Behaviors and Profit based Sales Campaign Design

Shukai Li
Data Analytics Department
Institute for Infocomm Research
Singapore
lisk@i2r.a-star.edu.sg

Xiao-Li Li
Data Analytics Department
Institute for Infocomm Research
Singapore
xlli@i2r.a-star.edu.sg

Whye Loon Tung
SG Innovation Lab
The Nielsen Company
Singapore
whyeloon.tung@nielsen.com

Kelvin Sim
OneAnalytiks Research
OneAnalytiks Pte Ltd
Singapore
kelvinsim@oneanalytiks.com

***Abstract*—In this paper, we propose a novel profit-based metric for pattern mining. The objective is for companies to maximize their profits by offering the correct assortment of products with high margins in marketing campaigns. Based on the proposed metric, a sales campaign can accurately target/recommend the products that will translate to higher profits for a company to its customers. Moreover, if a sales campaign also considers consumer behaviour like what time consumers want what kinds of products, it can achieve better recommendations that not only benefit the company financially but also increase the satisfaction level of the customers. In addition, we analyze different groups of customers through a segmentation analysis to help a company target customers with significant upside potential. This research seeks to improve the marketing decision-making process to launch effective sales campaigns to target high-value customers.**

***Keywords*-sales campaign; profit shift; association rule analysis**

I. Introduction

Traditional marketing research generally relies on surveying consumers' preferences and conducting focus groups to gather their opinions and sentiments of identified products and services [1]. However, it is expensive and time consuming, if possible, to collect rich data on consumer behaviors by these traditional ways. Moreover, consumers may be reluctant to provide their real thoughts/opinions and personal information due to their privacy concerns. In recent years, we have witnessed the emergence of the "*big data*" era [2], which has a significant impact on marketing research. Particularly, retailers and credit card companies are moving towards data driven marketing research. They are now collecting consumer transaction data in massive volumes. Data mining and analytics approaches are subsequently employed to discover novel patterns or insights from the collected data that can in turn help to drive their marketing efforts.

In this paper, we focus on the problem of identifying consumer purchasing patterns that can drive higher profits in a company's marketing campaigns [3]. This problem can be addressed by our proposed three-step approach. Firstly, consumer purchasing patterns are mined automatically from the transaction data that was collected by a company. Secondly, only significant purchasing patterns that are expected to drive high *profit* margins are selected to be used for the sales campaigns. Finally, purchasing time predictions are performed to provide a practical guidance to the timeliness of the product recommendations in a sales campaigns so as to increase its success rate and marketing effectiveness.

For the first step, association rule mining [4] and sequential rule mining [5, 6] algorithms are two of the widely used approaches in mining consumer purchasing patterns due to their efficiencies in handling large amount of transaction data [7]. In addition, it is easy for a non-technical layman to understand and interpret the identified rules and their patterns as they are a natural and intuitive way of knowledge representation [8]. This helps to enhance the applicability of the rule-mining algorithms in business analytics. In general, both association rules and sequential rules represent the purchasing patterns that occur frequently in the transaction data [9], with the sequential patterns containing additional time (sequence or order) information. The detailed survey papers for association rule mining and its interesting metrics are available in [10] and [11] respectively.

It is well known in the data mining community that the main drawback of association rule mining is that too many rules or patterns (e.g. thousands or even tens of thousands of rules) can be extracted from a large transaction data set, depending on the parameter settings of the rule mining algorithms employed. Hence, it is extremely important to sieve out the useful patterns that are considered significant. To this end, several *pattern metrics* [12] have been proposed to rank the patterns in terms of their usefulness, such as *support*, *confidence*, and *lift* [13] etc. Different metrics may be useful for different applications or objectives. For example, relevant patterns that involve frequent card purchases can be used in a credit card company's marketing campaign to drive high transaction card usage rates. However, existing rule evaluation metrics are not suitable in meeting our defined business objective, as they do not account for the *profitability* of the patterns or the propensity of the patterns to drive sales. With respect to the goal of enhancing the profit gains of sales campaigns, it is critical to design a new metric to rank highly those patterns that can drive sales and profit margins. From the view of the credit card company, profit means a commission of every card transaction performed.

In this paper, we propose novel profit-based rule eval-

1082-3409/15 $31.00 © 2015 IEEE
DOI 10.1109/ICTAI.2015.61

uation metrics, $Profit_{rule}$ and $Profit_{imp}$, to rank the association rules and patterns according to their effectiveness in driving sales (and hence financial success) in marketing campaigns. To the best of our knowledge, this is the *first* metric of its kind designed to enhance the profitability of sales/marketing campaigns. Moreover, the new metric has been formulated as a function of the existing rule evaluation measures, and this enables us to comprehend how these conventional metrics may impact profits as well as study the inter-relationship amongst them.

We investigate the effectiveness of $Profit_{rule}$ by working with one of the largest credit card companies in China. We mine both the association and sequential rules from their customer credit card transactions to discover the purchasing patterns. These consumer purchasing patterns are ranked using both the proposed $Profit_{rule}$ and the existing rule evaluation measures. The top ranked association rules/patterns identified by the different metrics are then used in different sale campaigns. We evaluate the profitability of these patterns via their corresponding financial performance of the marketing campaigns. Our results demonstrated that marketing campaigns using patterns selected by our proposed $Profit_{rule}$ generated the highest profits, significantly outperforming the financial returns of campaigns backed by the other existing rule evaluation metrics.

In summary, our main contributions are ternary:

1) We have proposed novel profit-based pattern evaluation metrics, $Profit_{rule}$ and $Profit_{imp}$, to identify consumer purchasing patterns that can be leveraged to drive higher sales in marketing campaigns;
2) We showed the effectiveness of $Profit_{rule}$ via a credit card marketing campaign. We have observed that the marketing campaign using association rules selected by $Profit_{rule}$ is able to generate more profits when benchmarked against other marketing campaigns using consumer insights selected by the existing metrics;
3) In-between purchasing time can be estimated using the historical transaction data, and this information can guide the business units on when we should execute those rules to improve both the profits and customers' satisfaction of a sales campaign.

The organization of this paper is as follows: Section I first provides an introduction about pattern rule mining and its application in marketing. Section II then discusses the advantages and strategies of a sales campaign to help technical readers better understand the business ground. Section III describes the existing metrics used to evaluate the usefulness of the rules extracted from historical transaction records. Section IV gives the theoretical formulation of our proposed profit driven metric, $Profit_{rule}$, to evaluate the association rules for the purpose of driving sales. Section V looks into the task of estimating the time interval between consecutive purchases to provide an informed guide on the timeliness of product recommendations to customers in a sales campaign, while Section VI presents the empirical analyses of the results from the sales experiments that we have conducted in collaboration with a credit card company in China. Finally, Section VII concludes this paper.

II. Sales Campaigns

As our objective is to derive higher profits in a company's sales campaign, we first look at what constitutes a sales campaign. Sales campaigns are a kind of promotional strategies, also known as sales promotions [3]. Business owners use them to provide consumers or business customers with discounts or incentives to encourage them to purchase products or services.

A. Strategies

Sales campaigns use channels such as coupons, price reductions, buy-one-get-one-free promotions, contests, loyalty programs and point-of-purchase displays. For example, loyalty programs are designed to reward customers according to the volume of their purchases, and airline companies have been using loyalty programs for years to reward frequent fliers. Sales campaigns usually target specific groups of customers, which results in better cost control. Sales campaign marketers, for example, run coupons or incentive programs for specific stores. Samples may be limited to certain quantities to keep costs in line with certain budget parameter constraints.

B. Improving Response Rates of Sales Campaigns

Data mining facilitates the prediction of customer behaviours to perform product recommendations through the analysis of historical transaction records. Using data mining models, marketers can turn customer data into actionable insights to increase campaign profitability and improve the return on marketing investment [14]. Specifically, data mining helps to provide answers to four important questions: What kind of customers should a company target? What products should a company offer to its customers? What time is the best to contact a customer? When should a company make a promotional offer to a customer?

C. Driving Revenue From Existing Customer Base Via Cross-selling and Up-selling

The customer base of a company is referred to as the group of consumers who repeatedly purchase certain kind of products (goods or services) from the company. These customers often become a primary source of revenue for the company; and compared to the approach of increasing the business turn-over via the acquisition of new customers, concentrating on the existing customers to drive higher sales is considered the most direct and effective route to higher revenue. This kind of conversion is often less expensive than attracting new customers, as the existing customers are familiar with the product offerings of the company and are receptive to new/improved product launches.

Table I
TRANSACTION TABLE SAMPLE

Transaction ID	Customer ID	Category Number	Date
10182	1287	23	06-09-2012
10183	1287	61	30-09-2012
10185	753	52, 60	06-10-2012
10188	187	23	06-10-2012
10191	126	55	07-10-2012
10193	1102	28	15-10-2012
10198	187	61	23-10-2012
10199	1538	28	25-10-2012
10215	3726	35	10-11-2012
10228	753	42, 28	18-11-2012
10232	2325	23	25-11-2012
10236	256	32	26-12-2012
10237	2325	61	29-12-2012
10239	328	58	30-12-2012
10243	2325	61	31-12-2012

Note: Each row corresponds to a transaction record, with the first column being the transaction ID, and the second column denotes the customer ID. In addition, the coded category identifiers of the items purchased by the customers are stored as the third column, with the dates of the transactions recorded in the fourth column.

III. RULE EVALUATION MEASURES

As mentioned in Section I, the main drawback for association and sequential rule mining is that the process may generate too many rules which potentially reduce the usability of the results in real-world applications. In this section, we will present a number of traditional rule evaluation measures that can control the number of the rules generated from the given customer transaction data set. In order to illustrate the different probabilistic measures clearly, we first define a transaction space as follows.

Definition 1 (Transaction space): Given a transaction data set T, the sample space Ω is the set of all possible kinds of products or services. The σ-algebra $\mathcal{F}$ is the set of all considered items, and M is a set of functions. Particularly, if a transaction $t \in T$ contains an item i, the value of $M_i(t)$ is 1; otherwise, $M_i(t)$ is 0. Then the transaction space W is defined as

$$W = (\Omega, \mathcal{F}, M)$$

There are two disjoint itemsets X and Y, which satisfy $X, Y \subset \mathcal{F}$ and $X \cap Y = \emptyset$. Then a rule R is defined as $X \Rightarrow Y$. The total number of transactions in T is assumed to be $N = |T|$ where $|\cdot|$ denotes the cardinality of a set.

For an itemset $X \subseteq \mathcal{F}$, $S(X)$ is to denote the set of transactions containing X, i.e., $S(X) = \{t \in T | \forall i \in X, M_i(t) = 1\}$.

Definition 2 (Support of an itemset): The support metric $supp(X)$ of an itemset X is defined as the percentage of transactions in the database which contains X.

$$supp(X) = \frac{|S(X)|}{N}$$

We illustrate the concept of *support* for an item and a rule using Table I. For example, the support of itemset $\{23\}$ in Table I is $1/5$ since it appears in 3 transactions out of the total 15 transactions.

The *support* of a *rule*, $X \Rightarrow Y$ where X and Y are non-overlapping itemsets, is defined as the percentage of transactions in T that contains $X \cup Y$. The rule *support* thus determines how frequent the rule is applicable in the whole transaction set T. The *support* of the rule $X \rightarrow Y$, $support(X \Rightarrow Y)$, is defined as follows:

Definition 3 (Support of a rule):

$$support(X \Rightarrow Y) = \frac{|S(X \cup Y)|}{N}$$

where $X \bigcup Y$ is the union of the itemset X and itemset Y. Basically, from support of a rule $X \Rightarrow Y$, we can infer the percentage of the population this rule covers, so that we can know how general or how representative the rule is.

Definition 4 (Confidence of a rule): The confidence metric $conf(X \Rightarrow Y)$ of the rule $X \Rightarrow Y$ is defined as

$$conf(X \Rightarrow Y) = \frac{supp(X \Rightarrow Y)}{supp(X)}$$

For example, the confidence of the rule $\{23 \Rightarrow 61\}$ in Table I is $\frac{3/15}{3/15} = 1$. The confidence of the rule $X \Rightarrow Y$ means how much confidence we have if we want to infer if a person will buy item Y when we know he/she has bought an item X. The larger the confidence, the more certain we are about the inference. Hence, the confidence determines the predictability and reliability of a rule.

Definition 5 (Lift of a rule): The lift metric $lift(X \Rightarrow Y)$ of the rule $X \Rightarrow Y$ is defined as

$$lift(X \Rightarrow Y) = \frac{supp(X \Rightarrow Y)}{supp(X) \times supp(Y)}$$

For example, the lift of the rule $\{23 \Rightarrow 61\}$ in Table I is $\frac{3/15}{(3/15)\times(4/15)} = \frac{15}{4}$. *Lift* is a measure to evaluate the predicting performance of an association rule (target) as having an enhanced response with respect to the population as a whole (random). An association rule will predict accurately if the response is much better than the average for the population as a whole. Lift is defined as the ratio of the target response divided by average response.

IV. A NOVEL PROFIT-BASED RULE EVALUATION METRIC

The three rule evaluation metrics introduced in the last section is definitely useful to control the number of the rules generated from the given transaction set T. In other words, we can mine those rules with their *support*, *confidence*, and *lift* greater than or equal to the user-specified minimum support, minimum confidence and minimum lift values. However, they do not sufficiently consider the profit factors that are crucial in the business world and thus they cannot be better leveraged in sales campaigns. As such, we need to design a novel evaluation metric for our purpose. In particular, given an association rule $X \Rightarrow Y$, we calculate its profit value, based on difference between its revenue and cost:

$$Profit_{rule} = Revenue_{rule} - Cost_{rule} \quad (1)$$

Given the rule $X \Rightarrow Y$, we want to evaluate if the rule can make significant profits by calculating the difference between its *rule profit* and the *random profit*. Let nc be the number of customers that are covered by the rule (or they have bought the products in the rule condition X). Let nr be the number of customers who buy the product Y after purchasing X. Clearly, $nr \leq nc$.

First of all, if we employ the rule for a sales campaign, then its *rule revenue*, *rule cost* and *rule profit* become

$$Revenue_{rule} = nr * arv \quad (2)$$

$$Cost_{rule} = nc * ac \quad (3)$$

$$\begin{aligned} Profit_{rule} &= Revenue_{rule} - Cost_{rule} \\ &= nr * arv - nc * ac \end{aligned} \quad (4)$$

where ac and arv is the average cost and average revenue per customer for the promotional strategy respectively.

Particular, the campaign cost for the rule, $nc * ac$, is the total cost of all the customers covered by the rule. The campaign revenue for the rule, $nr * arv$, is the total revenue of all the customers who buy both product Y and product X together.

As $conf(X \Rightarrow Y) = \frac{supp(X \Rightarrow Y)}{supp(X)} = \frac{nr/N}{nc/N} = \frac{nr}{nc}$, the profit becomes:

$$\begin{aligned} Profit_{rule} &= nr * arv - nc * ac \\ &= nc * (conf(X \Rightarrow Y) * arv - ac). \end{aligned} \quad (5)$$

Secondly, we may calculate the *random profit* for a promotional strategy without using any association rule. *Random profit* can be estimated as follows:

$$\begin{aligned} Profit_{norule} &= nr * arv - nc * ac \\ &= nc * p(Y) * arv - nc * ac. \end{aligned} \quad (6)$$

Note $nr * arv = nc * p(Y)$ where $p(Y)$ is the probability of customers buying Y which is equal to $supp(Y)$. Hence equation (6) becomes

$$Profit_{norule} = nc * supp(Y) * arv - nc * ac. \quad (7)$$

Finally, we consider the difference (or *improved profit*) between the above two profit values, namely *rule profit* and *random profit*:

$$\begin{aligned} Profit_{imp} &= Profit_{rule} - Profit_{norule} \\ &= (nc * conf(X \Rightarrow Y) - nc * supp(Y)) * arv \end{aligned} \quad (8)$$

It is obvious that

$$\begin{aligned} conf(X \Rightarrow Y) &= \frac{P(X \cup Y)}{P(X)} = \frac{P(Y)P(X \cup Y)}{P(Y)P(X)} \\ &= supp(Y) * lift(X \Rightarrow Y), \end{aligned} \quad (9)$$

so the equation (8) becomes

$$\begin{aligned} Profit_{imp} =& nc * supp(Y) * lift(X \Rightarrow Y) * arv \\ & - nc * supp(Y) * arv \\ =& nc * arv * supp(Y) * (lift(X \Rightarrow Y) - 1). \end{aligned} \quad (10)$$

Now we can observe from equation (10) that for a given rule $X \Rightarrow Y$, its *profit improvement* will be determined by three major factors, namely nc, $supp(Y)$ and $lift$ values, as arv is a constant. Particularly, bigger nc will guarantee that this rule is generic and can be applied to large amount of customers. Bigger $supp(Y)$, on the other hand, makes sure the rule's prediction is more accurate — once we applied it to make recommendation, there is a high chance that the customers will buy the recommended products. Moreover, the $lift$ is more important than nc and $supp(Y)$, because it decides whether the rule profit is increased or decreased compared with random profit. In other words, if $lift > 1$, it will bring positive benefits to the promotion strategy and thus *improved profit* $Profit_{imp}$ is positive; Otherwise it will lead to the negative impacts on the campaigns and we will not use this rule. Note that nc and $supp(Y)$ will influence the scale of $Profit_{imp}$ and thus also play important role to decide if the current rule is good enough to be selected for the promotional strategy.

In general, we will compute the *rule profit* $Profit_{rule}$ and *improved profit* $Profit_{imp}$ for all the rules mined from transaction data and select the rules with highest values $Profit_{rule}$(with $Profit_{imp} > 0$) for sales campaigns. The number of the rules selected will be determined by the total budget available for sales campaigns.

V. Sequential Patterns Mining with Estimated Purchasing Time

In the last two sections, we have discussed some existing metrics and our proposed profit based metric for rule selection. Basically, we will first mine sequential patterns from given transaction data and subsequently select a subset of rules with higher profit scores that can be leveraged for the sales campaigns. In other words, we will choose those customers covered by the rule "antecedent" (i.e. satisfy the rule conditions) and recommend the rule "consequent" (i.e. products in the rule conclusions). We use an example in Table I to illustrate this. Given a sequential patten $(23 \rightarrow 61)$, if a consumer has purchased a product of the category 23, we can then recommend the consumer the products of the category 61, since he/she is likely to buy it and we can thus make profit in our sales campaigns.

One important problem that we have yet to address is that *when* we should recommend the products in the rule to the customers? Clearly, this can provide a practical guidance for companies to recommend products in the right time (or predict when customers are likely to buy the products), to increase its success rate and marketing effectiveness. We can address this purchasing time prediction problem based on the historical transactions.

More specifically, for a sequential patten $(23 \to 61)$, we would like to know, on average, what is the usual time interval/difference between the purchasing time of products in 23 (or more generically, the time to purchase the last item in the antecedent of the sequential pattern) and products in category 61 (the first item of the consequent of the sequential pattern). We can choose *median*, instead of *average*, as the predicting purchasing time, as it is a better indication of central tendency and is less sensitive to outlier time values (e.g. exceptionally large values/long time intervals in the data) that could distort the prediction.

For example, the sequential pattern $(23 \to 61)$ in Table I has occurred three times. For customer $\{1287\}$, *event* (23) (i.e. buying products in 23) occurred on 06-09-2012 and *event* (61) occurred on 30-09-2012. Thus, its time interval is 24 days, indicating that this customer bought the products in (61), 24 days after purchasing the products in (23). As for customer $\{187\}$, *event* (23) occurred on 06-10-2012 and *event* (61) occurred on 23-10-2012, with time interval 17 days. Similarly, for customer $\{2325\}$, *event* (23) occurred on 25-11-2012 and *event* (61) occurred on 29-12-2012. Thus, its time interval is 34 days. Therefore, the median of the time intervals among the three events is 24 days, i.e. the median of 17 days, 24 days and 34 days. By this way, we can enhance the original sequential patten $(23 \to 61)$ with median time interval 24 days, i.e. $(23 \to 61, [24])$, which can be used as a guidance to indicate when we should apply the rule for product recommendation.

Note that we only consider the smallest interval between events in our calculation. For example, for customer $\{2325\}$, event (23) occurred on 25-11-2012 and event (61) occurred on both 29-12-2012 and 31-12-2012. We only consider the earliest value (i.e. 29-12-2012) and thus the smallest time interval between the two events, which is 34 days instead of 36 days. The reason is that we hope we can do the timely recommendation.

With sequential rule $(23 \to 61, [24])$ in targeted marketing, if a consumer purchased products in category 23, we can wait for about 24 days before recommending the consumer those products in the category 61. However, in practice we may not be able to contact many customers everyday based on our recommended median time interval, say $k-th$ day, obtained from large transaction data. Instead, there could be some fixed dates (e.g. in the 10th, 15th, 30th days every month) which are used to contact the customers.

Without loss of generality, given a sequential rule, we can compute all the time intervals between the purchasing time of the rule antecedent and the rule consequent across all the customers from the given transaction data. Assume that a customer is contacted at the day x, the probability to be the best contact time can be computed as

$$P(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-(x-k)^2/2\sigma^2} \tag{11}$$

where σ is standard variance of all the time intervals.

We can compute the absolute time interval value m between the x's nearest fixed contact day and k. Then the confidence level of this time x to contact becomes

$$\begin{aligned} Conf_{level} &= P\left(\frac{x-k}{\sigma/\sqrt{n}} < \frac{m}{\sigma/\sqrt{n}}\right) \\ &= P\left(x-k < m\right) \end{aligned} \tag{12}$$

Here, n is the number of all the time intervals for the rule. The higher the confidence level, we are more confident that x is the better time to contact customers. It is quite important to contact a proper customer in a proper time, which means the consumer behaviour has been understood.

The time complexity of our proposed methodologies is $O(n)$, so they will scale linearly with the number of transactions that the methodologies process and deal with large datasets.

VI. FIELD EXPERIMENTS

To evaluate the effectiveness of the proposed profit-based rule evaluation metric, $Profit_{rule}$, we analysed the customer credit card transaction records from one of the largest credit card companies in China. The transaction data has been recorded as a transaction table, and its format is as shown in Table I. The transaction data was observed over a time span of two years from June 2011 to May 2013, which consists of a one-year pre-experimental and a one-year experimental period, respectively. 1248 customers are chosen for the campaigns with their corresponding 367532 transaction records.

Demographic information of the selected customers such as their age, gender, income level, place of residence and so on, are retrieved from their credit card application forms. There are also some important customer-centric characteristics such as the purchase frequency and average amount spent etc. Of the entire data set, 416 customers (one-third of total customers) and their records are assigned to a control group while the remaining 832 customers (two-thirds of 1248) and their transactions are equally divided into two test groups referred to *Test Group A* and *Test Group B*(416 per group), respectively. Due diligence was performed to ensure the distributions of the demographics and customer-centric attributes are similar in the control and the two test groups during the customer assignment process. For instance, if the control group has 30% of the customers with income level between \$5000 and \$10000, then the two test groups will also have approximately 30% of the assigned customers in the same income range.

We have engaged the 1248 selected customers through direct mail campaigns, which is the most straightforward and targeted approach in meeting our experimental objective. Specifically, direct mails are intended to arrive at the doorsteps of consumers who have been identified to be great candidates for the recommended products and services. This contrasts against other broad methods of advertising that are designed to have wide and indiscriminating coverage, often addressing an entire population of consumers within a given geographic area.

All the 1248 identified customers in the study use the credit cards issued by the same bank in the same Chinese city. The field experiments are conducted based on the 67 predefined product categories, and the most popular product of each category is selected to represent the corresponding category in the respective sales/marketing campaigns. In the study, the product recommendation leaflets are distributed to the selected customers every 15 days. Three different types of campaigns were conducted in the study. They are *Profit Driven Campaign* (PDC), *Profit Driven Campaign with Time information* (PDCT) and the traditional *Confidence Driven Campaign* (CDC), which, are run during the experimental period of one year from June 2012 to May 2013. The pre-experimental period refers to the one year time frame prior to the commencement of the experiment and spanned from June 2011 through May 2012.

Specifically, CDC was run on the customers in the control group, while PDC and PDCT are evaluated on customers assigned to Test group A and Test group B, respectively. In the CDC design, the rules with the highest confidence are selected to drive the campaign as it focuses on feedback rates. However, in practice, high feedback rates do not necessarily translate to high profits as the targeted customers may be buying low-margin products. To enable a fair and unbiased comparison, all the three groups of customers received the promotional leaflets in the same time period.

A. Sales Campaigns Evaluation

We study and evaluate the effects of the three different campaigns across the control and the two test groups of customers by comparing and contrasting the means of a set of financial and relational metrics for the pre-experimental period before the launch of the target campaigns (June 2011 to May 2012) and those of the experimental period when the campaigns are in place (June 2012 to May 2013). The mean values of the various evaluation metrics for the control and test groups for the two periods have been tabulated as Table II. Values in parentheses represent the readings for the pre-experimental period. In the pre-experimental period, all the customers in the three groups received the same standard and un-optimized product recommendations.

The set of financial and relational metrics selected to evaluate the effects of the three campaigns are described as follows. The financial metrics captured the capital outlays on the various sales campaigns and the monetary gains the customers brought to the credit card company. This consists of *revenue* from the amount of card transactions performed, *investment* in the form of costs in conducting a particular campaign, and *return-on-investment* (ROI) as the ratio of $(revenue - investment)/investment$.

The relational metrics, on the other hand, are designed to quantify the qualitative information about how a customer feels towards a campaign, which includes how good the campaign is in meeting his/her needs, how much the customer plans to spend on buying the recommended products, and how likely the customer would recommend the suggested products to others. These indicators are selected based on previous research into customer relationship management [15, 16]. The relational metrics are measured with a ten-point interval scale anchored from "totally agree" (10) and "strongly agree" (8) to "strongly disagree"(2) and "totally disagree" (1). The reported financial metric readings have been scaled by an arbitrarily chosen constant to mask the actual values due to confidential consideration.

B. Financial Metrics Analysis

To assess if there is any significant difference in the financial characteristics of the customers in the three created groups (i.e., the control and the two test groups), Hotelling's T-square test is employed to evaluate the three financial metrics across the three customer groups. The results showed that for the pre-experimental period of June 2011 to May 2012, the customers in the three groups are similar in all the three financial metrics, eliminating any concerns of biasness in our evaluation of the effectiveness of the three different sales campaigns.

Next, t-test with Bonferroni adjustment is used to analyze the difference of each financial metric in the three groups across the pre-experimental and experimental periods. Specifically, we observe that the revenue and ROI metrics increased significantly (at $\alpha < 0.1$) in all the three groups across the two periods. This indicated that the three designed sales campaigns are able to identify the customers who are likely to purchase the recommended items in the selected product categories, leading to higher revenue and hence increased ROI.

To further ascertain which of the three campaigns is the most effective in improving profits, T-test with Bonferroni adjustment is used to perform the inter-group comparison where the mean values of the financial metrics for the control group are used as base levels. The field experimental results indicated that both the PDC and PDCT campaigns achieved significant improvements on the revenue and ROI metrics. This demonstrated that our proposed profit-based rule evaluation metric is more effective than the traditional confidence measure in identifying purchasing patterns that lead to higher profits. The confidence level of the inter-group comparison is computed as shown in Table II during the

Table II
COMPARISON FOR TEST AND CONTROL GROUPS

Metrics	Control Group: CDC	Test Group A: PDC	Test Group B: PDCT
Financial Metrics			
Revenue (¥)	34897 (32550)	38153 (32253)*	40892 (32762)***
Investment (¥)	12539 (14724)	11916(14537)	11729 (14338)
ROI (¥)	1.78 (1.21)	2.20 (1.22)*	2.49 (1.28)***
Relational Metrics			
Understanding customer's needs	7.15(5.37)	7.18 (5.34)	7.61 (5.28)**
Planning to buy the recommendations	6.83(5.33)	6.81 (5.21)	7.31(5.28)***
Recommending to others	6.95(5.14)	6.98 (5.23)	7.37 (5.19)*

Note: the financial values have been scaled by an arbitrary constant for confidential consideration.
*: Confidence level 97.5% **: Confidence level 99.0% ***: Confidence level 99.5%

experimental period. Furthermore, comparing the results of PDC and PDCT, we find that PDCT is significantly better than PCD in the reported revenue value and hence ROI (confidence level: 97.5%), suggesting there is notable impact in predicting the purchase time and contacting customers only when they are expected to make a purchase.

C. Relational Metrics Analysis

We conducted the analysis of the reported relational metrics using the same approach as the financial metrics. Hotelling's T-square test is employed to determine if there is any significant difference in the customers across the three created groupings in the pre-experimental period. The test result indicated there is no evidence of any notable differences across the three customer groups.

Subsequently, applying a T-test with Bonferroni adjustment revealed that in the experimental period, customers assigned to Test Group B and hence exposed to the PDCT campaign generally reported an improvement across all the relational metrics compared to the pre-experimental period. In contrast, no such significant improvement is observed for the customers in Test Group A. This clearly demonstrated that customer satisfaction and endearment was acquired only when the credit card company targeted customers who want to make a purchase with the list of recommended products. That is, customers in the Test Group B felt that the credit card company understood their needs and they appreciated the recommendations mailed to them.

Hence, the ablity to factor in purchase time estimation when making recommendations to the customers enabled the PDCT sales campaign to achieve improvements in both derived profits and customer endorsement.

D. Segmentation Analysis

Next, we performed a simple segmentation analysis on the customers in the various groups (i.e., the control and the two test groups) to further understand the profiles of the customers who responded to the respective sales campaigns. Segmentation analysis is a powerful tool that is often employed by a business to identify its most valuable customers. In segmentation analysis, an entire population of customers is partitioned into several subgroups or segments using a set of pre-defined factors. In this paper, three binary factors, namely, income level (high versus low), age (young versus old) and gender (male versus female), are employed to individually segment the customers in each of the three defined groups into two different subsets. The objective is to study how the customers in each subset of a group responded to the corresponding sales campaign by analysing the change in derived profits observed by the credit card company across the pre-experimental and experiment periods (see Figure 1).

With respect to Figure 1(a), one can observe that for Test Group A and Test Group B, there is a significant positive increase in the derived profits from the high-income customers as compared to their low-income counterparts over the pre-experimental and experimental periods. Specifically, the gap in derived profits from the high-income versus low-income customers is the most distinct for Test Group B (on which the PDCT campaign was conducted). In contrast, there is negligible difference in the derived profits across the two periods for the two income ranges in the control group. Thus, this suggested that higher-income customers are more sophisticated consumers and they expect to receive promotional leaflets only when they have an intention to make a purchase. Should that happens, the response rate to a sales campaign is expected to be favorable.

Furthermore, Figure 1(b) showed that younger customers in the two test groups are more receptive and responsive to the PDC and PDCT campaigns, thus contributing to a significant increase in derived profits across the pre-experimental and experimental periods. This behavior is distinctly different from the one exhibited by the older customers, who tend to be more conservative and hence less impulsive buyers.

Finally, the investigation into gender differences revealed that female customers have better response rates towards the PDC and PDCT campaigns than their male counterparts in Test Group A and Test Group B, respectively. As the recommended products in these two campaigns are offered in order to derive higher profits, the items are also generally more expensive. The results illustrated that female customers may be less sensitive to the pricing of the recommended products and hence are more willing spenders. The observations from this segmentation analysis suggested that in addition to the recommendation of carefully selected products to drive profits, the credit card company can also benefit financially by targeting sales campaigns at younger female customers with a high level of disposable income as this group of customers is considered high-value customers.

VII. Conclusions

In this paper, we proposed a novel rule evaluation metric for pattern mining to drive sales campaigns. Based on the new metric, a sales campaign can recommend carefully selected products to customers to derive higher profits. Furthermore, it has been shown that if a sales campaign considers the time interval of purchases, the response rate of the customers will improve in addition to better customer satisfaction, since the consumer behaviour has been analyzed and considered in the sales campaign design. Lastly, we applied the segmentation analysis to identify the customer profiles that are deemed to be receptive to a sales campaign and hence likely to purchase the recommended products.

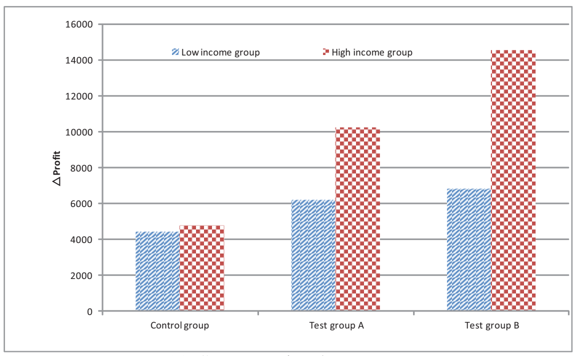

(a) Segmentation by Income

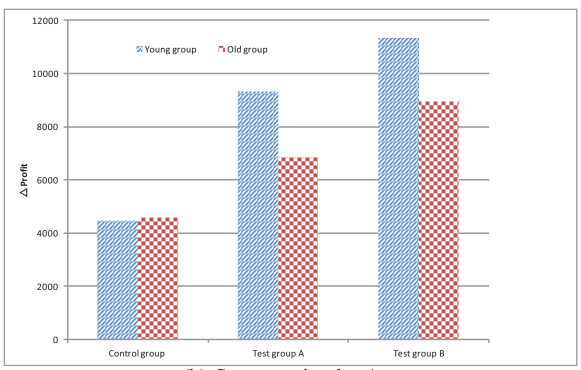

(b) Segmentation by Age

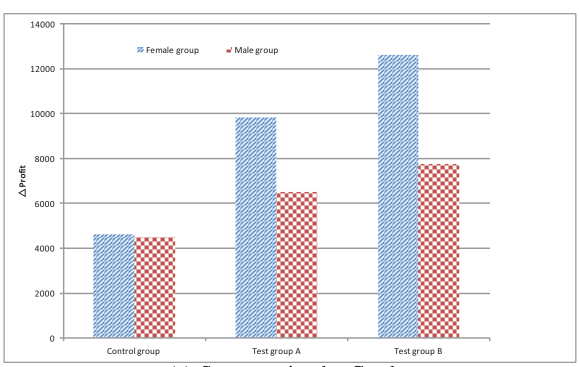

(c) Segmentation by Gender

Figure 1. Segmentation Analysis

References

[1] E. Lee, M. Y. Hu, and R. S. Toh, "Are consumer survey results distorted? systematic impact of behavioral frequency and duration on survey response errors," *Journal of Marketing Research*, vol. 37, no. 1, pp. 125–133, 2000.

[2] X. Wu, X. Zhu, G. Wu, and W. Ding, "Data mining with big data," *IEEE Trans. Knowl. Data Eng.*, vol. 26, no. 1, pp. 97–107, 2014.

[3] R. Venkatesan and P. W. Farris, "Measuring and managing returns from retailer-customized coupon campaigns," *Journal of Marketing*, vol. 76, no. 1, pp. 76–94, 2012.

[4] R. Agrawal, T. Imielinski, and A. Swami, "Mining association rules between sets of items in large databases," in *Proceedings of the ACM SIGMOD International Conference on Management of Data*, 1993, pp. 207–216.

[5] S. K. Harms and J. S. Deogun, "Sequential association rule mining with time lags," *Journal of Intelligent Information Systems*, vol. 22, no. 1, pp. 7–22, 2004.

[6] Y. Zhao, H. Zhang, L. Cao, C. Zhang, and H. Bohlscheid, "Mining both positive and negative impact-oriented sequential rules from transactional data," in *Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD)*, 2009, pp. 656–663.

[7] S.-H. Liao and C.-H. Wen, "Mining demand chain knowledge for new product development and marketing," *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, vol. 39, no. 2, pp. 223–227, March 2009.

[8] X.-L. Li and B. Liu, "Rule-based classification," *Data Classification: Algorithms and Applications, Charu Aggarwal (Eds.), CRC Press.*, 2013.

[9] C. Rudin, B. Letham, and D. Madigan, "Learning theory analysis for association rules and sequential event prediction," *Journal of Machine Learning Research*, vol. 14, no. 1, pp. 3441–3492, 2013.

[10] J. Hipp, U. Güntzer, and G. Nakhaeizadeh, "Algorithms for association rule mining - A general survey and comparison," *SIGKDD Explorations*, vol. 2, no. 1, pp. 58–64, 2000.

[11] Y. Zhang, L. Zhang, G. Nie, and Y. Shi, "A survey of interestingness measures for association rules," in *Business Intelligence: Artificial Intelligence in Business, Industry and Engineering, Proceedings of the Second International Conference on Business Intelligence and Financial Engineering*, 2009, pp. 460–463.

[12] E. R. Omiecinski, "Alternative interest measures for mining associations in databases," *IEEE Trans. on Knowledge and Data Engineering*, vol. 15, no. 1, pp. 57–69, Jan. 2003.

[13] S. Brin, R. Motwani, J. D. Ullman, and S. Tsur, "Dynamic itemset counting and implication rules for market basket data," in *Proceedings of the ACM SIGMOD International Conference on Management of Data*, 1997, pp. 255–264.

[14] G. Cui, M. L. Wong, and H.-K. Lui, "Machine learning for direct marketing response models: Bayesian networks with evolutionary programming," *Management Science*, vol. 52, no. 4, pp. 597–612, 2006.

[15] V. Kumar, J. A. Petersen, and R. P. Leone, "How valuable is word of mouth?" *Harvard Business Review*, vol. 85, no. 10, pp. 139–146, 2007.

[16] R. T. Rust, K. N. Lemon, and V. A. Zeithaml, "Return on marketing: Using customer equity to focus marketing strategy," *Journal of Marketing*, vol. 68, no. 1, pp. 109–127, 2004.