Recommender Systems for Product Bundling

Moran Beladev
Ben-Gurion University of the Negev
belachde@bgu.ac.il

Bracha Shapira
Ben-Gurion University of the Negev
bshapira@bgu.ac.il

Lior Rokach
Ben-Gurion University of the Negev
liorrk@bgu.ac.il

ABSTRACT

Recommender systems (RSs) enhance e-commerce sales by recommending relevant products to their customers. RSs aim at implementing the firm's web-based marketing strategy to increase revenues. Generating bundles is an example of a marketing strategy that aims to satisfy consumer needs and preferences, and at the same time, to increase customers' buying scope and the firm's income. Thus, finding and recommending an optimal and personal bundle becomes very important. In this paper we introduce a novel model of bundle recommendations that integrates collaborative filtering (CF) techniques, personalized demand functions, and price modeling. This model provides a recommendation list by finding pairs of products that maximizes both, the probability of their purchase by the user and the revenue received by selling this bundles.

Categories and Subject Descriptors

H.3.3 [Information Storage and Retrieval]: Information Search and Retrieval—Information Filtering.

Keywords

Bundle Recommendation, Recommender Systems, E-Commerce, Collaborative Filtering, SVD.

1. INTRODUCTION

Bundling refers to the practice of selling two or more items together as a package at a price that is below the sum of the independent prices. Optimal bundling would combine items into bundles that best fit the retailer's needs and the user's preferences, and maximize product compliance within the bundle. Thus, a single price $\mathbf{P_{A+B}} < \mathbf{P_A} + \mathbf{P_B}$ is set for the two products (A, B) if purchased jointly. One challenge is to suggest a price for a bundle that fits both the customer reservation price i.e., the maximal price buyers are accepted to pay, and the retailer's revenue [1]. Very few studies have combined bundling strategy with recommender systems (RSs). The field of frequent item set mining and association rules deals with finding a basket of items that are frequently bought together [2]. However, these techniques are not personalized, thus not applicable for RSs. The recommendation of bundles were presented as a tailored solution for the tourism domain using case-based reasoning where case models representing the travel plan bundle were matched against the user profile and preferences [3]. The authors of [4] presented a bundle optimization using a genetic algorithm to maximize the compatibility of the products within a bundle. However, these studies did not measure the recommendation aspect, i.e., if it is at all feasible and beneficial to predict bundle purchasing. The study presented in [5] introduces a bundle recommendation problem, in which its solution is a set of items that maximizes some total expected reward. However, the price aspect was not considered in the model. Our paper maximizes the expected revenue by considering the item-to-item cross dependencies, user-item collaborative filtering techniques and the demand-price function—resulting in recommendation of the best bundle and price proposal to the user.

2. BUNDLE RECOMMENDATION MODEL

We maximize the following retailer expected revenue function:

$$(1)\ ExpectedRevenue = P_i(A,B,T) \cdot (T - cost_A - cost_B)$$

where $P_i(A,B,T)$ is the probability that user i will purchase the bundle, which is composed of products A and B, at price T. The $cost_{\mathrm{A}}$ is the retailer's cost for product A and $cost_B$ is the retailer's cost for product B. The proposed bundle and the price T for user i is set to maximize the expected revenue:

$$(2)\ (A,B,T) = argmax_{\forall A,B,T} ExpectedRevenue(A,B,T)$$

In order to find $P_i(A,B,T)$ we find the corresponding prices C_A of product A and C_B of product B aggregated to the bundle price T:

$$(3)\ P_i(A,B,T) = max_{\forall c_A, c_B | c_A + c_B = T}\ P_i(A \cap B \cap C_A \cap C_B)$$

Thus, we have to find the prices C_A and C_B that maximize the probability of the user i to buy products A and B while paying those prices. According to Bayes' law:

$$(4)\ P_i(like\ A \cap willing\ to\ pay\ C_A\ for\ A) = P_i(like\ A) \cdot P_i(willing\ to\ pay\ C_A\ for\ A | like\ A)$$

According to the Jaccard measure: $(5)\ Jaccard = J_{A,B} = \frac{P(A \cap B)}{P(A \cup B)}$

Using combinatorial mathematics, the inclusion–exclusion principle: $(6)\ P(A \cap B) = P(A) + P(B) - P(A \cup B)$

Using equations (5) + (6): $(7)\ P(A \cap B) = \frac{P(A)+P(B)}{1+\frac{1}{J_{A,B}}}$

Using Bayes' law and equation (7):

$$(8)\ P_i((A \cap C_A) \cap (B \cap C_B)) = \frac{P_i(A) \cdot P_i(C_A|A) + P_i(B) \cdot P_i(C_B|B)}{1 + \frac{1}{J_{A,B}}}$$

We assume that the Jaccard measure, $J_{A,B}$, which denotes the products' compatibility, is not affected by the price. The $P_i(A)$, $P_i(B)$ probabilities are found using the CF technique; $P_i(C_A|A), P_i(C_B|B)$ is found by the upcoming personal demand.

2.1 Personalized Demand Graph

We would like to assume that each customer has its own demand graph for each product based on his/her preferences. Thus, we developed heuristics for estimating the "personalized" demand graph for user i and item j using very sparse data. Figure 1 demonstrates the demand of a generic customer versus an enthusiastic one (i.e., one that would pay high prices) as well as an indifferent one. We assume that the difference between the demand graphs can be reflected by the following:

$$(9)\ P_i(C_A|A) = min(P_{\cdot,A}(C_A) \times \alpha_{i,A}\ , 100\%)$$

where $P_{\cdot,A}(C_A)$ is the generic demand graph for item A given that the price C_A and $\alpha_{i,A}$ is the personalized bias factor for user i and item A. In order to find the personal bias factor, $\alpha_{i,A}$, we scan each customer's previous purchases or his/her highest bid on an item. We compare his/her price to the median of the generic graph. For example if customer i purchased item A for price C_A* then his/her bias factor is estimated as: (10) $\alpha_{i,A} = \frac{0.5}{P_{\cdot,A}(C_A^*)}$. For example (Figure 1), assume that a customer purchased the item for C_A*=1300; according to the generic graph, this price would be considered only by 35% of the interested population. Thus the personalized bias for this user is calculated as: $\alpha_{i,A} = \frac{0.5}{0.35} = 1.42$.

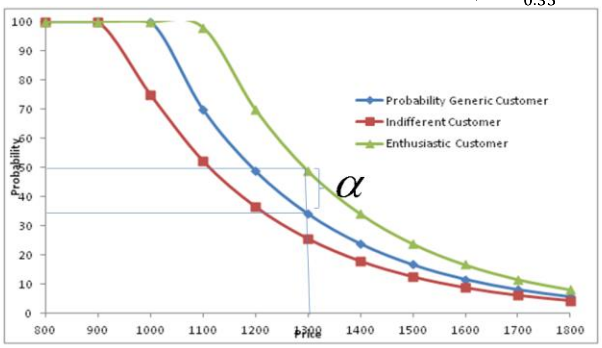

Figure 1. Personalized demand for various customer types

We can create a bias matrix for all purchases of items by users. The bias factor of products that have not been purchased by the customer can be predicted using the SVD method. Given the complete matrix, we can infer the personalized demand graph of each user i and item A from the generic demand graph calculated for item A and multiply it by the predicted alpha.

3. EVALUATION

Our model was evaluated based on two datasets. (Dataset 1) consists of transactions from a shopping website that sells electronics and furniture. (Dataset 2) is a supermarket dataset from Kaggle (*https://www.kaggle.com/c/acquire-valued-shoppers-challenge*). We used offline evaluation and compared our model to SVD and CF as baseline models. We evaluated: (i) The personal demand function by using a validation set in order to test the predicted alphas compared to the actual alphas, using the RMSE measure, and comparing the personal demand graph probability to 0.5 (median probability) of all purchased products in the test set- using the RMSE measure too; (ii) The product bundling recommendation by comparing the top 5 bundles to the top 5 items recommended by CF and SVD algorithms. For this we used precision, recall, the average quantity that was recommended and purchased, and the average price paid for the recommended and purchased products; (iii) The price bundling recommendation by comparing the recommended price to the actual price the user paid in the test set, measuring the sum of the absolute difference. The recommended price was compared to the mean price of the product. We also compared two strategies: (1) maximizing the bundle buying probability of the user, and (2) maximizing the expected revenue. For both datasets we evaluated our model on the top 1,000 customers and top 300 products. The first dataset resulted in 3,425 transactions and the second in 836,846 transactions. A recommended bundle is considered a hit in the test set if the two products have been purchased by the user within a week. In dataset 1 for the personal demand graph we received an RMSE of alpha of 0.072 and an RMSE error compared to the median of 0.261. Thus, the personal graphs are compatible to the users' preferences. In dataset 2, for the personal demand graph, we received an RMSE error of alpha of 1.067 and an RMSE of 0.34, compared to the median. The results for the product evaluation are presented in table 1 and 3 and the results for the price evaluation are presented in table 2 and 4. Bundle (1) and Bundle (2) represent the two strategies of maximizing probability and the expected revenue, respectively.

Table 1. Product bundling results for dataset 1

	Precision	Recall	Q	Price
CF	0.027	0.012	0.133	12.133
SVD	0.013	0.033	0.067	70.533
Bundle (1)	0.088	0.09	0.8	469.133
Bundle (2)	0.071	0.08	0.6	457.467

Table 2. Price bundling recommendation results for dataset 1

	Recommended price	Mean price
Bundle(1)	0.043	788.89
Bundle(2)	29.989	788.89

Table 3. Product bundling results for dataset 2

	Precision	Recall	Q	Price
CF	0.052	0.003	0.26	1.852
SVD	0.58	0.024	2.9	29.981
Bundle (1)	0.728	0.018	4.44	38.069
Bundle (2)	0.2	0.004	1.02	12.882

Table 4. Price bundling recommendation results for dataset 2

	Recommended price	Mean price
Bundle(1)	170.8	87.1
Bundle(2)	117.44	104.057

4. CONCLUSIONS AND FUTURE WORK

Our results demonstrate that bundles are predictable and may increase users' purchase scope. The first dataset is more difficult to predict, but the bundle model is at least comparable to state of the art algorithms and is even superior in some cases. The personal demand graph tends to be very accurate as was observed by the price recommendation accuracy. The second dataset contains commodities data, thus a personal demand graph is more difficult to predict. The recommended price was not as accurate as in the first dataset. Moreover, for dataset 2 the products are more predictable and the first bundle strategy yields the best results. For both datasets maximizing the probability of the user's purchase is more effective than maximizing the expected revenue. Future work will aim at improving the personal demand graph of dataset 2, examining more datasets and providing live user experiments.

5. REFERENCES

[1] Guiltinan, J. 1987. The price bundling of services: a normative framework. *The Journal of Marketing*, 74-85.

[2] Agrawal, R., Imielinski, T., and Swami, A.N. 1993. Mining association rules between sets of items in large databases. *ACM SIGMOD,* volume 22,2 of SIGMOD Record, 207–216.

[3] Ricci, F. 2002. ITR: a case-based travel advisory system. *Advances in Case-Based Reasoning*. Springer Berlin Heidelberg, 613-627.

[4] Birtolo, C. 2013. Searching optimal product bundles by means of GA-based Engine and Market Basket Analysis. *IFSA*

[5] Zhu, T. 2014. Bundle recommendation in ecommerce. SIGIR 2014, 657-666.