ISBN: 978-972-8939-09-0 © 2010 IADIS

INTERACTIVE AUDIENCE SELECTION TOOL FOR DISTRIBUTING A MOBILE CAMPAIGN

Talya Porat, Lihi Naamani-Dery, Lior Rokach and Bracha Shapira
Deutsche Telekom Laboratories at Ben Gurion University
Beer Sheva, 84105, Israel

ABSTRACT

An intelligent model for campaign management was developed as a collaborative research effort between Deutsche Telekom and Ben-Gurion University. The model segments and filters potential customers for solicitation according to a novel algorithm. However, a mathematical model, however cleverly designed, doesn't encompass the human knowledge, experience and specific requests of an expert campaign manager. Therefore campaign managers might be reluctant to use the model as it stands. The 'Interactive Audience Selection' solution proposed here attempts to bridge the gap between the campaign managers and the intelligent model by steering to combine and benefit from both the expert knowledge of the managers as well as the model's mathematical capabilities. The suggested information system tool enables the campaign manager to view the models' choice of potential customers for solicitation, understand it and improve it for the next audience sample. We propose that the transparency of the models' logic and the control the user has over the model's output will lead to the successful execution of a profitable campaign, without neglecting customer value.

KEYWORDS

Database Marketing, Marketing, Campaign Management, Interaction design, CRM, Mobile.

1. INTRODUCTION

Today, mobile customers are unaware of available mobile functionality and services. Huge investments in new service creation sometimes result in low returns. Many promotions for Telecom companies' services are mass communicated, instead of singling out specific audiences. The tool we propose is designed for the use of marketing personnel (e.g. campaign managers) for introducing relevant services and products to mobile customers. The proposed model applies a target marketing strategy: offering the right product to the right customer at the right time using the proper distribution channel. The model estimates the response probability of the potential customers, and helps the decision-maker assess the profitability of the different customers. Using the model, the marketing expert approaches the customers predicted to be most interested in the services and proposes a marketing offer. This strategy affords better efficiency than a mass marketing strategy, in which a product or service is mass-communicated to all mobile customers, usually resulting in low positive response rates.

From the campaign managers' perspective, there are two main advantages in the model's mathematical capabilities. The first is maximizing campaign profit; allowing the model to solicit customers with the highest profit potential as campaign audience, results in the most profitable campaign. The second advantage is revealing new audiences. Since the model is not limited to a specific audience, it samples customers in order to find the most profitable audiences. This may reveal new audiences that were not assumed to be potential customers.

Although the benefits of an intelligent model have been previously demonstrated (see for example Rokach et al., 2008), the value that a campaign manager provides cannot be automated out of existence. A Database marketing model lacks the human experience and intuition needed to recognize other considerations besides profitable customers. The critical role the human expert has in controlling the exposure of audience to a specific campaign may be illustrated in the following examples. The first is the campaign managers' necessity to avoid backlashes. More important than correctly assuming who will answer positively to a specific campaign, is to know whom to avoid marketing, even on account of loss of profit. An extreme case may be advertising cigars to children or teenagers. Obviously, marketing this audience may cause undesired outcomes or even a backlash. The second example is the need to preserve customer value by avoiding intrusiveness. The campaign manager may decide not to approach a certain customer profile, even if this profile has a high probability to accept the service. One case could be when a customer received many offers in the recent past.

In addition, the user may feel insecure using the model without understanding its logic and without the ability to supervise and control its output. In several application areas, security requirements demand that models can only be trusted if they can be understood by the user (Nauck et al., 2003). This result is in line with Payton and Zahay (2005) findings that lack of trust, low perceived data quality and unfulfilled marketing needs are the three main reasons for explaining why a corporate data warehouse (CDW) was not used by marketing managers to the extent that it was expected to be used.

Due to the above, the goal of our proposed 'Interactive Audience Selection' solution is to provide the campaign manager with an intelligent interactive system, which will combine the managers' expert knowledge and the model’s mathematical capabilities, thus benefiting from both.

Reviewing the literature reveals that there is almost no reference to the unique data needs of marketing personnel and to the user interface of database marketing solutions (Payton and Zahay, 2005). Only a minority of users have access to data-mining software, which is usually complicated to use and method-oriented instead of user-oriented (Nauck et al., 2003). Without detailed knowledge of the model's analysis methods, such tools are exposed to distribution-errors which may disrupt the efficiency of the audience targeting process.

Today, organizations have limited ability to monitor campaign performance during execution, and cannot make adjustments 'on the fly' for maximum effectiveness. The suggested tool goes one step further than today's database marketing solutions, by integrating the campaign manager as an active player in the model's procedure of finding the potential profitable customers for the campaign. The increase in the user's control over the campaign distribution consequently increases the users' trust in the model and thus increases the use of the model. We believe this will lead to a more profitable campaign, without neglecting customer value.

In the next sections, we describe the model's intelligence and the interactive audience selection solution.

2. THE INTELLIGENT MODEL

Pessimistic Active learning (PAL) is a novel algorithm for the discrete choice target marketing problem. It attempts to address the problem of which potential customers should be approached with a new product offer in order to maximize the net profit. The model focuses on binary discrete choice problems, where the customer’s response is binary, such as acceptance or rejection of a marketing offer. The model uses active learning (Cohn et al., 1996), a data mining policy which actively selects unlabeled instances for labeling. In a marketing context, unlabeled instances are customers with an unknown response and labeling means acquiring the customers’ response. Active learning has been previously used for facilitating direct marketing campaigns (Saar-Tsechansky and Provost, 2007). We assume that a marketing campaign doesn’t approach its customers one by one, but rather in sessions. In each session a batch of customers is simultaneously approached. Thus the marketing department can concentrate its exploitation of resources, such as marketing personnel and equipment. The number of customers that is courted in each session, if not defined otherwise by the campaign manager, will be computed according to the total number of potential customers. A detailed description of PAL’s innovations and their assessments can be found in Rokach et al., 2008. As a more recent example, Figure 1 shows the profit gained from running PAL on the newly published real-world dataset of Orange [1] customers.

[1] http://www.kddcup-orange.com/

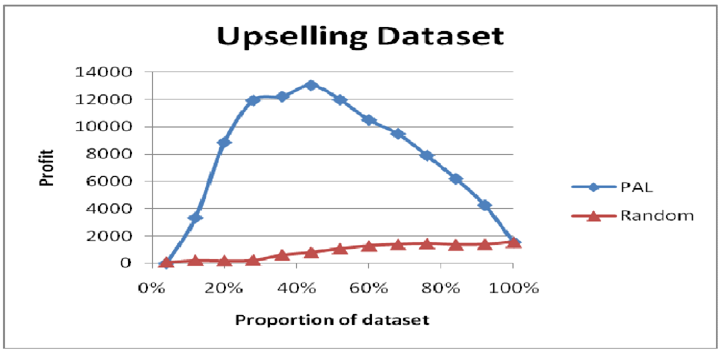

Figure 1. PAL vs. a Random approach.

Axis x shows the proportion of the dataset, that is the amount of customers approached. Axis y displays the profit gained, assuming a fixed income of 14 for a positive response and loss of 1 for a negative response. The baseline PAL is compared against a random method, which randomly samples customers to be courted. This method displays a linear profit. If the objective is to approach all customers (100%) of the dataset, then no model is required. However, PAL's advantage can be clearly seen when only a percentage of the customers are selected. PAL displays a profit peak when around 50% of the dataset is exploited.

2.1 Design Challenge

When designing the implementation of the model into a real-world working system the question was how to design an intelligent, easy and intuitive system which will combine both ends: the user and the model. Thus, the success of the project hinges on taking advantage of the model's intelligence while retaining user control over the model's output, together with a careful and creative design of the interactive process.

As a first step, we defined the user profile and generated several possible user scenarios. The two main scenarios are described. The first scenario is searching for the most profitable solution. This can be the case when marketing a more common service or product that hasn't many boundaries in its distribution (such as an Internet service), and therefore has more tolerance for distribution errors. In this case, the user may activate the model with or without pre-defined audience filtering and by this complete the audience selection interaction. This solution also supports the growing need of data analysis tools to run largely unsupervised, since businesses do not have the time to deploy a team of experts to analyze data about service usage (Nauck et al., 2003). Another possible scenario is marketing services or products that have a more limited target audience (e.g. children's games, chat services, etc.), and therefore less tolerance for distribution errors. In this case, in order to increase customer value and prevent backlashes and intrusiveness, the profit model cannot be accepted as the sole decision authority. The campaign manager needs the option to review the model's output in order to understand it and improve it if necessary.

In the following section, we describe some of the suggestions we propose to satisfy the different scenarios.

2.2 Design Solution

The Interactive Audience Selection solution contains three iterative steps (Figure 2). First, the user may define preliminary audience filtering. Second, the model runs on the pre-defined audience filters and finally, the combined output is reflected to the user for review, and if necessary for fine-tuning. This process may be done for each distribution session until the campaign distribution is complete. The mean number of sessions or iterations for a campaign is around ten.

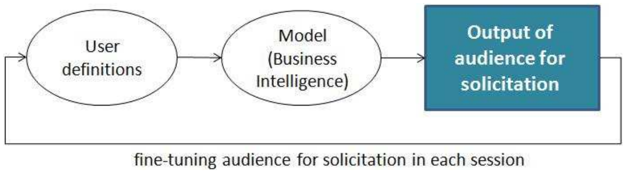

Figure 2. Interactive Audience Selection Solution.

A PC prototype was developed to demonstrate the viability of the Interactive Audience Selection concept. The prototype was designed according to the three iterative steps: (1) preliminary audience filtering (2) campaign distribution (model) (3) review and fine-tune output (see Figure 3).

From a practical point of view, certain restrictions have to be imposed on the model, such as running the model only in a specific region (e.g. Berlin), or only on customers that have a specific device (e.g. iPhone) or on the contrary, exclude specific customers (e.g. children), therefore the first step enables the user to define preliminary audience filters to be considered by the model.

In the second step, the model runs on the pre-defined audience filters. In the first iteration, the model randomly samples customers to be courted under the pre-defined audience restrictions. In the following iterations the campaign distribution is more educated, since it learns how to classify future sets of data from customers' responses. In this step a dynamic diagram displaying the customers that are being approached 'online' is displayed. General information such as the number of customers that are being approached and the number of responses (positive/negative) may also be displayed. When the current distribution is done and all relevant customers were approached, the third step is displayed. The third step is for reviewing and improving the model's output. It enables the campaign managers to analyze 'on-the-fly' the first group of customers which already received the campaign ('Current iteration') and fine-tune the next group which will be approached in the next session ('Future iteration').

Thus, the user may track the feature selection process and figure out which predictor, or a set of predictors, was eliminated at which step and why. The importance of each predictor to explaining the response is also summarized, in addition to a summary of explanations about patterns detected in the data. The 'Future iteration' display, which displays the customers that will be approached in the next iteration, enables users to add a new audience filter, remove an existing one or change the ratio between the different filters. Analyzing and editing the data is done using direct manipulation (drill-down and drag & drop). The user may select the desired display (diagram, pie-chart, table, etc.) both in the current iteration and future iteration views.

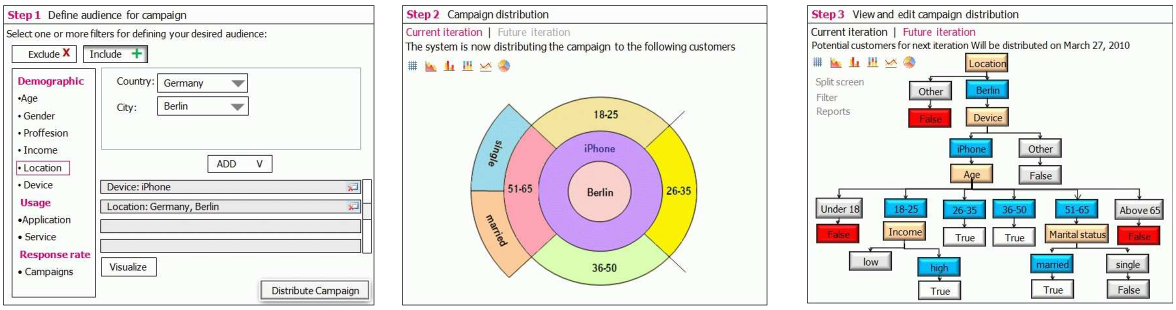

Figure 3. Interactive Audience Selection Design.

3. CONCLUSION

We have presented our interactive audience selection solution for optimizing the identification of prospective customers for a mobile marketing campaign. Our ongoing research project concentrates on evaluating the interactive system. We plan to conduct an experiment which will compare the three different audience selection methods: (1) manual user filtering – the user defines and filters the desired campaign audience (2) model filtering - automatic model filtering for finding the desired campaign audience (3) Interactive audience selection – our combined interactive solution for finding the desired campaign audience. We plan to measure the three methods by subjective measures (e.g. user satisfaction) and objective measures (e.g. campaign profit, backlashes, intrusiveness). The uniqueness of our interactive audience selection solution lies in giving the user the ability to monitor campaign distribution during execution and make adjustments "on the fly".

REFERENCES

Cohn, D. A., Ghahramani, Z., Jordan, M. I. 1996. Active learning with statistical models. *Journal of Artificial Intelligence Research*, Vol. 4, pp. 129–145.

Nauck, D., Spott, M., and Azvine, B. 2003. SPIDA – A Novel Data Analysis Tool. *BT Technology Journal*, Vol. 21, No. 4, pp. 104-111.

Payton, F.C., and Zahay, D. 2005. Why doesn't marketing use the corporate data warehouse? The role of trust and quality in adoption of data-warehousing technology for CRM applications. *Journal of Business & Industrial Marketing*, Vol. 20, No. 4/5, pp. 237-244.

Rokach L., Naamani L., Shmilovici A. 2008. Pessimistic Cost-sensitive Active Learning of Decision Trees for Profit Maximizing Targeting Campaigns. *Data Mining and Knowledge Discovery*, Vol. 17, No. 2, pp. 283-316. DOI= http://dx.doi.org/10.1007/s10618-008-0105-2

Saar-Tsechansky M. and Provost F. 2007. Decision-Centric Active Learning of Binary-Outcome Models. *Information systems research*, Vol. 18, No. 1, pp. 4-22.