Visualizing Recommender System Results via Multidimensional Scaling

Wolfgang Gaul[1], Patrick Thoma[1], Lars Schmidt-Thieme[2], and Sven van den Bergh[3]

[1] Institut für Entscheidungstheorie und Unternehmensforschung, Universtät Karlsruhe (TH)
[2] Institut für Informatik, Universtität Freiburg
[3] Mentasys GmbH, Karlsruhe

Abstract. Web site visitors who look for desired items can formulate search queries which are taken by recommender systems to provide support within the underlying buying situation (e.g., enabling users to view recommended items and buy the ones they find most appropriate). Data from a large German retail online store is used to visualize products viewed most frequently together with search profiles that represent identical search queries of larger subgroups of site users. Comparisons between products viewed most frequently and those purchased most frequently can be used to improve the generation of recommendations. The results give interesting insights concerning the searching, viewing, and buying behavior of online shoppers.

1 Introduction

"Recommender Systems" is the label for a methodology that can, e. g., be designed (among others) to guide online users through complex web sites, huge online-stores or any other kind of information that can not be overviewed or searched through completely. As this is only one of the many situations in which recommender systems can be applied a framework that can be used to classify such systems would be of help [2]. ¿From an empirical point of view, we analyzed data gathered from a large German retail online store where a rule-based recommender system for a special product class was offered to support potential buyers. Customers received a sorted list of items of the underlying product class that "best" matched their search queries. We conducted multidimensional scaling (MDS)[1] to visualize the different search queries together with recommendable items of the product domain. This allowed us to analyze the searching behavior of site visitors on the basis of the most frequently chosen search queries, to compare user preferences with respect to the most appropriate products displayed within the MDS approach according to the distances between the positions of search queries and product locations in the underlying space, and to find out distinctions between viewing behavior (products viewed most frequently) and buying behavior (products purchased most frequently). The results revealed important hints on how to further improve the generation of recommendations.

2 Recommender Systems

A recommender system is software that collects and aggregates information about site visitors (e.g. buying histories, products of interest, hints concerning desired/desirable search dimensions or other FAQ) and their actual navigational and buying behavior and returns recommendations (e.g. based on customer demographics and/or past behavior of the actual visitor and/or user patterns of top sellers with fields of interest similar to those of the actual contact) [3]. Based on such a general description, different target groups for the application of recommender systems can be mentioned, e. g., web site visitors (help with respect to site navigation), online shop operators (support with respect to site structure optimization), or product managers (recommendations with respect to product line design). Taking these considerations as a starting point for our analyses, we were particulary interested in visualizing the lists of available, recommended, viewed and/or purchased products of a retail online shop to support the operation of a recommender system designed to assist site visitors with respect to their buying process. Multidimensional scaling seemed to be an appropriate methodology for this purpose.

3 Property Fitting and Search Query Positioning Within MDS

Multidimensional Scaling (MDS) is the label for a class of methods that represent similarity or dissimilarity values with respect to pairs of objects as distances between corresponding points in a low-dimensional metric space [1]. Graphical display of the object representations as provided by MDS enables to literally "look" at the data and to explore structures visually. A popular and classical technique to construct such object representations is the Kruskal method [4]. The goodness of fit of the solution obtained within the different iterative steps of the method can be assessed by the so called *Kruskal stress*.

One weakness of classical MDS is the difficulty to interpret the object representations in the low-dimensional space. This problem can be overcome by *property fitting*. Let $O = \{o_1, \ldots, o_N\}$ be the set of objects and $b_n = (b_{n1}, \ldots, b_{nM})$, $n = 1, \ldots, N$, the representation of object o_n in the underlying M-dimensional target space. If additional information about the objects is given, e.g., in form of attribute vectors $a_p = (a_{1p}, \ldots, a_{Np})'$ for property p, $p = 1, \ldots, P$, one can construct property vectors $c_p = (c_{p1}, \ldots, c_{pM})$ so that the projections

$$\hat{a}_{np} = \sum_{m=1}^{M} c_{pm} b_{nm} \tag{1}$$

of b_n onto c_p approximate the actual attribute values of the objects as good as possible with respect to the least squares criterion

$$\sum_{n=1}^{N} (\hat{a}_{np} - a_{np})^2. \tag{2}$$

Vector notation leads to $c_p = (B'B)^{-1}B'a_p$ with $B = (b_{nm})$. Quality of fit can be measured by correlation coefficients between $\hat{a}_{np}$ and a_{np}.

Similarly we performed the subsequent transformation of the search queries $s_q = (s_{q1}, \ldots, s_{q\tilde{P}_q})'$ where the notation search profile is used for a set of at least 10 identical search queries by different web users. Here, $\tilde{p} = 1, \ldots, \tilde{P}_q, \tilde{P}_q \leq P$, indicates properties specified in the query. In cases where a range of values was stated, e.g., for the price, we set $s_{q\tilde{p}}$ equal to the mean obtained from the lower and upper boundaries of the specified range. With given property vectors $c_{\tilde{p}}$ we looked for the representation $z_q = (z_{q1}, \ldots, z_{qM})'$ of s_q so that the projections

$$\hat{s}_{q\tilde{p}} = \sum_{m=1}^{M} z_{qm} c_{\tilde{p}m} \tag{3}$$

of z_q onto $c_{\tilde{p}}$ approximate the actual attribute values of the search profiles as good as possible with respect to the least squares criterion

$$\sum_{\tilde{p}=1}^{\tilde{P}_q} (\hat{s}_{q\tilde{p}} - s_{q\tilde{p}})^2. \tag{4}$$

Vector notation leads to $z_q = (C_q'C_q)^{-1}C_q's_q$ with $C_q = (c_{\tilde{p}m})$.

4 Empirical Analysis

4.1 Prerequisites

For our empirical analysis we used a set of transaction data gathered over a time period from February 5, 2003 until Juli 28, 2003 from a recommender system installed to support the searching behavior with respect to the product class "digital cameras" of website visitors of a large German retail online store. Products were described by price, manufacturer, optical zoom, digital zoom, resolution, memory, manageability and 19 other features which we aggregated to flash, manual control and extras. In this way we ended up with a bunch of ten product properties. Users can formulate a search query by specifying their preferred values in these 10 dimensions. They are then presented a sorted list of products that best match their needs according to an internal, rule-based algorithm. Implicit customer feedback is measured by counting the number of *product views* which are indicated by clicks on the product image to receive further, more detailed information about the product, and the number of purchase events which are indicated by clicks on the "to-market-basket"-icon. The product catalog consisted of 259 products. Based on the 10 dimensions discussed earlier only 181 different positionings of products remained Thus, it may happen in a few cases that the location of a frequently viewed or purchased product actually describes more than one product (with differences in those features that we aggregated to flash, manual control or extras). Nearly 70000 search queries appeared in the underlying time period (of which unfortunately about 49000 were empty and may result from web robots). Approximately 80000 product views and 1200 purchase events were recorded. To calculate dissimilarities between products we used the results of the *dist* function from the "R" software package [5]. Multidimensional scaling was done with the help of the function *isoMDS* from the library "MASS" [6] which is an implementation of Kruskal's non-metric MDS [4]. For the solution in the selected two-dimensional space *stress* was 0.167 and still "sufficient". Property fitting was implemented by ourselves as described in section 3. The correlation coefficients for the property vectors were high (> 0.7) for most of the properties. Best fits were achieved for price (0.87) and resolution (0.8), the worst for digital zoom (0.54) and brand (0.47).

4.2 Searching Behavior

As a starting point, we wanted to visualize the products from the product class just described together with the search queries that users mainly specify in the underlying situation. We aggregated identical queries to create so-called search profiles and selected those which represented at least 10 queries. This resulted in 144 search profiles representing 4991 of 22779 search queries (21.9%). These search profiles were mapped into the selected perceptual space as described in section 3. If only one search criteria was specified, we represented this profile as a point on the corresponding property vector instead of depicting a whole subspace (line through that point orthogonal to the considered property vector) since the space representation would have become unreadable otherwise. The result is shown in figure 1.

Two characteristics are remarkable: First, many search profiles are positioned on distinguishable lines. The ones on the negative prolongation of the price vector represent search profiles in which only the desired price was specified. The ones on the other dominant lines result from profiles, where mostly optical zoom and price were specified but only the price varied. Optical zoom was normally set to "two-fold" in these cases. Second, many search profiles are situated in an area which indicates demand for high memory where in contrast only few products are shown. This could be a hint, that customers look for cameras with higher memory than the actually available cameras can offer which could lead to dissatisfaction with the product assortment and reduce willingness to buy. This was confirmed by the observation of product

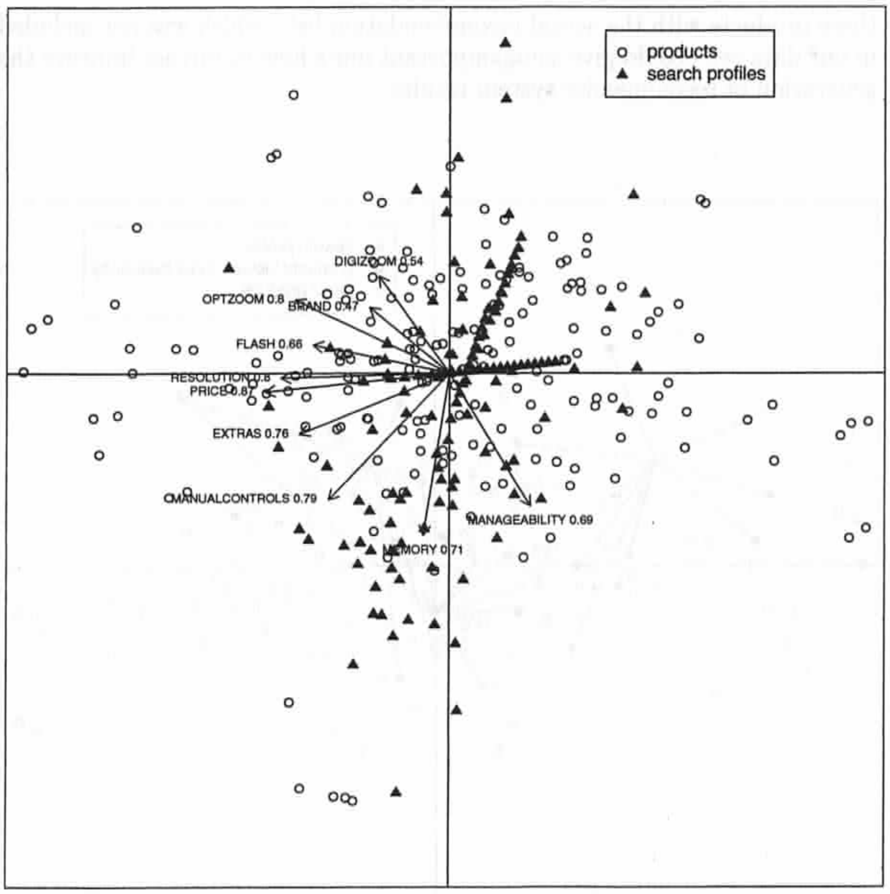

Fig. 1. Representation of search profiles and products

experts that most customers who purchased a digital camera also bought additional memory sticks/cards together with the selected camera.

4.3 Viewing and Buying Behavior

Next, we wanted to examine relations between searching, product viewing and buying behavior of the underlying web visitors. For each search profile we identified the products viewed most frequently and connected graphically their locations with the corresponding search profile position by lines. For three of the most frequent search profiles this is shown in Figure 2. One can see that in general users view products which are positioned relatively close to their specified search profile. However, for all search profiles there are products which are even closer and have not been viewed frequently and others that are relatively far away but had been viewed very often. Comparing these products with the actual recommendation list - which was not included in our data set - could give some important hints how to further improve the generation of recommender system results.

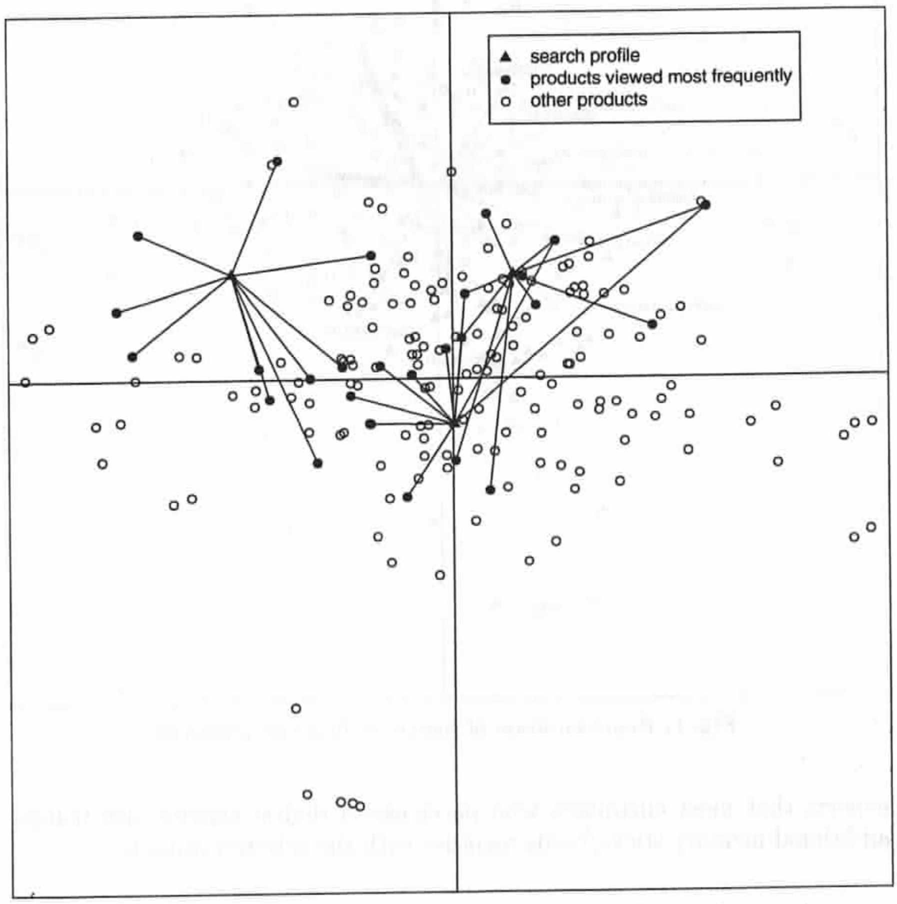

Fig. 2. Selected search profiles and products viewed most frequently

The same is valid for the corresponding analysis for the purchased products. Figure 3 shows that the products most frequently purchased are not necessarily the ones closest to the specified profiles. They might not even be among the most frequently viewed ones. This could, perhaps, be explained by what we sometimes call "Porsche effect": People are very interested in outstanding products but they buy the standard and cheaper alternatives.

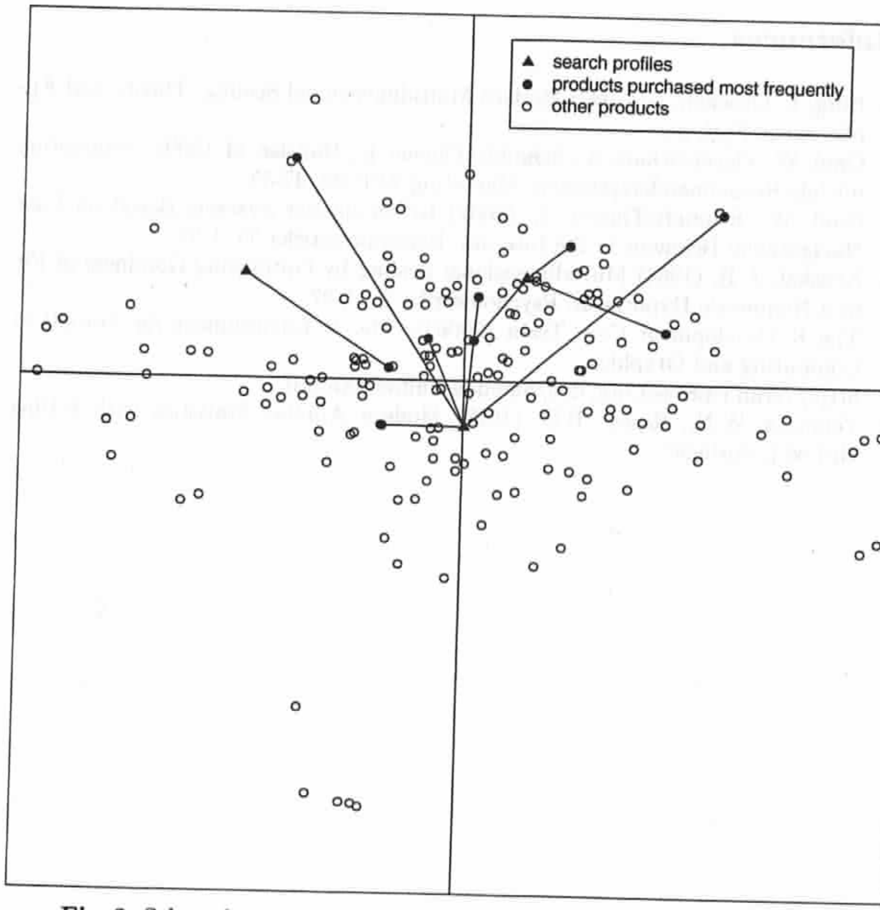

Fig. 3. Selected search profiles and products purchased most frequently

5 Conclusions

In this paper we argued that twodimensional representations of products of a given product class, search queries and recommender system results as well as those products viewed and/or purchased most frequently can be obtained by MDS and displayed together in appropriate spaces to help to visualize the performance of complex recommender systems and the usage of recommender system results by web site visitors. We showed that these visualizations can create valuable insights for product managers as well as for recommender system engineers. We restricted the presentation in different aspects. First, representations are based on distances in the underlying space. Second, only an actual snapshot of the system is represented. Further work will have to address both problems.

References

1. Borg, I., Groenen, P. (1997) Modern Multidimensional Scaling: Theory and Applications, Springer
2. Gaul, W., Geyer-Schulz, A., Schmidt-Thieme, L., Hahsler, M. (2002) eMarketing mittels Recommendersystemen, Marketing ZFP 24, 47-55
3. Gaul, W., Schmidt-Thieme, L. (2002) Recommender Systems Based on User Navigational Behavior in the Internet, Behaviormetrika 29, 1-22
4. Kruskal, J. B. (1964) Multidimensional Scaling by Optimizing Goodness of Fit to a Nonmetric Hypothesis. Psychometrika 29, 1-27
5. The R Development Core Team (20003), The R Environment for Statistical Computing and Graphics, http://cran.r-project.org/doc/manuals/fullrefman.pdf
6. Venables, W.N., Ripley, B.D. (1997) Modern Applied Statistics with S-Plus [3rd ed.], Springer