Measuring and Visualising Similarity of Customer Satisfaction Profiles for Different Customer Segments

Frank Klawonn[1], Detlef D. Nauck[2] and
Katharina Tschumitschew[1]

[1] Department of Computer Science University of Applied Sciences BS/WF
Salzdahlumer Str. 46/48, D-38302 Wolfenbuettel, Germany
[2] BT Group, Chief Technology Office, Research and Venturing Intelligent Systems
Research Centre Adastral Park, Orion Building pp1/12, Ipswich IP5 3RE, UK

Abstract. Questionnaires are a common tool to gain insight to customer satisfaction. The data available from such questionnaires is an important source of information for a company to judge and improve its performance in order to achieve maximum customer satisfaction. Here, we are interested in finding out, how much individual customer segments are similar or differ w.r.t. to their satisfaction profiles. We propose a hybrid approach using measures for the similarity of satisfaction profiles based on principles from statistics in combination with visualization techniques. The applicability and benefit of our approach is demonstrated on the basis of real-world customer data.

Keywords: customer satisfaction; rank correlation; MDS; cluster analysis.

1 Introduction

Customer satisfaction is a key issue for a company to maintain and improve its position in the market. Questionnaires filled in by customers via telephone interviews, direct interviews, mail or the Internet provide a very important source of information on customer satisfaction. Such questionnaires usually contain questions concerning different aspects of customer satisfaction as well as other questions regarding other general or specific information like age of the customer or which item or service they have purchased from the company. There are, of course, many ways to analyse the data available from such questionnaires, depending on the kind of question or information the company is interested in [3, 4, 8]. This paper focuses on the following aspect. Customers are usually grouped into different customer segments. Customer satisfaction might be similar or differ among these customer segments. When significant differences among customer segments can be identified, the company can use this information to take appropriate actions in order to improve the customer satisfaction, especially for those customer segments where the satisfaction is lower. Differences found in customer satisfaction can also help to estimate the impact of possible company campaigns or actions on the customer segments.

2 Problem Description

Here, the only questions of interest within a questionnaire are those asking customers directly about their satisfaction. Nevertheless, besides the overall satisfaction a questionnaire will usually contain questions regarding customer satisfaction with respect to different criteria, for instance concerning different services, quality of products, prices or information provided by the company. The answers to a question regarding customer satisfaction are usually limited to a specific ordinal scale with varying granularity for different criteria. In the simplest case the ordinal scale might only contain two answers, i.e. "Are you satisfied with ...? Yes/No". But in most cases a set of more refined answers is provided, for instance, extremely satisfied, very satisfied, satisfied, ... , extremely dissatisfied.

In addition to the ordinal scale that allows specifying the degree of customer satisfaction explicitly, there are usually additional answers like "don't know", "not applicable" or "refuse to answer". It can also happen that a customer does not answer a question. Here, all these cases that do not provide an explicit evaluation of the customer satisfaction, are considered in the same way as a null answer. In the following, we will refer to a null answer as a missing value.

Finding similarities and especially differences concerning customer satisfaction for different customer segments is the focus of this paper. We assume that the customer segments are given. The segmentation might depend on the customer's age, income, area of residence and other aspects. How the customer segmentation is defined exactly is not relevant for this paper.

We assume that altogether a number of q different customer satisfaction questions are considered. Each question has an individual ordinal scale of possible answers plus a specific category for "missing value" as described above. We also consider a number of c different customer segments. We assume that statistics for each customer segment and each customer satisfaction question are available. This means for each customer segment and each question, we know either the absolute or the relative frequencies of the possible answers, including missing values, to the question. It is not required that exactly the same questionnaire was presented to customers from different customer segments. It is only necessary that at least the same q questions concerning customer satisfaction were contained in each questionnaire. Furthermore, all questionnaires must use the same granularity for the ordinal scale of corresponding customer satisfaction questions while the ordinal scales for different questions may vary.

3 Measuring Similarities between Customer Satisfaction Profiles

In order to compare the satisfaction profiles of different customer segments with respect to any of the questions, we first have to compare the corresponding distributions over the possible answers. Initially, we restrict the comparison of two customer segments to a single question. The combination of a number of customer satisfaction questions will be considered later in this section.

A very naive approach for comparing two segment on one question would be to simply compare the distributions on all answers including the missing values. However, this can be misleading if the proportions of missing values are not identical in the two customer segments. To illustrate this effect, consider the following artificial example. Assume that in both customer segments all customers who have provided an answer on the ordinal scale have voted for the same degree of satisfaction. However, in the first customer group there are 20% missing values, whereas there are no missing values in the second group. The difference in customer satisfaction for these two groups lies only in the proportion of missing values, but not in the distribution of those who have provided an evaluation of their customer satisfaction. For this reason, we consider the distributions on the ordinal scale and the proportions of missing values separately.

Ignoring the missing values means we first have to normalise the two distributions over the values of the ordinal scale, so that the frequencies add up to 100%. In the above simple example, this would mean that the two distributions over the non-missing answers would be identical after normalisation.

The similarity or difference of two probability distributions over an ordinal scale could be measured on the basis of the differences of the frequencies or in terms of the Kullbach-Leibler entropy (see for instance [2]). However, in this way the ordinal scale would be considered as a finite set of discrete values without any ordering structure. As an extreme example consider three distributions. For the first distribution 100% of the probability mass is concentrated on the largest value of the ordinal scale, for the second one 100% of the probability mass is concentrated on the second largest value, for the third 100% of the probability mass is concentrated on the smallest value. Comparing these distributions in terms of frequency differences or in terms of the Kullbach-Leibler entropy would tell us that they differ in the same way. However, it is obvious that the the first distribution is more similar to the second one than to the last one, for example.

Therefore, we propose to compare the cumulative distribution functions over the ordinal scale in a manner not identical, but similar to the Wilcoxon rank test, also called Mann-Whitney-U-test, known from statistics (see for instance [5, 7, 11]). When the ordinal scale for the question X has the values (possible answers) $v_1, \ldots, v_h$ and the probability distribution is given as $P(X = v_k) = p_k$, $(k = 1, \ldots, h)$, then the cumulative distribution function is $P(X \leq v_k) = F_k = \sum_{i=1}^{k} p_i$. The pointwise difference

$$d_0\left(P^{(1)}, P^{(2)}\right) \;=\; \sum_{k=1}^{h} \left|F_k^{(1)} - F_k^{(2)}\right| \;=\; \sum_{k=1}^{h-1} \left|F_k^{(1)} - F_k^{(2)}\right| \qquad (1)$$

between the cumulative distribution functions seems to be more appropriate to measure the difference between two probability distributions on an ordinal scale. Note that $F_h^{(1)} = F_h^{(2)} = 1$ always holds.

The distance measure d_0 will have a tendency to higher values, when the ordinal scale has more values, i.e. h is large. This means that questions with finer granularity tend to contribute much more to the difference between satisfaction profiles. We take this effect into account as follows. Consider the two cases:

1. Assume an ordinal scale with just two values (i.e. $h = 2$, for instance, when a question with the only answers yes and no is considered). For the two extreme distributions where 100% of the probability is put on one answer and 0% on the other, the distance measure (1) will yield the value $d_0 = 1$.
2. Now assume an ordinal scale with $h = 2r$ values. Consider the distribution $P^{(1)}$ where the answers are uniformly distributed over the first r values, i.e. $p_1^{(1)} = \ldots = p_r^{(1)} = 1/r$ and $p_{r+1}^{(1)} = \ldots = p_{2r}^{(1)} = 0$, and the distribution $P^{(1)}$ where the answers are uniformly distributed over the last r values, i.e. $p_1^{(2)} = \ldots = p_r^{(2)} = 0$ and $p_{r+1}^{(2)} = \ldots = p_{2r}^{(2)} = 1/r$. Then we obtain $d_0\left(P^{(1)}, P^{(2)}\right) = r = \frac{h}{2}$.

We require that the distance (dissimilarity) between the distributions in the first case should be the same as the distance between the two distribution in the second case. Therefore, we introduce a correction factor and use $d_{\mathrm{ord}}\left(P^{(1)}, P^{(2)}\right) = \frac{2}{h} \cdot d_0\left(P^{(1)}, P^{(2)}\right)$ as the distance between two probability distributions on an ordinal scale with h values (possible answers).

So far we have only compared probability distributions over an ordinal scale ignoring missing values. In order to take the missing values into account, we compute the difference $d_{\mathrm{miss}}\left(p_{\mathrm{miss}}^{(1)}, p_{\mathrm{miss}}^{(2)}\right) = \left|p_{\mathrm{miss}}^{(1)} - p_{\mathrm{miss}}^{(2)}\right|$ between the relative frequencies of the missing values. The overall distance between the (normalised) probability distributions $P^{(1)}$ and $P^{(2)}$ with a proportion (relative frequency) of missing values $p_{\mathrm{miss}}^{(1)}$ and $p_{\mathrm{miss}}^{(2)}$, respectively, is a convex combination of the distances d_{ord} and d_{miss}: $d_{\mathrm{ord+miss}}\left(\left[P^{(1)}, p_{\mathrm{miss}}^{(1)}\right], \left[P^{(2)}, p_{\mathrm{miss}}^{(2)}\right]\right) =$

$$\left(1 - \max\{p_{\mathrm{miss}}^{(1)}, p_{\mathrm{miss}}^{(2)}\}\right) \cdot d_{\mathrm{ord}}\left(P^{(1)}, P^{(2)}\right) + \max\{p_{\mathrm{miss}}^{(1)}, p_{\mathrm{miss}}^{(2)}\} \cdot d_{\mathrm{miss}}\left(p_{\mathrm{miss}}^{(1)}, p_{\mathrm{miss}}^{(2)}\right).$$

This way, the influence of the difference between the probability distributions on the ordinal scale is reduced when at least one of them has a high proportion of missing values.

So far, we have only discussed measuring the difference between two distributions over an ordinal scale incorporating missing values. In the case of customer satisfaction profiles, we may use this approach for comparing the distributions of answers in two customer segments with respect to one question. For comparing two customer segments with respect to a number of questions we simply add up the distances obtained for the single questions.

4 Visualisation

The previous section provides a method to compute the dissimilarity between two customer segments with respect to their answers to selected questions. Although a pairwise comparison of customer satisfaction profiles will already provide important insights in the relation between customer segments and customer satisfaction, it is even more interesting to have an overall overview about how similar or different the satisfaction profiles of a collection of customer segments are. In order to provide this overview, we visualise the distances (dissimilarities) in the plane. We compute the pairwise distances between the customer segments on the basis of the considerations described in the previous section. Then we represent each customer segment by a point in the plane. These points should be positioned in the plane in such a way that the distances between them are as close to the computed dissimilarities of the customer satisfaction profiles as possible. In general, it will not be possible to place the points so that the computed dissimilarities exactly coincide with the geometric distances. It is, for instance, impossible to place four points in the plane with the same (non-zero) distance between each of them.

Nevertheless, there are well-known techniques to position points in the plane so that the distances between the points approximate given (abstract) distances (or distances in a higher-dimensional space). Multidimensional scaling (MDS) and especially Sammon mappings (see [1]) belong to these techniques.

Given a collection of customer segments and a set of satisfaction questions, we compute the pairwise differences (dissimilarities) between the customer segments according to the method described in Section 3. Then we apply MDS based on these pairwise distances to visualise the dissimilarities in satisfaction profiles between all customer segments.

5 Results

The proposed approach was tested on data from over 10,000 customer questionnaires from eight different customer segments marked by the numbers 0,...,7. The customer segments have been found by a typical marketing analysis where demographic and product data are run through a cluster analysis and the identified clusters are later identified and labelled by marketing experts. The actual meaning of the segments is confidential. Each of the customer segments contains between 1300 and 1600 customers. The satisfaction profile for the customer segments is defined on the basis of four questions concerning satisfaction with different ordinal scales with 6-8 values plus a null answer, comprising no answer given and "don't know's".

Figure 1 shows the result of MDS applied to the computed dissimilarities for the eight customer segments on the left hand side. Each spot represents a customer segment. The closer two spots are the more similar are the satisfaction profiles of the corresponding customer segments. An alternative to the MDS approach is hierarchical clustering as it is shown in figure 2.

Figures 3 and 4 show the distributions over the answers to questions 0 and 3 (the distributions over the answers to questions 1 and 2 are not shown in this paper for rasons of limited space). Missing values are ignored for these distributions. Only the distribution over the ordinal scales are shown as relative frequencies. Each figure contains the distribution of all customer segments with respect to one question. Eight neighbouring bars represent the freqencies of the eight customer segments for one ordinal value of the corresponding question.

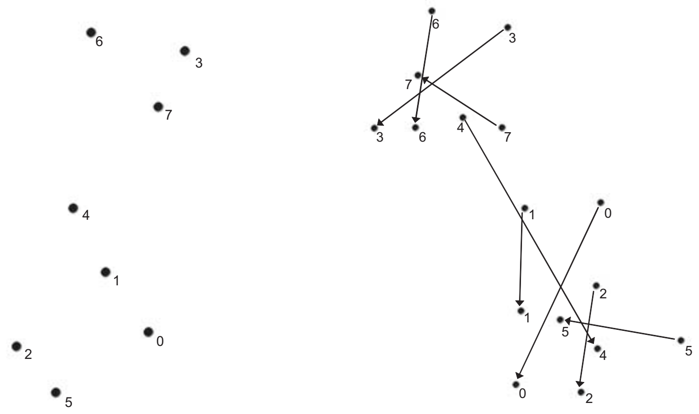

Fig. 1. Visualisation of the similarities of customer profiles (left) and changes of customer profiles (right) based on MDS.

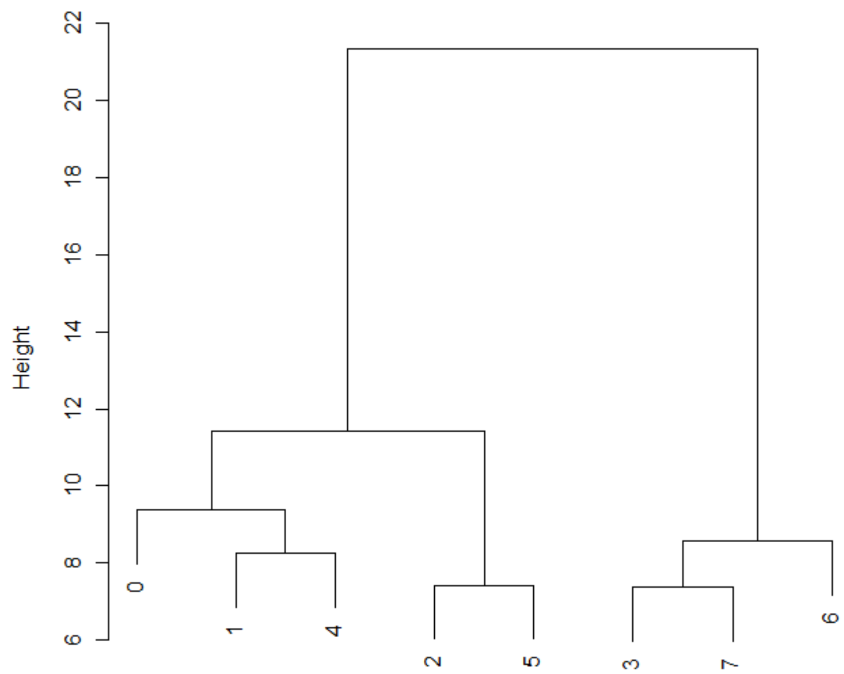

Fig. 2. Visualisation of the similarities of customer profiles based on hierarchical clustering.

It is also interesting to see how the satisfaction profiles of customer segments change over time. We apply the same technique as above, but simultaneously

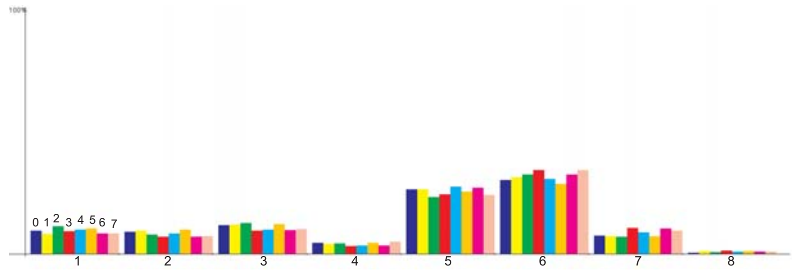

Fig. 3. Distributions over the answers to question 0.

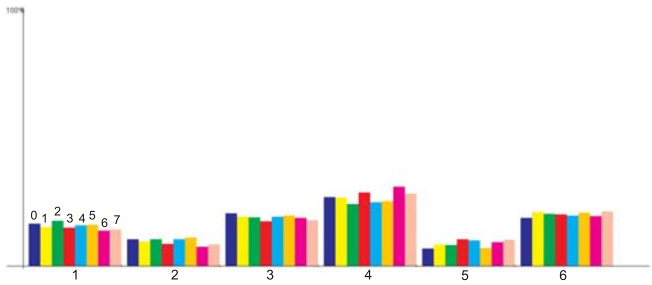

Fig. 4. Distributions over the answers to question 3.

for different time periods. Here we consider again eight customer segments with questionnaire result from two time periods. Therefore, instead of eight points in the MDS scatterplot, we now have 16 points, two points for each customer segment. Figure 1 shows the result of this analysis on the right hand side. The arrows in the diagram point from the earlier time period to the later one. The satisfaction profile of customer segment 4 has changed most. To a marketing analyst this is not surprising, because this segment is used to group the most volatile customers in terms of their lifestyle, attitudes and product usage.

The visualisation of similarity between customer segments and the changes of segments in regard to dimensions like customer satisfaction give very important feedback to business analysts. Customer segmentation projects a business view onto customers and represents true customer behaviour only to some extent. It is therefore important to constantly verify if the interpretation of segments has to be adjusted over time. The analysis presented here provides a quick and cheap alternative to re-segmentation. Customer segmentation is typically an expensive activity because it involves running in-depth surveys and sometimes purchasing additional marketing or demographic data. Businesses typically run some form of regular customer analysis on a smaller scale, for example, customer satisfaction surveys directly after engaging with individual customer. By analysing the similarity of customer segments in relation to available process or survey data we can quickly establish relationship and their changes between segments. This analysis can reveal interesting, previously unknown information or prompt a required re-segmentation because discovered relationships no longer align with the interpretation of segments.

6 Conclusions

The proposed approach to analysing similarities between customer satisfaction profiles of different customer segments has shown interesting results and justifies further investigation. The visualisation technique can also be used to track historical or hypothetical (what-if-analysis) changes in the satisfaction profiles of customer segments. For tracking purposes, instead of MDS more sophisticated methods like NeuroScale [6], MDS$_{polar}$ [9] or its extensions [10] that construct an explicit mapping from the high-dimensional space to the visualisation plane could be used.

References

1. Borg, I. and Groenen, P. (1997) Modern multidimensional scaling: Theory and applications. Springer, Berlin.
2. Cover, T. and Thomas, J. (1991) Elements of information theory. Wiley, New York.
3. von Hagen, F., Baaken, T., Holscher, V. and Carolin Plewa (2006) International research customer satisfaction surveys (Germany and Australia) and research provider surveys (Germany and Europe). Int. Journ. of Technology Intelligence and Planning 2, 210–224
4. Hwang, H.-G., Chen, R.-F. and Lee, J.M. (2007) satisfaction with internet banking: an exploratory study. Int. Journ. of Electronic Finance 1, 321–335
5. Lehmann, E.L. (2006) Nonparametrics: Statistical methods based on ranks. Springer, Berlin.
6. Lowe, D. and Tipping, M.E. (1996) Feed-forward neural networks topographic mapping for exploratory data analysis. Neural Computing and Applications, 4, 83-95.
7. Mann, H. B. and Whitney, D. R. (1947) On a test of whether one of two random variables is stochastically larger than the other. Annals of Mathematical Statistics 18, 50–60.
8. Nauck, D.D., Ruta, D., Spott, M. and Azvine, B. (2006) A Tool for Intelligent Customer Analytics. In: Proceedings of the IEEE International Conference on Intelligent Systems. IEEE Press, London, 518–521.
9. Rehm, Klawonn, F. and Kruse, R. (2005) MDS$_{polar}$: A new Approach for Dimension Reduction to Visualize High Dimensional Data. In: Famili, A.F., Kook, J.N., Peña, J.M. and Siebes A. (eds.) Advances in intelligent data analysis VI. Springer, Berlin, 316–327.
10. Rehm, Klawonn, F. and Kruse, R. (2006) POLARMAP – Efficient visualisation of high dimensional data. In: Banissi, E., Burkhard, R.A., Ursyn, A., Zhang, J.J., Bannatyne, M., Maple, C., Cowell, A.J., Tian, G.Y. and Hou, M. (eds): Information visualization. IEEE, London, 731–740.
11. Wilcoxon, F. (1945) Individual comparisons by ranking methods. Biometrics Bulletin 1, 80–83.