha)\frac{\sum_{v \in V_m} f(\sum_{u \in S_m} \hat{x}_{uv}[u \in C_v]) R_v T_v}{\overline{E}},$$
$$s.t. \quad S_m \subseteq U_m, |S_m| = k.$$

2) Related Models: We name the proposed model considering the above-mentioned factors as IPS (short for **I**nterested and **P**rofitable customers **S**election). Since the IPS model is a general model, we can obtain a series of related models by changing the settings of IPS. Specifically, when we do not consider the difference of the product profits, i.e., $R_v T_v = 1$, we name the model as IPS_EP (**I**nterested and **P**rofitable customers **S**election with **E**qual product **P**rofit). When we treat $\hat{x}_{uv}$ as 1 when computing $E(S_m)$, we get another model named IPS_EI (**I**nterested and **P**rofitable customers **S**election with **E**qual **I**nterest score). Finally, when we ignore the product profit $R_v T_v$ and the interest score $\hat{x}_{uv}$ both, we get the model IPS_EPI (**I**nterested and **P**rofitable customers **S**election with **E**qual product **P**rofit and **I**nterest score).

3) Solution: We first explore some properties of the k-MVC problem in the following theorems.

Theorem 1. *It is a NP-hard problem to search the exact solution of the k-MVC selection problem.*

The proof of this theorem can be achieved by reducing the k-MVC problem to the NP-hard Weighted Maximum Coverage Problem [18].

Next, we present Theorem 2 proposed in [19], and then prove the submodularity of the object $F(S_m)$.

Theorem 2. *Given functions $F : 2^V \to \mathbb{R}$ and $f : \mathbb{R} \to \mathbb{R}$, the composition $F^{'} = f \circ F : 2^V \to \mathbb{R}$ (i.e., $F^{'}(S) = f(F(S))$) is nondecreasing submodular, if f is nondecreasing concave and F is nondecreasing submodular.*

Theorem 3. *The objective function $F(S_m)$ of the k-MVC problem is submodular.*

Proof. For any $\forall A \subseteq B \subseteq U \setminus u$, $I(A \cup u) - I(A) = \frac{r_u}{\overline{I}} = I(B \cup u) - I(B)$, so we conclude that $I(S_m)$ is submodular. Similarly, for any $\forall A \subseteq B \subseteq U \setminus u$, we can prove that $\sum_{u \in S_m} \hat{x}_{uv}[u \in C_v]$ is submodular. Then, based on Theorem 2, we can prove the submodularity of $f(\sum_{u \in S_m} \hat{x}_{uv}[u \in C_v])$. Because f is nondecreasing concave when $\sum_{u \in S_m} \hat{x}_{uv}[u \in C_v]$ is submodular, $f(\sum_{u \in S_m} \hat{x}_{uv}[u \in C_v])$ is also nondecreasing submodular with respect to S_m. Finally, since the submodularity is closed under nonnegative linear combinations, $F(S_m)$ is submodular. □

Below, we will give the performance guarantee of maximizing submodular functions using a simple greedy algorithm.

Guarantee on Solution Quality. As proved above, although the k-MVC problem is NP-hard, it can be solved by maximizing a non-negative monotone submodular function. Specifically, for a non-negative monotone submodular function $F : 2^U \to \mathbb{R}$, let $S \subseteq U$ be a set of size k obtained by selecting elements from U one by one, where the element that leads to the greatest increase of the margin of the function value is selected each time. Let $S^* \subseteq U$ be the set of the maximum value of F over all the k-element sets. Then S provides a $(1 - \frac{1}{e})$-approximation compared to the optimal solution [20].

Scaling Up by Reducing Function Computations. Inspired by [21], we attempt to improve the computation performance of the algorithm. Assume that we have obtained the marginal increments $\delta_A(u) = F(A \cup u) - F(A)$ for $u \in U \setminus A$. The key idea is that, for $A \subseteq B \subseteq U$, it holds that $\delta_A(u) \geq \delta_B(u)$ for $u \in U \setminus B$. Thus, instead of recomputing $\delta(u) = \delta_A(u)$ for each customer after adding a new element into the customer set, we perform lazy evaluations: we preserve a list recording the marginal gain $\delta(u)$ for each candidate customer u of the merchant m in a decreasing order. When searching the next customer for adding it into the customer set S_m, we traverse the nodes in the list. If $\delta(u)$ for the top node is invalid, we recompute it, and insert it into the list according to the existing order; otherwise, we move the top node to the customer set S_m. Therefore, we can find a new customer, without recomputing $\delta(u)$ for each candidate customer, consequently, improving the computation performance.

TABLE II
THE STATISTICS OF THE DATESET BEFORE AND AFTER PROCESSING

	Original Data	Pruned Data
#Customers	9,774,184	1,233,341
#Items	8,133,507	220,291
#merchants	86,799	6,413
#consumption records	12,627,634	1,633,105
#avg. items per seller	93.71	34.35
#avg. records per product	1.55	7.41
#avg. records per customer	1.29	1.32

V. EXPERIMENTS

We conduct experiments on a real-world dataset to demonstrate (1) the impact of the parameter α on valuable customer selection, (2) the performance of our approach compared to other models and (3) the correlation of the selection results from different models.

A. Experimental Setup

1) Dataset: The dataset Rec-$Tmall$[1] and the platform $TianChi$ are provided by Tmall[2], a well-known commerce site in China. Rec-$Tmall$ contains enormous logs of online customer behaviors. $TianChi$ is running on an Open Data Processing Service (ODPS), which is developed to deal with big data.

To ensure the reliability of the experimental results, we first preprocess the dataset. The statistics of the dataset before and after preprocessing are presented in Table II. We only keep the products that have more than 5 purchase records and the merchants having more than 10 products. For each product, we use the latest 20% of its records for testing and the rest records for training.

Due to the privacy issue, the dataset Rec-$Tmall$ doesn't provide the profit rate and price of each product. Therefore, without loss of generality, we set the profit rate of each product from 5% to 30%, and the price of each product from the same

[1] https://tianchi.shuju.aliyun.com/datalab/dataSet.htm?spm=5176.100073.888.15.W14WTh&id=2

[2] www.tmall.com

merchant to be equal, so *the profit R_vT_v of each product can be reduced to R_v*. Specifically, we use the following three strategies to generate the profit rate for each product: (1) randomly; (2) in inverse proportion to the sales volume ; (3) in direct proportion to the sales volume

2) Benchmark Methods: To evaluate the effectiveness of the IPS model, we compare it with three related models: the **IPS_EP** model, the **IPS_EI** model, and the **IPS_EPI** model, mentioned above. Specifically, we use different nondecreasing concave functions $f(x)$ for these methods. Moreover, to evaluate the customer selection performance of our approach, we compare it with the following methods:

IS. In this method, only the part of interested customer selection is considered. Thus, it is equivalent to set α to 1 in our IPS model.

Traditional Recommendation Methods. We implement user-based CF [6] and item-based CF [5] for customer selection. In both UCF and ICF, the Customer-Merchant matrix is derived from the interaction between customers and merchants (i.e., purchasing behaviors of customers in merchants).

3) Evaluation Metrics: We evaluate the proposed model from three aspects: accuracy, average profit per customer, and average profit of purchased products.

Accuracy. Given the selected set S_m for the merchant m, pre_m measures the rate of selected customers that really purchase products in the merchant, which is defined as $Pre_m = \frac{|S_m \cup SP_m|}{|S_m|}$, where SP_m is the customers who really purchase products of merchant m.

Customer Average Profit. To measure the profit that the selected customers can make for the merchant, we propose the following metric to measure the average profit of each customer which is defined as $Pro\_u_m = \frac{\sum_{v \in I_m} Q_v R_v}{|SP_m|}$, where I_m is the products that are sold in the merchant m, and Q_v denotes the sales volume of the product v.

Product Average Profit. We use the average profit of the purchased products in the test set to measure the customers' profit preference, which is defined as $Pro\_v_m = \frac{\sum_{v \in I_m} Q_v R_v}{\sum_{v \in I_m} Q_v}$.

B. Sensitivity of Controlling Parameter

Before the experiments, we first set the size of the interested customer set for each product to be 40 since it can get the best precision result. Specifically, for each merchant and α, we run Equation 6 to obtain a set of 20 valuable customers from the candidate customers. Besides, the performance of benchmark methods IPS_EP, IPS_EI and IPS_EPI also can be changed by control parameters.

From the experimental results, we find that different functions $f(x)$ perform similarly, so we only demonstrate the results when $f(x) = \frac{x}{1+x}$. Figure 2 shows the experimental results on the precision and customer average profit metrics with different profit rate generation strategies.

From the figure, we have the following findings: (1) all the methods based on different product profit rate generation strategies have similar performance; and (2) our proposed approach IPS could increase the overall quality of selected customers under a relatively acceptable compromise on accuracy, and outperform other related models.

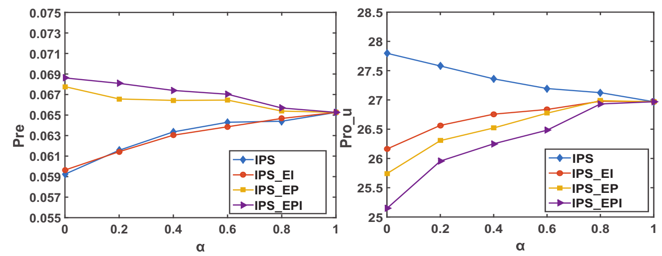

(a) Random Profit Rate

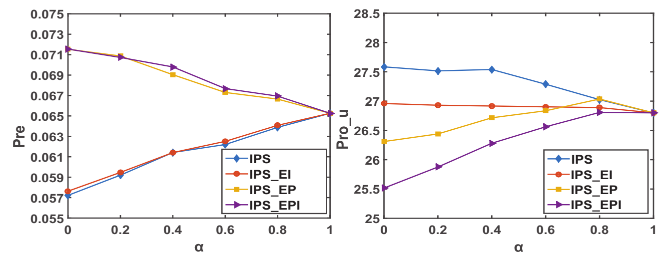

(b) Inverse Profit Rate

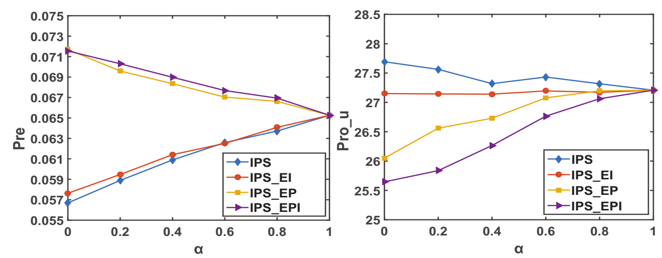

(c) Direct Profit Rate

Fig. 2. Trade-off between interest and profit by tuning α with different profit functions.

C. Customer Selection Performance Comparison

In the second group of experiments, we compare the customer selection performance of our proposed model with other models on different customer set sizes with different $f(x)$.

Customer Selection Performance. As mentioned above, IPS outperforms IPS_EP, IPS_EI and IPS_EPI in terms of the metrics of precision and customer average profit. Thus, on the two metrics, we here only show the results of IPS, IS, UCF and ICF. Besides, for the product average profit metric, we show the results of IPS, IPS_EI and IS, since the results of IPS_EP and IPS_EPI are much worse than those of IPS and IPS_EI, and the performance of ICF and UCF is almost the same to that of IS. From the experimental results, we find that different profit rate generation strategies have similar performance, so we only demonstrate the results when we generate each product's profit rate randomly. Here, we choose $f(x)$ as $\ln(x+1)$ and $x/(1+x)$, and set the value of α to 0.5. The precision and profit results are shown in Figure 3(a).

First, when comparing the results of IPS and IS, we can observe that IPS could improve the customer average profit and product average profit under a small compromise on the precision. Then, from the comparison between IPS and IPS_EI, we observe that IPS is always better than IPS_EI on

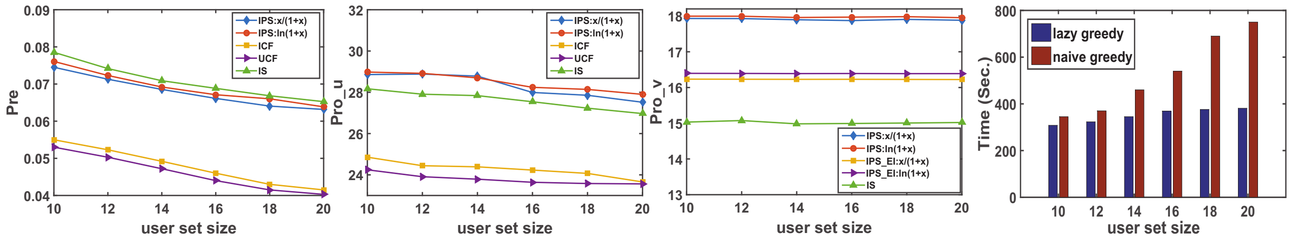

(a) The overall performance.

(b) The computational costs.

Fig. 3. Customer selection performance Comparison.

the metrics of customer average profit and product average profit, and slightly worse than IPS_EI on the precision metric. Second, from the comparison between our approach and other traditional CF models, we observe that both UCF and ICF perform poorly on the three metrics, and ICF is slightly better than UCF. In summary, the above results show that it is effective to use the two-stage strategy combined with the profit measure to select the valuable customers for merchants.

Time Efficiency. We compare the proposed lazy greedy algorithm with the exhaustive search and the naive greedy algorithm. Due to the impact of other factors (e.g., the scheduling of ODPS platform, the network conditions of the machine group cluster, the time consumption of I/O), the running times of IPS, IPS_EP, IPS_EI and IPS_EPI for the same solution are almost the same to each other. Here, we only present the average running time of lazy greedy algorithm and naive greedy algorithm, because exhaustive search is too time-consuming. The results are shown in Figure 3(b). From the results, we see that the proposed lazy forward algorithm can greatly improve the running efficiency compared to the naive greedy algorithm.

VI. Conclusion

To conclude, in this paper, we studied the problem on valuable customer selection, under the goal of maximizing the customer interested quality and profitable quality simultaneously. Along this line, we first generated many candidate customers using a Bayesian personalized pairwise ranking method based on the customer implicit feedback. Second, we selected the valuable customers who are not only interested in the merchant, but also capable of making good profits for the merchant. Finally, extensive experimental results on a real-world dataset demonstrated the effectiveness of our proposed approach. We hope this study could lead to more future work.

Acknowledgement. This research was partially supported by grants from the National Key Research and Development Program of China (Grant No. 2016YFB1000904), the National Science Foundation for Distinguished Young Scholars of China (Grant No. 61325010), the National Natural Science Foundation of China (Grant No. 61403358 and No. 61602147) and the Zhejiang Provincial Natural Science Foundation of China (No. LY15F020020). Qi Liu gratefully acknowledges the support of the CCF-Intel Young Faculty Researcher Program (YFRP).

References

[1] J. McAuley, R. Pandey, and J. Leskovec, "Inferring networks of substitutable and complementary products," in *SIGKDD*. ACM, 2015, pp. 785–794.

[2] W. Zhong, R. Jin, C. Yang, X. Yan, Q. Zhang, and Q. Li, "Stock constrained recommendation in tmall," in *SIGKDD*. ACM, 2015, pp. 2287–2296.

[3] X. Amatriain and B. Mobasher, "The recommender problem revisited: morning tutorial," in *SIGKDD*. ACM, 2014, pp. 1971–1971.

[4] P. Jiang, Y. Zhu, Y. Zhang, and Q. Yuan, "Life-stage prediction for product recommendation in e-commerce," in *SIGKDD*. ACM, 2015, pp. 1879–1888.

[5] B. Sarwar, G. Karypis, J. Konstan, and J. Riedl, "Item-based collaborative filtering recommendation algorithms," in *WWW*, 2001, pp. 285–295.

[6] J. L. Herlocker, J. A. Konstan, A. Borchers, and J. Riedl, "An algorithmic framework for performing collaborative filtering," in *SIGIR*. ACM, 1999, pp. 230–237.

[7] J. D. Rennie and N. Srebro, "Fast maximum margin matrix factorization for collaborative prediction," in *ICML*. ACM, 2005, pp. 713–719.

[8] S. Rendle and C. Freudenthaler, "Improving pairwise learning for item recommendation from implicit feedback," in *WSDM*. ACM, 2014, pp. 273–282.

[9] S. Rendle, C. Freudenthaler, Z. Gantner, and L. Schmidt-Thieme, "Bpr: Bayesian personalized ranking from implicit feedback," in *UAI*. AUAI Press, 2009, pp. 452–461.

[10] Y. Yao, H. Tong, G. Yan, F. Xu, X. Zhang, B. K. Szymanski, and J. Lu, "Dual-regularized one-class collaborative filtering," in *CIKM*. ACM, 2014, pp. 759–768.

[11] E. W. Ngai, L. Xiu, and D. C. Chau, "Application of data mining techniques in customer relationship management: A literature review and classification," *Expert systems with applications*, vol. 36, no. 2, pp. 2592–2602, 2009.

[12] J. Y. Woo, S. M. Bae, and S. C. Park, "Visualization method for customer targeting using customer map," *Expert Systems with Applications*, vol. 28, no. 4, pp. 763–772, 2005.

[13] C.-H. Wu, S.-C. Kao, Y.-Y. Su, and C.-C. Wu, "Targeting customers via discovery knowledge for the insurance industry," *Expert Systems with Applications*, vol. 29, no. 2, pp. 291–299, 2005.

[14] G. Ma, Q. Liu, L. Wu, and E. Chen, "Identifying hesitant and interested customers for targeted social marketing," in *PAKDD*. Springer, 2015, pp. 576–590.

[15] F. Tang, Q. Liu, H. Zhu, E. Chen, and F. Zhu, "Diversified social influence maximization," in *ASONAM*. IEEE, 2014, pp. 455–459.

[16] Q. Liu, Z. Dong, C. Liu, X. Xie, E. Chen, and H. Xiong, "Social marketing meets targeted customers: A typical user selection and coverage perspective," in *ICDM*. IEEE, 2014, pp. 350–359.

[17] L. Wu, Q. Liu, E. Chen, N. J. Yuan, G. Guo, and X. Xie, "Relevance meets coverage: A unified framework to generate diversified recommendations," *ACM Transactions on Intelligent Systems and Technology*, vol. 7, no. 3, p. 39, 2016.

[18] D. S. Hochba, "Approximation algorithms for np-hard problems," *ACM SIGACT News*, vol. 28, no. 2, pp. 40–52, 1997.

[19] H. Lin and J. Bilmes, "A class of submodular functions for document summarization," in *ACL-HLT 2011-Volume 1*. Association for Computational Linguistics, 2011, pp. 510–520.

[20] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher, "An analysis of approximations for maximizing submodular set functionsli," *Mathematical Programming*, vol. 14, no. 1, pp. 265–294, 1978.

[21] J. Leskovec, A. Krause, C. Guestrin, C. Faloutsos, J. VanBriesen, and N. Glance, "Cost-effective outbreak detection in networks," in *SIGKDD*. ACM, 2007, pp. 420–429.