Predictive Modeling of Campaigns to Quantify Performance in Fashion Retail Industry

Chandadevi Giri
Email: chandadevi.giri@hb.se

Dep. of Business Administration and Textile Management,
University of Boras,
Automation, Computer Engineering Signal & Image Analysis GEMTEX, ENSAIT
University of Lille
Dep. of Textile Engineering College of Textile and Clothing Engineering
Soochow University

Ulf Johansson
Email: ulf.johansson@ju.se

Dept. of Computer Science and Informatics, Jönköping University

Tuwe Löfström
Email: tuwe.lofstrom@ju.se

Dept. of Computer Science and Informatics, Jönköping University,

***Abstract*— Managing campaigns and promotions effectively is vital for the fashion retail industry. While retailers invest a lot of money in campaigns, customer retention is often very low. At innovative retailers, data-driven methods, aimed at understanding and ultimately optimizing campaigns are introduced. In this application paper, machine learning techniques are employed to analyze data about campaigns and promotions from a leading Swedish e-retailer. More specifically, predictive modeling is used to forecast the profitability and activation of campaigns using different kinds of promotions. In the empirical investigation, regression models are generated to estimate the profitability, and classification models are used to predict the overall success of the campaigns. In both cases, random forests are compared to individual tree models. As expected, the more complex ensembles are more accurate, but the usage of interpretable tree models makes it possible to analyze the underlying relationships, simply by inspecting the trees. In conclusion, the accuracy of the predictive models must be deemed high enough to make these data-driven methods attractive.**

Keywords—Fashion retail, Campaign Prediction, Machine Learning, Predictive Modeling, Decision Trees, Random Forest

I. Introduction

Promotional campaigns have gained popularity with the growing multi-channel consumer interaction. With this in mind, it is not surprising to see the importance put on understanding promotional campaigns in the marketing literature. Both quantitative and qualitative works have been conducted in this domain, with applications in different retail industries. The background of this research is briefly outlined in [1], [2]. It has been observed from many studies that retailers segment their customers and offer different promotions in personalized ways to achieve maximum retention. Some papers e.g., [3][4], investigate scenarios of targeting faithful customers, while [5] suggests that long-term customers should be offered lower prices than newer. Dynamic pricing has been adopted by the industry that uses data from their databases to optimize the prices of the product based on market demands [6]. Several online fashion retailers offer promotions based on subscriptions and referrals [7], [8]. Conditional promotions are also popular, for example, a consumer will be offered a 1€ discount on the purchase of a minimum of 5€. This was analyzed in [9], where it was concluded that conditional campaigns were successful in driving unplanned shopping. Technology has also enabled RFID-based systems which could help to understand the important elements in retailing [10], [11]. This has allowed the industry to store information about the consumer buying patterns and behavior in a better way. Many studies have been conducted to model consumer behavior that relies on the pricing factor [12][13]. Artificial intelligence [14] and machine learning [15] have gained importance in the fashion and clothing industry mainly focusing on the manufacturing sector and for improving supply chain management. However, the industry could benefit from using data analytics in several other scenarios. In this paper, we present a data-driven campaign prediction model to identify the profitability and success rate of campaigns.

II. Problem Statement and Research Aim

Attracting and retaining customers is challenging for most industries. However, for the fashion retailers, it is even harder, since there are many influencing factors for customers in the digital world. So, continuously working for retaining customers is vital for these companies, to gain or keep a competitive advantage. The promotional campaigns help fashion retailers to retain customers by offering attractive discounts or services like free delivery, free returns and gifts. Promotional campaigns are, however, very expensive for the retailers, in particular if the conversion rates, and consequently the revenues, are low. Therefore, fashion retailers need to have an appropriate and well-defined strategy for designing their campaigns. In practice, it is fair to say that the fashion retail industry still struggles to even identify the parameters driving customer attention. In this paper, we propose a data-driven method to identify the profitability and ultimately success rate of different campaigns. The suggested method uses predictive modeling where two different models are created:

1. **Profitability**: This regression model predicts the average profit from a campaign.
2. **Success**: This classification model predicts the overall success of a campaign

978-1-7281-0858-2/19/$31.00 ©2019 IEEE

Fig. 1 below describes the business problem of fashion retailers and states some possible solutions investigated here. In the empirical study, we present a case study on 826 promotional campaigns from a leading Swedish fashion retailer. In more detail, we study campaign data features and use machine learning to model *profitability* and *success*. The results of this study are promising enough to argue that fashion retailers could use this or similar data-driven methods to optimize campaigns. This could also help them significantly in deciding the discount levels (e.g., 10%, 20%, etc.) and included services, e.g., free delivery, free return, etc.

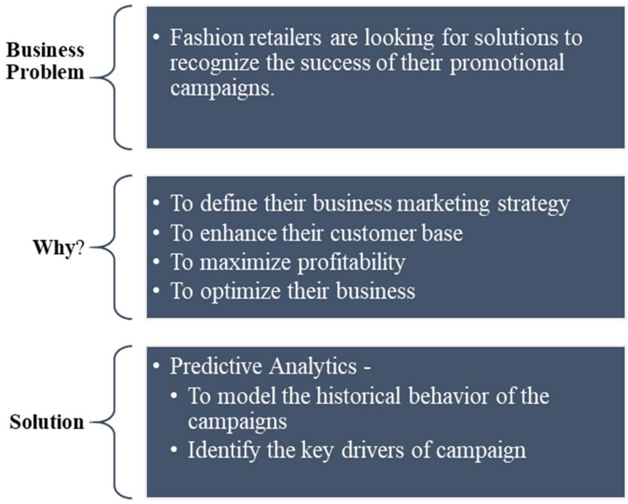

Fig. 1. Overview of Research Problem and Objective

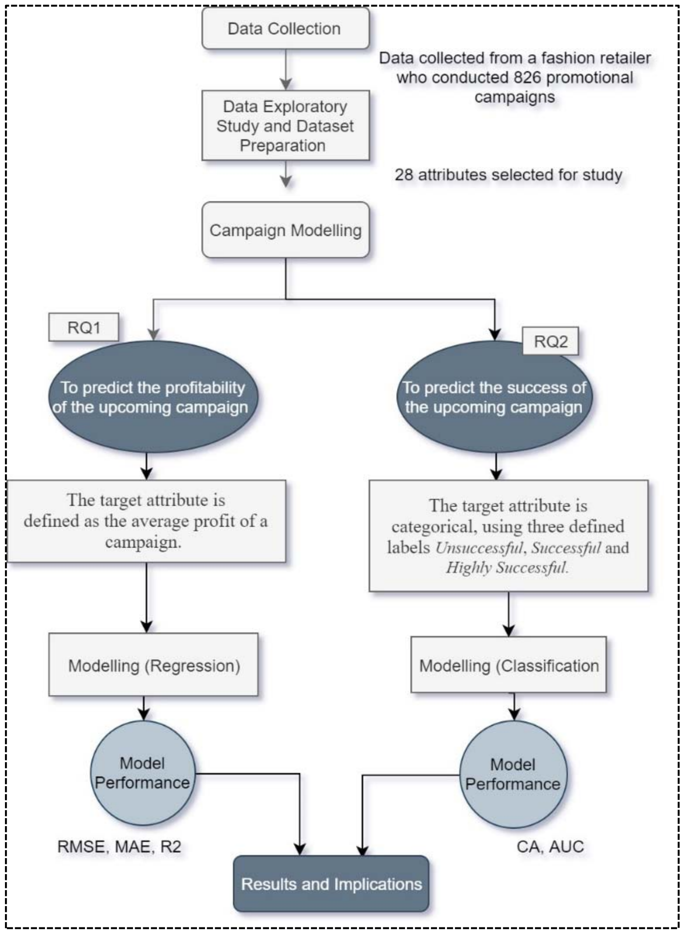

Fig. 2. Research Framework

The research framework used in this study is explained in Fig. 2, which outlines the research methodology of the work, describing the steps from data collection, data analysis and modeling. A detailed description of these steps is given in the following sections.

The succeeding sections are organized like this: First section III describes the methodology and the model evaluation metrics used in the research work. Section IV briefly discusses the steps employed to formulate the models, with the subsequent subsections discussing data collection, exploratory study, creating new features and developing models. Section V discusses the prediction results of the models. Section VI concludes this study.

III. PREDICTIVE MODELING AND EVALUATION

This section describes the regression and classification models and the evaluation metrics used in this work.

A. Predictive Models

Two different kinds of predictive models are used for modeling the campaign behavior in the experimentation described below:

Decision trees is a tree-like representation of a model that is formed by the set of rules splitting at the nodes, and making predictions in the leaves. The most popular decision tree models are CART[16] and C4.5/C5.0 [17]. The main advantage of decision trees over other machine learning modeling techniques is their interpretability, thus allowing explanations of individual predictions and inspection of overall underlying relationships.

Random forest [18], is an ensemble learning technique that can be used to develop both classification and regression models. While random forests consist of decision trees (or regression trees), these base models (called *random trees*) differ slightly from their standard counterparts. More specifically, in order to introduce the necessary diversity, bootstrap sampling is used, and the attributes available when optimizing a split are restricted to a randomized subset of the attributes.

B. Model Evaluation Metrics

In the empirical investigation, several evaluation metrics are used.

Regression

- RMSE is defined as the square root of the averaged squared errors.
- MAE is the mean absolute error of the predictions.
- R^2 is the proportion of the variance in the explained variable explained by the model.

Classification

- Classification accuracy indicates the proportion of accurately classified examples.
- Area under ROC curve is a measure of the ordering capability of the classifier.

IV. METHOD

This section briefs the steps followed in conducting the research aligned with the campaign prediction as explained in the research framework (Fig. 2). This section is divided into three parts, beginning with the data collection, then exploring the campaign data attributes followed by modeling.

TABLE I. CAMPAIGN DATA ATTRIBUTES

No.	Swedish	Description
1	check	True when campaign type was check
2	firstLine	True when campaign type was first line
3	firstAndSecondLine	True when campaign type was first and second line
4	combDiscount	True when campaign type was combination
5	allOrder	True when campaign type was all order
6	ladder	True when campaign type was ladder
7	threeForTwo	True when campaign type was three for two
8	discCheck	The percent discount when payed by internal check. Can only occur together with check campaigns
9	discFirst	The percent discount on the first item. Firstline, firstandsecondline, combDiscount, allorder and ladder must have a value for DisFirst
10	discRest	The percent discount on the second item. CombDiscount, and ladder must have a value for DiscRest. Firstandsecondline can have a DisSecond.
11	marketingDiscount	The average amount discount received per order (in SEK) for the entire campaign
12	freeGift	True when offered a free gift
13	payedGift	True when offered a payed gift
14	freeShipping	True when offered free shipping
15	freeExressShipping	True when offered free express shipping
16	freeReturn	True when offered free return
17	reqSale	True when the campaign required sale items
18	reqReducedPrice	True when the campaign required items with reduced price
19	reqBrandSelection	True when the campaign required items from specific brands
20	reqValue	True when the campaign required a minimum value
21	req#Items	True when the campaign required a minimum number of items
22	reqTime	True when the campaign required to be used within a limited time
23	reqOrdinaryPrice	True when the campaign required items with ordinary prices
24	reqRedOrdPrice	True when the campaign requireditems with reduced or ordinary prices
25	grossDemand	The average demand created for a campaign
26	NumRecipients	The number of reciepients exposed to a campaign
27	NumOrders	The number of orders resulting from a campaign
28	profit	The average profit of the order, the target value for the regression modelling

A. Data Collection

Data from 826 unique conducted campaigns were collected from a Swedish fashion retailer. Every campaign is described using the 28 features described in Table I. The first seven attributes represent the campaign types, which are mutually exclusive. The following four attributes (8-11) are the discount(s) offered in the campaign. The attributes 12-16 are addons, offering some additional services. The attributes 17-24 are all different requirements on the order to allow the campaign offerings to apply. Attribute 25, finally, is the average demand created for a campaign. The three attributes at the end (26-28) are either a target attribute or used to define a target attribute.

B. Exploratory Data Analysis and Data Preparation

This section discusses the data attributes in more detail while visualizing their distributions.

1. Campaign types: There are 7 campaign types (see attributes 1-7 in Table I). The frequency distribution of the order type is shown in the Fig. 3. Almost 50 % of all campaigns are *first line* campaigns, followed by *all order* and *check* campaigns, with more than 100 campaigns each. There are also 26 unclassified campaigns, lacking campaign types.
2. The discount attributes of the campaigns are attributes 8-10, with attribute 11 representing the average discount for the entire order. Check discounts (*discCheck*) are only combined with *check* campaigns, whereas discounts on the first item (*discFirst*) are included in all types of campaigns except the *check* and *threeForTwo* campaigns. The attribute *discRest* are used in campaigns offering different discount levels, like *firstAndSecondLine*, *combination* and *ladder*. Distribution of the first discount attribute is shown in Fig. 3 below. As expected, since most campaigns include a discount on the first item, the number of campaigns with *discFirst* is much higher, with a majority of campaigns having discounts with 30% or more on the first item. For the same reason, most campaigns do not include check discounts or more than one discount, so the majority of campaigns have 0% on *discCheck* and *discRest.* The campaigns that do include these discounts offer discounts ranging from 10% to 50%.

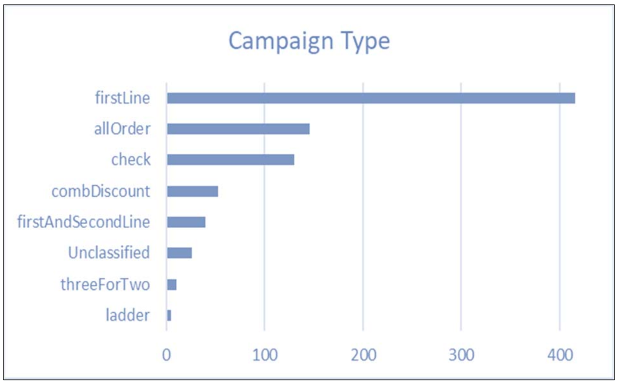

Fig.3. Distribution of order type features

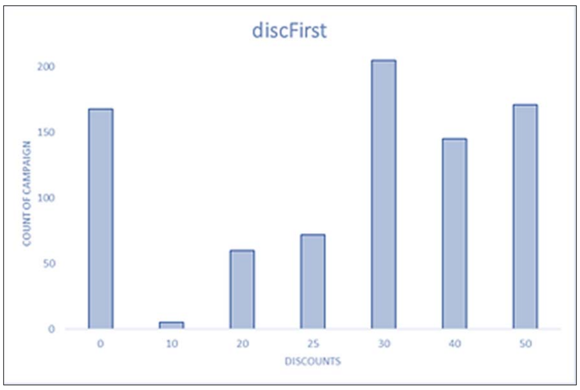

Fig. 4 Distributions of DiscFirst features

3. Addon attributes (12-16) offer some extra service to the customers, like free shipping or a free gift.
4. Requirement attributes (17-24) limits the applicability of the campaign to only apply if you meet the requirements.
5. The gross demand attribute (25) is the average demand created for a campaign.
6. Fig. 5 shows the distribution of average profit per order within campaign. The average profit is rightly skewed, with most orders having a positive profit. The profit numbers are in Swedish currency (SEK) and the majority of the data points lie between 0 and 200.

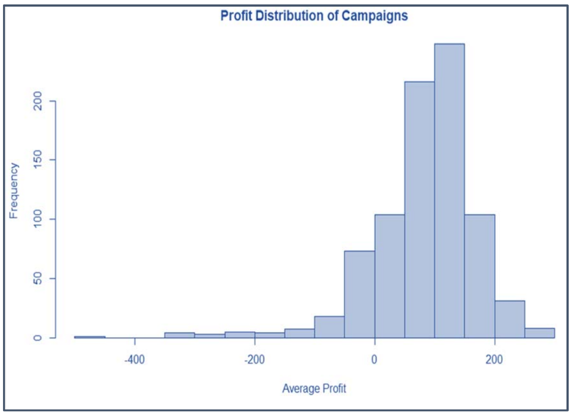

Fig.5. Average profit distribution of campaigns

For the regression experiment (RQ1) the target attribute is the average profit, as described in Table I. For the classification (RQ2) a new feature called *activation* was first created using

$$Activation = \frac{Total\ number\ of\ orders\ received}{Total\ number\ of\ receipients} \quad (1)$$

Looking at the distribution in Fig. 6 below, it can be observed that it is heavily left-skewed. Most of the data points fall between 0 and 0.02, i.e., the response rate is typically less than 2%.

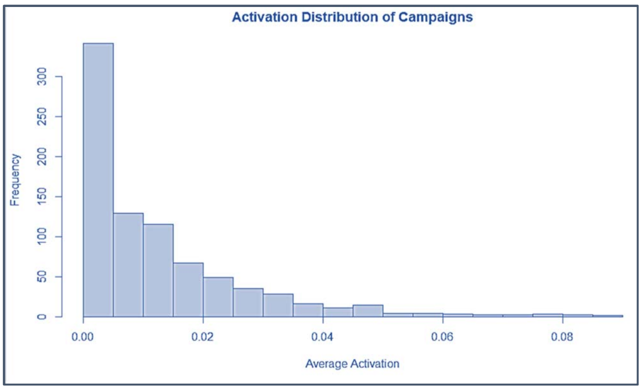

Fig.6. Activation distribution

Now, the classification target for *success* was defined based on both *profit* and *activation* as shown in Fig. 7. Three labels {*Unsuccessful*, *Successful*, *Highly Successful*} were introduced using the criteria in Table II. The obvious logic is that a truly successful campaign should have both high activation and profit.

TABLE II CAMPAIGN PERFORMANCE DEFINING CRITERIA

Model	Values (A=Activation and P=Profit)	Campaign Performance
LA and LP	A <=0.02 & P <0	Unsuccessful
LA and HP	A <=0.02 & P >0	Successful
HA and HP	A > 0.02 & P >0	Highly Successful

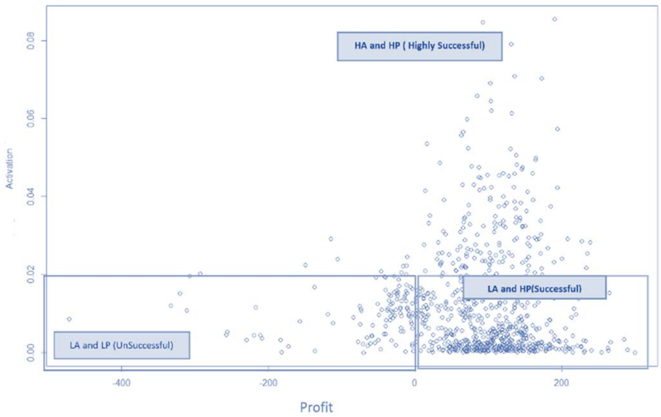

Fig.7. Labels of "Campaign Performance" using Profit and Activation

Out of 826 campaigns, 115 campaigns are classified as *Unsuccessful*, 546 campaigns as *Successful* and 165 campaigns as *Highly Successful*.

C. Modeling

To predict the profit of the campaigns, regression trees and random forest were applied. Here the features 1 to 25 in Table 1 are used as input variables with the target *profit*. Standard 10-fold cross-validation was used for the evaluation. Default model parameters were used for the decision trees and the random forest, i.e., the forests consist of 100 trees.

To predict the success of the campaign, classification trees and random forest were applied. Again, the inputs consist of attributes 1-25 in Table I, with *success* as target. Learning parameters were the same as in the regression experiment.

V. EXPERIMENTAL RESULTS AND DISCUSSION

This section presents and discuss the results.

A. Profitability Model Evaluation

Table III shows the performance of the regression models predicting profitability. First of all, looking at the error metrics, it should be noted that the predictions are quite accurate, on average. Actually, as seen by the MAEs, the mean prediction error is lower than 40 SEK. As expected, the random forest outperformed the regression trees, but, again looking at MAE, the differences in actual numbers are not that large. Also from looking at R2 values, it is obvious that the models are fairly accurate, explaining more than 50% of the relationship between inputs and target.

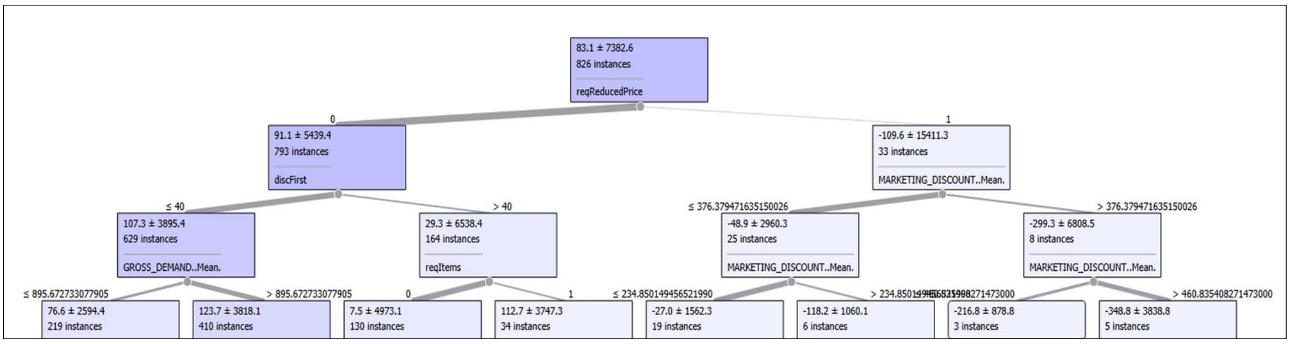

Fig. 8. Regression Tree of Profitability Model

TABLE III Prediction Results on Test data for Profitability Model

Model	RMSE	MAE	R2
Regression Trees	57.458	38.531	0.553
Random Forest	45.783	32.164	0.716

Fig. 8 shows the top four levels of the regression tree built for the profitability task. The depth of the tree is limited for ease of understanding. A few general insights that can be drawn from the tree are that required reduced price results in negative profit, just as discounts higher than 40%, while a required minimum number of items results in increased profit.

B. Campaign Success Model Evaluation

The results for the Campaign Success modelling explained in section IV.C are presented in Table IV. As can be expected, the random forest performs better than the classification tree, even if the difference is not that large.

TABLE IV Classification results on campaign test data

Model	AUC	CA
Classification Trees	0.802	0.741
Random Forest	0.843	0.768

Confusion matrices were created for both models to illustrate classifier predictions and is shown in Fig. 9. It can be seen that Random Forest has predicted all target class labels better than Classification trees with prediction accuracy of 87.2% for *Successful*, 52.7% for *Highly Successful* and 76.5% for *Unsuccessful*. It is worth noting that both models are fairly good at identifying the unsuccessful campaigns and that most misclassifications done for the unsuccessful campaigns were classified as successful rather than highly successful.

To demonstrate the model results, ROC curves are plotted for both the models with the target class instances as shown in Fig. 10. This represents the plot of 'True Positive' (Sensitivity) with 'False Positive' (Specificity), and it could be observed that the ROC curve for the class label *Unsuccessful* has better accuracy for both models than other two class labels, corresponding to the reflection done above. Furthermore, both models have a very steep beginning of the curve, indicating that they both can sort out most of the unsuccessful campaigns rather accurately and quickly.

Classification Tree is plotted to get full insights into the decision rules created by the model (see Fig. 11). The depth of the tree is kept to 4 levels for ease of interpretation. A successful campaign accounts for 66.1 % of the total campaigns. The insights from the tree can be summarized as: High discount on the first item (> 40%) together with a free gift identifies the majority of the unsuccessful campaigns; If it is not a *firstLine* campaign and the discount on the first item is high (> 40%), then the campaign is unsuccessful; Required reduced price will most often result in unsuccessful campaigns; Inclusion of a free gift seems to be a distinguishing mark for highly successful campaigns.

VI. Concluding Remarks

In this paper, we have proposed and evaluated a data-driven solution for fashion retailers to model the success of promotional campaigns. In more detail, two different situations were modeled, (i) to predict the profitability of the promotional campaigns and (ii) to classify the campaign success level. To achieve this, we collected data from one of the leading Swedish fashion retailers, where data attributes describe discount features, order type features, promotional features, number of recipients, number of orders received,

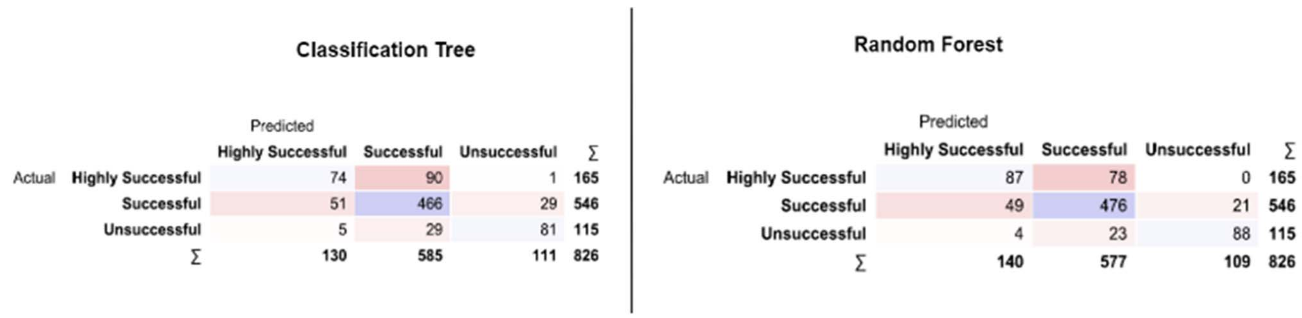

Classification Tree

Actual \ Predicted	Highly Successful	Successful	Unsuccessful	Σ
Highly Successful	74	90	1	165
Successful	51	466	29	546
Unsuccessful	5	29	81	115
Σ	130	585	111	826

Random Forest

Actual \ Predicted	Highly Successful	Successful	Unsuccessful	Σ
Highly Successful	87	78	0	165
Successful	49	476	21	546
Unsuccessful	4	23	88	115
Σ	140	577	109	826

Fig.9. Confusion Matrix with the number of instances

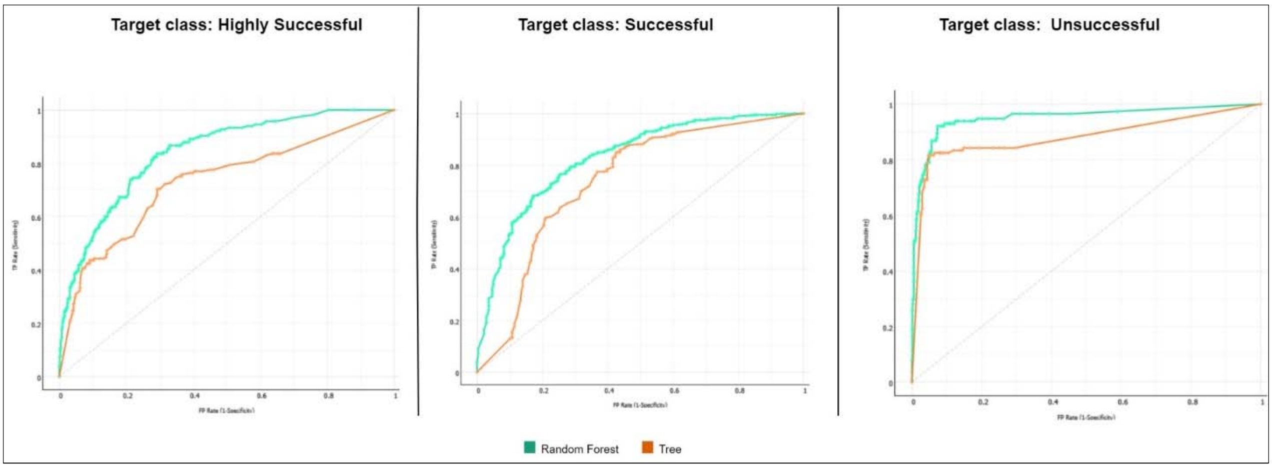

Fig.10. ROC Curve

and profit gained. For the modeling, both tree models and random forest were used. As expected, the random forest models outperformed the tree models, for both classification and regression, regarding the accuracy, but by analyzing the regression and classification trees some valuable insights were learned. Specifically, a high discount on the first item leads to highly profitable campaigns. Besides, if this was combined with a free gift and marketing discounts, it would lead to highly successful campaigns.

Thus, this case study demonstrated that data-driven methods can be used to understand, and ultimately optimize campaigns and promotions. Obviously, such models could be used to simulate campaigns prior to going live, potentially giving the retailers a sophisticated tool for campaign planning.

VII. Acknowledgment

Chandadevi Giri's work is carried out under the framework of Sustainable Management and Design for textiles (SMDTex). Ulf Johansson and Tuwe Löfström are part of the Data-Driven Innovation: Algorithms, Platforms and Ecosystems project, funded by the Knowledge foundation (Grant number: 20160035). The authors are grateful to the Swedish fashion retailer for supplying the data.

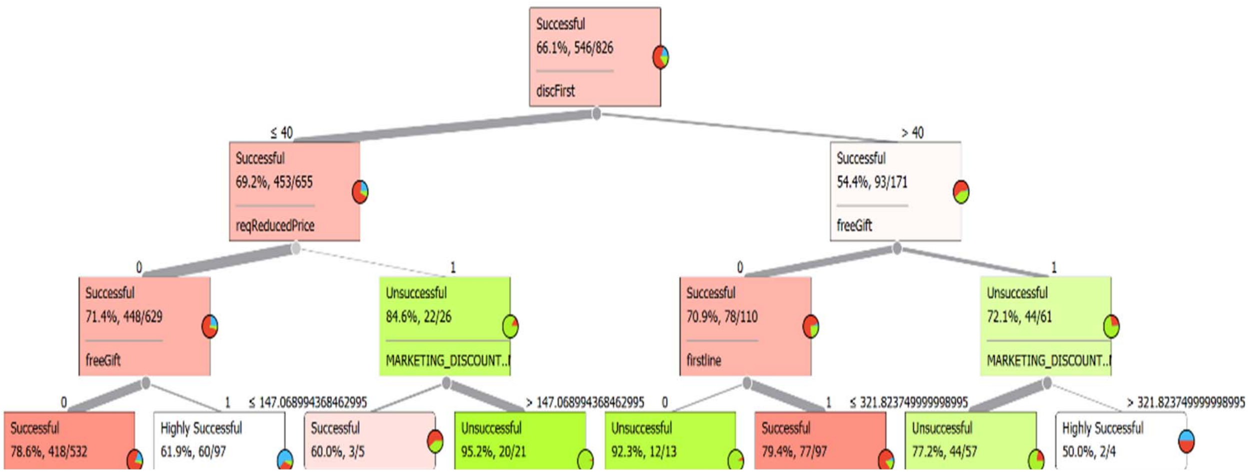

Fig. 11 Classification Tree for Campaign Success Model

REFERENCES

[1] D. Grewal, K. L. Ailawadi, D. Gauri, K. Hall, P. Kopalle, and J. R. Robertson, “Innovations in Retail Pricing and Promotions,” *J. Retail.*, vol. 87, pp. S43–S52, Jul. 2011.

[2] K. L. Ailawadi, J. P. Beauchamp, N. Donthu, D. K. Gauri, and V. Shankar, “Communication and Promotion Decisions in Retailing: A Review and Directions for Future Research,” *J. Retail.*, vol. 85, no. 1, pp. 42–55, 2009.

[3] F. M. Feinberg, A. Krishna, and Z. J. Zhang, “Do we care what others Get? A Behaviorist Approach to Targeted Promotions,” *J. Mark. Res.*, vol. 39, no. 3, pp. 277–291, Aug. 2002.

[4] C. Giri, S. Thomassey, and X. Zeng, “Customer Analytics in Fashion Retail Industry,” in *Functional Textiles and Clothing*, Singapore: Springer Singapore, 2019, pp. 349–361.

[5] D. Grewal, D. M. Hardesty, and G. R. Iyer, “The effects of buyer identification and purchase timing on consumers’ perceptions of trust, price fairness, and repurchase intentions,” *J. Interact. Mark.*, vol. 18, no. 4, pp. 87–100, Jan. 2004.

[6] T. T. Nagle and J. E. (College teacher) Hogan, *The strategy and tactics of pricing : a guide to growing more profitably*, 4th ed. Upper Saddle River, N.J. : Pearson/Prentice Hall, 2006.

[7] M. J. Barone and T. Roy, “Does Exclusivity Always Pay Off? Exclusive Price Promotions and Consumer Response,” *J. Mark.*, vol. 74, no. 2, pp. 121–132, Mar. 2010.

[8] G. Ryu and L. Feick, “A Penny for Your Thoughts: Referral Reward Programs and Referral Likelihood,” *Journal of Marketing*, vol. 71. Sage Publications, Inc., pp. 84–94.

[9] L. Lee and D. Ariely, “Shopping Goals, Goal Concreteness, and Conditional Promotions,” *J. Consum. Res.*, vol. 33, no. 1, pp. 60–70, Jun. 2006.

[10] S. Ganesan, M. George, S. Jap, R. W. Palmatier, and B. Weitz, “Supply Chain Management and Retailer Performance: Emerging Trends, Issues, and Implications for Research and Practice,” *J. Retail.*, vol. 85, no. 1, pp. 84–94, Mar. 2009.

[11] S. K. Hui, P. S. Fader, and E. T. Bradlow, “Path Data in Marketing: An Integrative Framework and Prospectus for Model Building,” *Mark. Sci.*, vol. 28, no. 2, pp. 320–335, Mar. 2009.

[12] J. Lindsey-Mullikin and D. Grewal, “Imperfect Information: The Persistence of Price Dispersion on the Web,” *J. Acad. Mark. Sci.*, vol. 34, no. 2, pp. 236–243, Apr. 2006.

[13] S.-F. S. Chen, K. B. Monroe, and Y.-C. Lou, “The effects of framing price promotion messages on consumers’ perceptions and purchase intentions,” *J. Retail.*, vol. 74, no. 3, pp. 353–372, Sep. 1998.

[14] R. Nayak and R. Padhye, *Artificial intelligence and its application in the apparel industry*. Elsevier Ltd, 2017.

[15] C. M. Bishop and C. M., *Pattern recognition and machine learning*. Springer, 2006.

[16] L. Breiman, J. Friedman, R. Olshen, C. S.- Group, and undefined 1984, “Classification and regression trees. Wadsworth Int.”

[17] J. R. (John R. Quinlan, *C4.5 : programs for machine learning*. Morgan Kaufmann Publishers, 1993.

[18] L. Breiman, “Random Forests,” *Mach. Learn.*, vol. 45, no. 1, pp. 5–32, 2001.