ISM: Item Selection for Marketing with Cross-Selling Considerations

Raymond Chi-Wing Wong, Ada Wai-Chee Fu
Department of Computer Science and Engineering
The Chinese University of Hong Kong
cwwong, adafu@cse.cuhk.edu.hk

Summary. Many different algorithms are studied on association rules in the literature of data mining. Some researchers are now focusing on the application of association rules. In this paper, we will study one of the application called Item Selection for Marketing (ISM) with cross-selling effect consideration. The problem ISM is to find a subset of items as marketing items in order to boost the sales of the store. We prove a simple version of this problem is NP-hard. We propose an algorithm to deal with this problem. Experiments are conducted to show that the algorithms are effective and efficient.

1 Introduction

In the literature of data mining, there are a lot of studies on association rules [2]. Such studies are particularly useful with a large amount of data in order to understand the customer behavior in their stores. However, it is generally true that the results of association rule mining are not directly useful for the business sector. Therefore there has been research in examining more closely the business requirements and finding solutions that are suitable for particular issues, such as marketing and inventory control. Recently, some researchers [10] studied the utility of data mining such as association and clustering, on decision making for revenue-maximizing enterprises. They have formulated the general problem as an optimization problem where a profit is to be maximized by determining a best strategy. The profit is typically generated from the customer behaviour in such an enterprise. More specific problems for revenue-maximizing are considered in more recent works [6,14,9,5,4,17]. The related problem of mining user behaviour is also of much research interest recently and a number of results can be found in [18].

In this paper we consider the problem of selecting a subset of items in a store for marketing in order to boost the overall profit. The difficulty of the problem is that we need to estimate the cross-selling effect to determine the influence of the marketed items on the sales of the other items. It is known that the records of sales transactions are very useful [3] and we determine the cross-selling effect with such information. We call the problem defined this way *Item Selection for Marketing (ISM)*. We show that a simple version of this problem is NP-hard. We propose a hill climbing approach to tackle this problem. In our experiment, we apply the proposed approach to a set of real data and the approach is found to be effective and efficient.

2 Related Work

One major target of data mining is solving decision making problems for the business sector. A study of the utility of data mining for such problems is investigated in [10], published in 1998. A framework based on optimization is presented for the evaluation of data mining operations. In [10] the general decision making problem is considered as a maximization problem as follows

$$\max_{x \in \mathcal{D}} \sum_{i \in C} g(x, y_i) \tag{1}$$

where $\mathcal{D}$ is the set of all possible decisions in the domain problem (e.g. inventory control and marketing), $\mathcal{C}$ is the set of customers, y_i is the data we have on customer i, and $g(x, y_i)$ is the utility (benefit) from a decision x and y_i. However, when we examine some such decision problems more closely, we find that we are actually dealing with a maximization problem of the form

$$\max_{x \in \mathcal{D}} g(x, Y) \tag{2}$$

where Y is the set of all y_i, or the set of data collected about all customers. The above is more appropriate when there are correlations among the behaviours of customers (e.g. cross-selling, the purchase of one item is related to the purchase of another item), or when there are interactions among the customers themselves (e.g. viral marketing, or marketing by word-of-month among customers). This is because we cannot determine $g()$ based on each single customer alone.

We illustrate the above in two different problems that have been studied. The first problem is about optimal product selection [5,4,16,17] (in SIGKDD 1999,2000,2002, and ICDM 2003, respectively). The problem is that in a typical retail store, the types of products should be refreshed regularly so that losing products are discarded and new products are introduced. Hence we are interested to find a subset of the products to be discontinued so that the profit can be maximized. The formulation of the problem considers the important factor of cross-selling which is the influence of some products on the sales of other products. The cross-selling factor is embedded into the calculation of the maximum profit gain from a decision. This factor can be obtained from an analysis of the history of transactions kept from previous sales which corresponds to the set Y in formulation (2). [1]

The second such problem is about viral marketing where we need to choose a subset of the customers to be the targets of marketing so that they can influence more of other customers. Some related work can be found in [6,14,9,1] (in SIGKDD 2001,2002,2003 and WWW2003, respectively). Again the profit gain from any decision relies on an analysis based on the knowledge collected about all customers.

[1] The problem is related to inventory management which has been studied in management science, however, previous works are mostly on the problems of when to order, where to order from, how much to order and the proper logistics [15].

The problem that we tackle here is of a similar nature since we also consider the factor of cross-selling when calculating the utility or benefit of a decision. In our modeling, we adopt concepts of the association rules to model the cross-selling effects among items.

Suppose we are given a set I of items, and a set of transactions. Each transaction is a subset of I. An *association rule* has the form $X \rightarrow I_j$, where $X \subseteq I$ and $I_j \in I - X$; the *support* of such a rule is the fraction of transactions containing all items in X and item I_j; the confidence for the rule is the fraction of the transactions containing all items in set X that also contain item I_j. The problem is to find all rules with sufficient support and confidence given some thresholds. Some of the earlier work include [13,2,12].

3 Problem Definition

In this section we introduce the problem of ISM. To the best of our knowledge, this is the first definition of item selection problem for marketing with the consideration of cross-selling effects. Item Selection for Marketing (ISM) is a problem to select a set of items for marketing, called *marketing items*, so as to maximize the total profit of marketing items and non-marketing items among all choices. In ISM, we assume that the sales of some items are affected by the sales of some other items. Given a data set with m transactions, $t_1, t_2, ..., t_m$, and n items, $I_1, I_2, ..., I_n$. Let $I = \{I_1, I_2, ..., I_n\}$. The profit of item I_a in transaction t_i before marketing is given by $prof(I_a, t_i)$. Let $S \subset I$ be a set of selected items. In each transaction t_i, we define two symbols, t'_i and d_i, for the calculation of the total profit.

$$t'_i = t_i \cap S, \qquad d_i = t_i - t'_i$$

Definition 1 (Profit Before Marketing). The original profit $Profit_0$ before marketing for all transactions is defined as:

$$Profit_0 = \sum_{i=1}^{m} \sum_{I_a \in t_i} prof(I_a, t_i) \tag{3}$$

Suppose we select a subset S of marketing items. Marketing action such as discounting will be taken on S. Let us consider a transaction t_i containing the marketing items I_a and non-marketing items I_b. If we market item I_a with a cost of $cost(I_a, t_i)$ (e.g. discount of item), the profit of item I_a after marketing in transaction t_i will become $prof(I_a, t_i) - cost(I_a, t_i)$. After the marketing actions are taken, more of the marketing items, says I_a, will be purchased. We define the changes in the sales by $\alpha(T)$, where T is a set of items:

$$\alpha(T) = \frac{\text{sale volume of } T \text{ after marketing}}{\text{sale volume of } T \text{ before marketing}}. \tag{4}$$

In the above the sale volume of T is measured by the total amount of the items in T that are sold in a fixed period of time. [2] If $\alpha(\{I_a\}) = 1$, then there is no increase of the sales of items I_a. If $\alpha(\{I_a\}) = 2$, then the sales of I_a is doubled compared with the sales before marketing.

On the other hand, without the consideration of cross-selling effect due to marketing, the profit of non-marketing items I_b is still $prof(I_b, t_i)$. With the consideration of cross-selling effects, some of the non-marketing items I_b will be purchased more if there is an increase of sales of marketing items I_a. The cross-selling factor is modelled by $csfactor(T, I_b)$, where T is a set of marketing items I_a, and $0 \leq csfactor(T, I_b) \leq 1$. That is, more customers may come to buy item I_b if some other items in T are being marketed. The increase of the sale of item I_b is modelled by $(\alpha(T) - 1)csfactor(T, I_b)$. [3] If $\alpha(T) = 1$, then there is no increase of sales of marketing items in set T. So, there is no increase of sales of non-marketing item I_b. The term $(\alpha(T) - 1)csfactor(T, I_b)$ becomes zero. Similarly, if $\alpha(T) = 2$, the sales of items in set T is doubled. Thus, the increase of sales is modelled by $csfactor(T, I_b)$.

Definition 2 (Profit After Marketing). The profit after marketing $Profit_1$ is defined as follows.

$$\begin{aligned} Profit_1 = \textstyle\sum_{i=1}^{m} \Big[& \textstyle\sum_{I_a \in t'_i} \alpha(\{I_a\})(prof(I_a, t_i) - cost(I_a, t_i)) \\ & + \textstyle\sum_{I_b \in d_i} (1 + (\alpha(t'_i) - 1)csfactor(t'_i, I_b))prof(I_b, t_i) \Big] \end{aligned} \tag{5}$$

Recall that t'_i is the set of items in transaction t_i that are selected to be marketed. For each transaction t_i, we compute the profit from the marketing items (discounted by $cost(I_a, t_i)$), and the profit from the non-marketing items whose sales are influenced by $csfactor()$. $Profit_1$ is the sum of the profits from all tranasctions. The objective of marketing is to increase the profit gain compared with the profit before marketing. The profit gain is defined as follows.

Definition 3 (Profit Gain). Profit gain is :

$$\text{Profit Gain} = Profit_1 - Profit_0 \tag{6}$$

From the above definitions, we can rewrite the profit gain as follows.

$$\begin{aligned} & \text{Profit Gain} = Profit_1 - Profit_0 \\ & = \textstyle\sum_{i=1}^{m} \Big[\textstyle\sum_{I_a \in t'_i} [(\alpha(\{I_a\}) - 1)prof(I_a, t_i) - \alpha(\{I_a\})cost(I_a, t_i)] \\ & \quad + \textstyle\sum_{I_b \in d_i} (\alpha(t'_i) - 1)csfactor(t'_i, I_b)prof(I_b, t_i) \Big] \end{aligned} \tag{7}$$

[2] We note that different items may have their different increase ratio of the sales (i.e. $\alpha(\{I_i\})$). However, it is difficult to predict this parameter $\alpha(\{I_i\})$ for each item I_i. For simplicity, we set all $\alpha(\{I_i\})$ to be the same (e.g. α_0) in this paper, which is the same as [6,14].

[3] If $\alpha(\{I_i\}) = \alpha_0$ for all i, then it is easy to see that $\alpha(T) = \alpha_0$ for any T (a subset of I) .

Next we can formally define the problem of ISM:

ISM: Given a set of transactions with profits assigned to each item in each transaction and the cross-selling factors, $csfactor()$, pick a set S from all given items which gives a maximum profit gain.

This problem is at least as difficult as the following decision problem.

ISM Decision Problem: Given a set of items and a set of transactions with profits assigned to each item in each transaction, a minimum profit gain G, and cross-selling factors, $csfactor()$, can we pick a set S such that $Profit\ Gain \geq G$?

Note that the cross-selling factor can be determined in different ways, one way is by the domain experts. Let us consider the very simple version where $csfactor(t'_i, I_a) = 1$ for any non-empty set of t'_i. That is, any selected items in the transaction will increase the sale of the other items with the same volume. This may be a much simplified version of the problem, but it is still very difficult.

Theorem 1 (NP-hardness). *The item selection for marketing (ISM) decision problem where $csfactor(t'_i, I_a) = 1$ for $t'_i \neq \phi$ and $csfactor(t'_i, I_a) = 0$ for $t'_i = \phi$ is NP-hard.*

Proof: We shall transform the problem of MAX CUT to the ISM problem. MAX CUT [7] is an NP-complete problem defined as follows: *Given a graph - (V, E) with weight $w(e) = 1$ for each $e \in E$ and positive integer K, is there a partition of V into disjoint sets V_1 and V_2 such that the sum of the weights of the edges from E that have one endpoint in V_1 and one endpoint in V_2 is at least K?* The transformation from MAXCUT to ISM problem is described as follows. Let $G = K$, $\alpha(\{I_a\}) = 2$, and $\alpha(t'_i) = 2$. For each vertex $v \in V$, construct an item. For each edge $e \in E$, where $e = (v_1, v_2)$, create a transaction with 2 items $\{v_1, v_2\}$. Set $prof(I_j, t_i) = 1$ and $cost(I_j, t_i) = 0.5$, where t_i is a transaction created in the above, $i = 1, 2, ..., |E|$, and I_j is an item in t_i. It is easy to check that Profit Gain $= \sum_{i=1}^{m} \sum_{I_b \in d_i} csfactor(t'_i, I_b)$. The above transformation can be constructed in polynomial time. When the problem is solved in the transformed ISM, the original MAX CUT problem is also solved. Since MAX CUT is an NP-complete problem, ISM problem is NP-hard.

□

4 Association Based Cross-Selling Effect

In the previous section, we see that the cross-selling factor is important in the problem formulation. The factor is indicated by $csfactor(t'_i, I_j)$, where t'_i is a set of items selected for marketing and I_j is another item. This factor can be provided by domain experts if they can estimate the impact of t'_i on I_j. However, in typical application, the amount of items would be large and it would be impractical to expect purely human analysis on these values. We suggest that the factor is to be determined by data mining technique based on the history of transactions collected for the application. We shall adopt the concepts of association rules for this purpose.

Definition 4. Let $d_i = \{Y_1, Y_2, Y_3, ..., Y_q\}$ where Y_i refers to a single item for $i = 1, 2, .., q$, then $\diamond d_i = Y_1 \vee Y_2 \vee Y_3 \vee .... \vee Y_q$. □

In our remaining discussion, $csfactor(t'_i, I_j)$ is equal to $conf(\diamond t'_i \rightarrow I_j)$, where $conf(\diamond t'_i \rightarrow I_j)$ is the confidence of the rule $\diamond t'_i \rightarrow I_j$ The definition of confidence here is similar to the definition of association rules. That is,

$$csfactor(t'_i, I_j) = conf(\diamond t'_i \rightarrow I_j)$$
$$= \frac{\text{number of transactions containing any item in } t'_i \text{ and } I_j}{\text{number of transactions containing any items in } t'_i} \tag{8}$$

The reason for the above formulation is given as follows. A transaction can be viewed as a customer behavior. In transaction t_i, there are the cross-selling effect between any marketing items I_a in t'_i and non-marketing items in set d_i. Let us consider some cases. If all items in t_i are being marketed, then there are no non-marketing items, and the profit gain is the difference between the profit of marketing items after marketing and that before marketing. If all items in t_i are not marketed, as there are no marketing items, in transaction t_i, there is no cross-selling effect from marketing items in transaction t_i. Thus, the profit gain due to marketing becomes zero. Now, consider the case of a transaction containing both marketing items and non-marketing items. Suppose the customer who purchases any marketing items in set t'_i always purchases non-marketing items I_b. This phenomenon is modelled by a gain rule $\diamond t'_i \rightarrow I_b$. The greater the confidence of these rules is, the greater the cross-selling effect is. That is, if this confidence is high, then when more of t'_i are sold, it means that very likely more of I_b will also be sold. [4]

5 Hill Climbing Approach

The ISM problem is likely to be very difficult. We propose here a hill climbing approach to tackle the problem. [5]

Let $f(S)$ be the function of the profit gain of the selection S of marketing items. Initially, we assign $S = \{\}$. Then, we will calculate $f(S \cup \{I_a\})$ for each item I_a. We choose the item I_b with the greatest value of $f(S \cup \{I_b\})$ and insert it into set S. The above process repeats for the remaining items whenever $f(S \cup \{I_b\}) > f(S)$.

[4] The rule $I \rightarrow \diamond d_i$ is called a *loss rule* in [17], because in [17], the problem is to determine a set of items to be discontinued from a store, d_i refers to some items to be removed, and it may cause some loss in profit from other items.

[5] We have also tried to apply the well-known optimization technique of quadratic programming. However, we could only approximate the problem by a quadratic programming problem and the approximation is not very accurate since we need to throw away terms in a Taylor's series which may not be insignificant. The resulting performance is not as good as the hill climbing method and hence are not shown.

5.1 Efficient Calculation of the Profit Gain

As the formula of the profit gain is computationally intensive, an efficient calculation of this formula is required. The hill climbing approach chooses the item with the greatest profit gain for each iteration. Suppose S now contains k items at the k-th iteration. At this iteration, we store the value of $f(S)$ in a variable f_S. At the $(k+1)$-th iteration, we can calculate $f(S \cup \{I_x\})$ from f_S efficiently for all $I_x \notin S$. Let $\mathcal{T}$ be the set of transactions containing item I_x and at least one item in selection set S. We can calculate $f(S \cup \{I_x\})$ as

$$f(S \cup \{I_x\}) \leftarrow f_S + g(I_x) - h(S, \mathcal{T}) + h(S \cup \{I_x\}, \mathcal{T}) \tag{9}$$

$$\text{where } g(I_x) = \textstyle\sum_{i=1}^{m} \left[(\alpha(\{I_x\}) - 1)prof(I_x, t_i) - \alpha(\{I_x\})cost(I_x, t_i)\right]$$

$$h(X, \mathcal{T}) = \sum_{t_i \in \mathcal{T}} \sum_{I_b \in d_i} (\alpha(t'_i) - 1)csfactor(t'_i, I_b)prof(I_b, t_i)$$

For $h(X, \mathcal{T})$ we assume all items in set X are selected for marketing, i.e. $t'_i = t_i \cap X$, and $d_i = t_i - t'_i$. Function $g(I_x)$ is the profit gain of marketing item I_x in all transactions. Function $h(X, \mathcal{T})$ is the profit gain of non-marketing items for the selection X in all transactions in set $\mathcal{T}$.

Let us consider the calculation of $f(S \cup \{I_x\})$. For $g(I_x)$, we need to add the profit gain of the newly added marketing item I_x after marketing (i.e. $g(I_x)$) to f_S. For the remaining parts, we only deal with the transactions in set $\mathcal{T}$. We need to subtract the profit gain of non-marketing items for the selection S in all the transactions in set $\mathcal{T}$ (i.e. $h(S, \mathcal{T})$) and then add the profit gain of non-marketing items for the new selection $S \cup \{I_x\}$ in all the transactions in set $\mathcal{T}$ (i.e. $h(S \cup \{I_x\}, \mathcal{T})$). As the set $\mathcal{T}$ is typically small compared with the whole database, we can save much computation by restricting the scope of search to $\mathcal{T}$. In the actual implementation the scope restriction is realized by a special search procedure of a special FP-tree as described below.

5.2 FP-tree Implementation

The transactions in the database are examined for computation whenever the confidence term $conf(\diamond t'_i \to I_j)$ is calculated. So, we need to do this operation effectively. If we actually scan the given database, which typically contains one record for each transaction, the computation will be very costly. Here we make use of the FP-tree structure [8].

We construct an FP-tree $\mathcal{FPT}$ once for all transactions, setting the support threshold to zero, and recording the occurrence count of itemsets at each tree node. With the zero threshold, $\mathcal{FPT}$ retains all information in the given set of transactions. Then we can traverse $\mathcal{FPT}$ instead of scan the original database. The advantage of $\mathcal{FPT}$ is that it forms a single path for transactions with repeated patterns. In many applications, there exist many transactions with the same pattern, especially when the number of transactions is large. These repeated patterns are processed only once with $\mathcal{FPT}$. By traversing $\mathcal{FPT}$ once, we can count the number of transactions containing any items in set t'_i and item I_b and number of transactions containing any items in set t'_i.

The details of the procedure can be found in the description of the function parseFPTree(N,D) in [17]. From our experiments this mechanism can greatly reduce the overall running time.

6 Empirical Study

We have used Pentium IV 2.2GHz PC to conduct our experiments. In our experiments, we study the resulting profit gain of marketing using the proposed algorithm. After the execution of our algorithm, there will be a number of selected marketing items. Let there be J resulting items (or marketing items). Note that J is not an input parameter. We compare our results with the naive approach of marketing by choosing J items with the greatest values of profit gain in Definition 3 in Section 3 as marketing items, assuming no cross-selling effect (i.e. $csfactor(t'_i, I_b) = 0$ for any set t_i and item I_b). This naive approach is called *direct marketing*.

6.1 Data Set

We adopted the data set from BMS WebView-1, which contains clickstream and purchase data collected by a web company and is part of the KDD-Cup 2000 data [11]. There are 59,602 transactions and 497 items. The average transaction size is 2.5. The profit of each item is generated similarly as [17].

6.2 Experimental Results

For the data set, we study two types of marketing method - discounted items and free items. For discounted items, the selling price is half of the original price. Free items are free of charge. As remarked in Section 3, we shall assume a uniform change in the sale volume for all marketing items, i.e. $\alpha(I_i) = \alpha$ for all items I_i. This set up is similar to that in [6]. For the real data set, the experimental results of profit gains and execution time against α for the situation of discount items are shown in Figure 1 and Figure 2. Those for the situation of free items are shown in Figure 3 and Figure 4. In the graphs showing the profit gains, we show the number of resulting marketing items next to each data point of the hill climbing method. This number is also the number of iterations in the hill climbing method.

In all the experiments, the profit gain for the hill climbing approach is always greater than that for direct marketing. This is because the proposed algorithm considers the cross-selling effect among items, but the direct marketing does not.

The execution time of direct marketing is roughly constant and is very small in all cases. For the hill climbing approach, the execution time increases significantly with the increase in α. This is explained by the fact that when α is increased, the marketing effect increases, meaning that the increase in sale of marketing items will be greater, which also increases the sale of non-marketing items by cross-selling effect. The combined increase in sale will be able to bring more items to be profitable for marketing since they can now counter the cost of marketing. This means that the hill climbing approach will have more iterations as α increases since the introduction of each marketing item requires one iteration, and this means longer execution time.

Note that in the scenario of free marketing items, direct marketing leads to zero or negative profit gain. This is because the items are free and generate no profit, and hence when compared to the profit before marketing, the profit gain is zero or negative. In the synthetic data, it is found that the gain is zero in most cases, since the marketing items are chosen to be those with no recorded transaction. The gain becomes negative for the real data set. Such results are similar to those for direct marketing in [6].

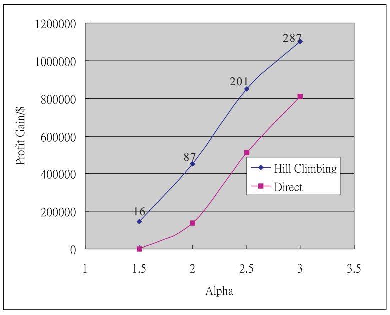

Fig. 1. Profit Gains against α for Real Data Set (Discount)

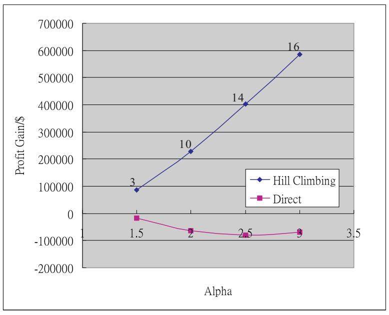

Fig. 3. Profit Gains against α for Real Data Set (Free)

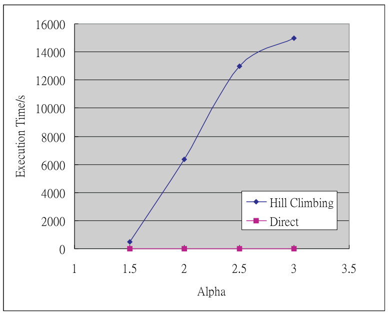

Fig. 2. Execution Time against α for Real Data Set (Discount)

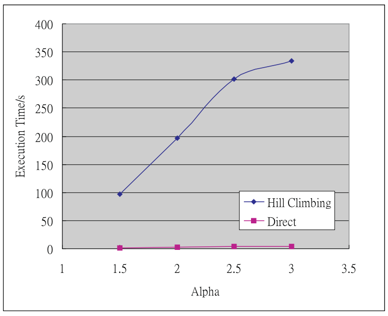

Fig. 4. Execution Time against α for Real Data Set (Free)

7 Conclusion

In this paper, we have formulated the problem Item Selection for Marketing (ISM) with the consideration of cross-selling effect among the items. We proved that a simple version of this problem is NP-hard. We adopt the concepts of association rules to the determination of the cross-selling factor. Then we propose a hill climbing approach to deal with this problem. We have conducted some experiments on both real data and synthetic data to compare our method with the results of a naive marketing method. The results show that our algorithm is highly effective and efficient.

Acknowledgements: This research was supported by the RGC Earmarked Research Grant CUHK 4179/01E.

References

1. R. Agrawal, S. Rajagopalan, R. Srikant, and Y. Xu. Mining newsgroups using networks arising from social behavior. In *Proc. of WWW*, 2003.
2. R. Agrawal and R. Srikant. Fast algorithms for mining association rules. In *Proc. of VLDB*, 1994.
3. T. Blischok. Every transaction tells a story. In *Chain Store Age Executive with Shopping Center Age, 71 (3)*, pages 50–57, 1995.
4. T. Brijs, B. Goethals, G. Swinnen, K. Vanhoof, and G. Wets. A data mining framework for optimal product selection in retail supermarket data: The generalized profset model. In *Proc. of SIGKDD*, 2000.
5. T. Brijs, G. Swinnen, K. Vanhoof, and G. Wets. Using association rules for product assortment decisions: A case study. In *Proc. of SIGKDD*, 1999.
6. P. Domingos and M. Richardson. Mining the network value of customers. In *Proc. of SIGKDD*, 2001.
7. M.R. Garey and D.S. Johnson. Computers and intractability: A guide to the theory of np-completeness. Freeman, 1979.
8. J. Han, J. Pei, and Y. Yin. Mining frequent patterns without candidate generation. In *Proc. of SIGMOD*, 2000.
9. D. Kempe, J. Kleinberg, and E. Tardos. Maximizing the spread of influence through a social network. In *Proc. of SIGKDD*, 2003.
10. J. Kleinberg, C. Papadimitriou, and P. Raghavan. A microeconomic view of data mining. In *Knowledge Discovery and Data Mining*, volume 2:4, pages 254–260. Kluwer Academic Publishers, 1998.
11. Ron Kohavi, Carla Brodley, Brian Frasca, Llew Mason, and Zijian Zheng. KDD-Cup 2000 organizers' report: Peeling the onion. *SIGKDD Explorations*, 2(2):86–98, 2000.
12. H. Mannila. Methods and problems in data mining. In *Proc. of Int. Conf. on Database Theory*, 1997.
13. H. Mannila, H. Toivonen, and A. I. Verkamo. Efficient algorithms for discovering association rules. In *Proc. of KDD*, 1994.
14. M. Richardson and P. Domingos. Mining knowledge-sharing sites for viral marketing. In *Proc. of SIGKDD*, 2002.
15. B.W. Taylor. Chapter 16: Inventory management. In *Introduction to Management Science, 7th Edition.* Prentice Hall, 2001.
16. K. Wang and M.Y. Su. Item selection by "hub-authority" profit ranking. In *Proc. of SIGKDD*, 2002.
17. R.C.W. Wong, A.W.C. Fu, and K.Wang. Mpis: Maximal-profit item selection with cross-selling considerations. In *ICDM*, 2003.
18. X. Wu, A. Tuzhilin, and J. Shavlik (Ed.). Third IEEE international conference on data mining, ICDM. 2003.