Dynamic Micro Targeting: Fitness-Based Approach to Predicting Individual Preferences

Tianyi Jiang, Alexander Tuzhilin

New York University
tjiang, atuzhili@stern.nyu.edu

Customer segmentation, such as customer grouping by the level of family income, education, or any other demographic variable, is considered as one of the standard techniques used by marketers for a long time [10]. Its popularity comes from the fact that segmented models usually outperform aggregated models of customer behavior [11]. More recently, there has been much interest in the marketing and data mining communities in learning *individual* models of customer behavior within the context of ***1-to-1*** marketing [9] and personalization [4], when models of customer behavior are learned from the data pertaining only to a particular customer. These learned individualized models of customer behavior are stored as parts of customer profiles and are subsequently used for recommending and delivering personalized products and services to the customers [2].

As was shown in [7], it is a non-trivial problem to compare segmented and individual customer models because of the tradeoff between the sparsity of data (bias) for individual customer models and customer heterogeneity (variance) in aggregate models: individual models may suffer from sparse data, while aggregate models suffer from high levels of customer heterogeneity.

A typical approach to customer segmentation is based on the *statistics-based* approach that computes the set of statistics from customer's demographic and transactional data [3, 7, 12], such as the average time it takes the customer to browse the Web page describing a product, maximal and minimal times taken to buy an online product, RFM statistics [8], etc. After such statistics are computed for each customer, the customer base is partitioned into customer segments by using various clustering methods on the space of the computed statistics [7]. It was shown in [7] that while the best statistics-based approaches can be effective and even outperform the ***1-to-1*** case under certain conditions, the approach can also be very ineffective as different customer statistics calculations result in different *n*-dimensional spaces and various distance metrics or clustering algorithms would yield very different clusters.

Recent research [6] proposes the *direct grouping* segmentation approach that partitions the customers not based on computed statistics and particular clustering algorithms, but in terms of directly combining *transactional data* of several customers, such as Web browsing and purchasing activities, and building a single model of customer behavior on this combined data. This approach avoids the pitfalls of the statistics based-approach in that it does not require selection of arbitrary statistics and grouping customers based on these statistics. Instead, it provides a more direct approach to customer segmentation by combining customers' data collectively resulting in better model for this group of customers. We have shown in [6] that the *direct grouping* segmentation approach dominates the statistics-based segmentation and the ***1-to-1*** approaches [5, 6].

In this paper we aim to increase the performance of previous customer *segmentation* approaches [5-7] via the method of *micro targeting*, where predictive models of customer behavior are built not on the segments of customers but rather on the customer-product groups. For example, we may want to build a model predicting whether a customer is going to purchase a particular product during a visit to a certain website, and we may want to build this model for a certain segment of customers *and* a certain category of products (e.g., consumer electronics). Therefore, this approach no longer groups customers into a fixed set of segments, but rather identifies *micro-targeting regions* in the Customer × Product space that are the most suitable for building the best *local* predictive models of consumer behavior for a particular customer and product. Moreover, we can build multiple local models for a particular customer and product and select the best-performing one when trying to predict customer behavior for a particular product. The proposed approach was based on the observation that a customer may possess different underlying utility functions across different types of products, and therefore should be modeled separately for different product types. The advantage of this approach is that the *micro-targeting region* is generated dynamically from a fitness maximization perspective, and that we could improve the predictive accuracies by generating models from individual and pooled customers' product category specific preferences. Since the identification of best-performing micro-targeting regions is a strictly more general problem than the *direct grouping* segmentation, *micro targeting* should perform at least as good as for the case of segmentation.

The proposed approach is related to the reduction-based method for providing multi-dimensional recommendations [1] where certain segments of ratings are selected from the multi-dimensional cube of ratings and recommendation algorithms use these and only these segments of ratings. However, unlike [1], we focus on building customers' local models in this paper, rather than on the recommendation problem.
Similarly to *direct grouping* methods, *micro targeting* is a type of Optimal Customer Segmentation (OCS) problem [6], which was shown to be NP-hard in [6]. Therefore, as in [6], we propose a suboptimal polynomial time heuristic for constructing *micro-targeting regions* using the *direct grouping* approach advocated in [5, 6]. The general heuristic of *micro targeting* works as follows. Starting from single-customer's product specific transaction set, we iteratively seek to merge existing customer transaction sets by combining data from two sets SetA and SetB when 1) the predictive model based on the combined data for transaction sets SetA and SetB performs better than respective models on SetA and SetB and 2) combining SetA with any other existing transaction set would have resulted in worse performance than the combination of both SetA and SetB. *Micro targeting* deploys a greedy search strategy since it determines the best pair of customer product specific transaction sets at each iteration and merges them together resulting in the best local solution. The algorithm terminates when there are no more improvements to be made from pairing existing customer product specific sets.

We expect the proposed suboptimal *micro targeting* approach to outperform previous *direct grouping* methods because of its granular Customer x Product search space. *Micro targeting* methods should also outperform the statistics-based segmentation and the ***1-to-1*** approaches since [6] has shown the dominance of *direct grouping* methods.

The current status of this research is in early prototype phase. We have designed several algorithms that implement the concepts of our proposed *micro targeting* approach, and will be conducting tests to verify the validity of our proposed methods across a wide range of empirical settings including different (a) panelist datasets (real and synthetic), (b) micro-targeting approaches, (c) prediction tasks, (d) performance measures, and (e) predictive classifiers built on the *micro targeting* sets of customers and products.

In summary, the contribution of this research lies in achieving better predictive performance of local predictive models by increasing the granularity of the customer preference analysis without sacrificing predictive accuracy and generalizability of previous customer segmentation approaches. Our proposed algorithm also runs within the same polynomial bounds as that of the best performing *direct segmentation* approaches so as to keep computation expense within practical limits.

References

[1] Adomavicius, G., R. Sankaranarayanan, S. Sen, and A. Tuzhilin, *Incorporating contextual information in recommender systems using a multidimensional approach.* ACM TOIS, 2005. **23**(1): p. 103-145.
[2] Adomavicius, G. and A. Tuzhilin, *Personalization technologies: A process-oriented perspective,* in *CACM*. 2005.
[3] Brijs, T., T. Swinnen, K. Vanhoof, and G. Wets. *Using shopping baskets to cluster supermarket shoppers*. in *AARTF*. 2001. Amelia Island Plantation, FL.
[4] CACM, *Communications of ACM*, in *Special Issue on Personalization*. 2000.
[5] Jiang, T. and A. Tuzhilin. *Grow Your Customers through Time: Efficient Online Segmentation of a Preparation: Growing Customer Base*. in *WITS*. 2006.
[6] Jiang, T. and A. Tuzhilin. *Improving Personalization Solutions through Optimal Segmentation of Customer Bases*. in *ICDM*. 2006.
[7] Jiang, T. and A. Tuzhilin, *Segmenting Customers from Population to Individual: Does 1-to-1 Keep Your Customers Forever?* IEEE TKDE, 2006. **18**(10).
[8] Novo, J., *Drilling Down: Turning Customer Data Into Profits with a Spreadsheet*. 2004: Booklocker.
[9] Peppers, D. and M. Rogers, *Enterprise One to One*. 1997, New York: Bantam Pub. Group Inc.
[10] Smith, W., *Product Differentiation and Market Segmentation as Alternative Marketing Strategies.* Journal of Marketing, 1956. **21**: p. 3-8.
[11] Wedel, M. and W. Kamakura, *Market Segmentation: Conceptual and Methodological Foundations*. 2nd ed. 2000: Kluwer Publishers.
[12] Yang, Y. and B. Padmanabhan. *Segmenting Customer Trans. Using a Pattern-Based Clustering Approach*. in *ICDM*. 2003.