Preference Mapping for Automated Recommendation of Product Attributes for Designing Marketing Content

Moumita Sinha and Rishiraj Saha Roy
Adobe Research Labs, India
Bangalore, India - 560029.
{mousinha, rroy}@adobe.com

ABSTRACT

Identification of relevant product attributes is critical to the success of any marketing campaign. This task can be conceptualized as an attribute recommendation problem based on the *product's content or features*, where the goal of a solution would be to automatically recommend relevant features to the marketer for highlighting in a campaign. In this research, we try to solve this problem by using *preference mapping*, a powerful technique for associating feature preferences with users. We perform preference mapping with *sentiment scores* associated with product attributes mined from user reviews on the Web. As a result of this process, we are able to visualize a set of compared products and the appropriateness of the attributes on *the same two-dimensional space*, enabling us to easily recommend important features to a marketer. Finally, we show that expert recommendations or ratings for product features do not necessarily correlate with preference maps based on user sentiments.

Categories and Subject Descriptors

Information retrieval [**Retrieval tasks and goals**]: Recommender systems

General Terms

Algorithms, Experimentation, Human factors

Keywords

Preference Mapping, Sentiment Scores, Product Attributes

1. INTRODUCTION

Motivation. Product manufacturers are always faced with the dilemma of identifying which attribute(s) of their products they should highlight in their targeted marketing campaigns. For example, a digital camera has several defining aspects like power of zoom, size of display and image size in megapixels. A release of a new camera model by a manufacturer like Nikon will usually be followed by a marketing campaign to potential customers that will try to highlight certain aspects or attributes of the model. This attribute recommendation problem is critical to the success of the campaign. Focusing on features that do not appeal to users can result in a loss of large amount of ad spend and potential losses in product revenue for a manufacturer. In this paper, we address this challenge by proposing a principled technique called *preference mapping* [6], used in a novel way to automate the process of product attribute recommendation.

Related research. Alpert [1] presents one of the relatively early works emphasizing the importance of identifying relevant product attributes, and compares the effectiveness of direct and indirect questioning techniques. Cropper et al. [3] finds that a linear hedonic price function performs as well as a linear logit model in estimating consumer preferences for product attributes. But their analysis is based on simulations and does not draw connections between preferred attributes and campaign design. Zhang and Liu [12] try to identify product features that are associated with user sentiment by analyzing the contextual text associated with the mention of the product feature. While it could be meaningful to further scrutinize such attributes while designing product campaigns, the authors do not propose any method towards that end. Lehdonvirta [10] aims to discover product attributes that are likely to drive purchase decisions for virtual goods like online games and engaging activities on social media. However, the analysis presented by the author is purely from a sociological perspective and the author does not provide an algorithm for automating the above process. Recommendation algorithms similar to collaborative filtering have been used for designing campaigns, but they rely heavily on large amounts of existing customer preference data available with the advertiser [11]. On a related note, they are also known to have limitations such as data sparsity and model scalability, which leads to poor recommendations [2]. We provide a method for associating products with their marketable attributes that relate to each other based on publicly available sources. Such data sources may become accessible much before the advertiser receives direct information about customers' preferences based on product view or product purchase data. *Preference mapping* is an approach to identify customer preferences based on users' surveys of product attributes. Individual user differences are not averaged, but are directly incorporated into the mapping model and play vital roles in the preference fitting process [5]. As of date, the technique has only been used for understanding user preferences for diverse food items like lamb sausages [7], lager beer [6] and vanilla ice cream [4]. We believe that this

CBRecSys 2014, October 6, 2014, Silicon Valley, CA, USA.

method has a far greater potential and can be readily extended to unexplored application areas.

Approach. In this research, after specifying our product and attribute set, we acquire sentiment scores of *user reviews* that mention attributes for the products in our set. Following this, we associate *user sentiments with the attributes* mentioned in the reviews (instead of the product as a whole) and average them over reviewers who have written reviews concerning the attributes. We perform *preference mapping* on this processed dataset involving products, attributes and average sentiment scores and generate a *biplot visualization* that can be used for attribute recommendation. Finally, we compare our recommendations with *expert opinion* and show that there is no perfect correlation with what experts believe to be good features and what consumers like in a marketed product.

Organization. The rest of this paper is organized as follows. In Sec. 2, we describe our method of applying preference mapping to this situation. Next, we describe our data in Sec. 3 followed by experimental results and discussion in Sec. 4. Finally, we summarize our contribution and provide directions for future work in Sec. 5.

2. METHOD

We analyze a set of products p and a set of product attributes k. Customers who have bought these products often go to the product or retailer website to provide feedback about the product in the form of textual reviews. Most of these reviews generally contain mentions of product attributes. Further, positive or negative sentiments usually accompany the above mentions of the attributes. In our approach, we collect reviews where each sentence talks about only one attribute. Appropriate anaphora resolution is performed for review sentences when the attribute name is not directly mentioned [8]. Each sentence in each review is then assigned a sentiment score. Since each sentence mentions exactly one attribute, the sentiment score associated with the sentence is assumed to be the score associated with the attribute. Note that the effectiveness of our algorithm is not affected by the scale or range of this sentiment scoring. Next, the scores are averaged over the reviewers for each attribute for each product.

A preference mapping is then performed with the reviewer-averaged scores of each of the various attributes for the different products. We now explain how this is performed. As the first step, sentiment scores for all the product attributes are scaled to the same range so that variances are comparable across attributes of each product. Consider $X = (X_1, X_2, \ldots, X_p)^T$ as the matrix of the reviewer-averaged scores for the p products (say, different camera models) and the k attributes (like battery life, size of display and shutter delay). Thus each X_i is a vector with its elements as X_{ij}, which is the reviewer-averaged sentiment score for attribute j of product i. The principal component (PC) transformation of the feature vector X is the linear transformation $Y = \Gamma^T(X - \mu)$ where $\mu = E(X)$ and $\Sigma = Var(X) = \Gamma\Delta\Gamma'$. The transformation is such that $Var(Y)$ is maximized and the following holds:

$$\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_p$$

where, $Var(Y_j) = \lambda_j$, $j = 1, 2, \ldots, p$, $E(Y_j) = 0$ and $Cov(Y_j, Y_i) = 0$ when $i \neq j$.

Functions $Var(\cdot)$, $E(\cdot)$ and $Cov(\cdot)$ refer to the variance, expectation, and covariance functions, respectively, and the λ_j's represent the eigenvalues of the matrix X. These eigenvalues have the corresponding eigenvectors as $\gamma_1, \gamma_2, \ldots, \gamma_p$ (the number of eigenvectors is equal to the rank of the matrix X). Then the i^{th} PC for each product is the weighted sum of the scores of the product across the attributes, the weights being obtained from the i^{th} eigenvector. A biplot graph can be plotted for PC1 and PC2 with the weighted scores of each of the products and the eigenvector values for each attribute. The resultant graph provides an easily interpretable visualization that shows how products compare among each other based on customer reviews and the relative proximity of each attribute to their respective products with respect to associated positive user sentiment. Based on this multivariate visualization, marketing contents can be designed, highlighting favorable attributes for products. A schematic of the steps a marketer will undergo to utilize statistical analysis of social reviews to design product specific marketing campaigns is shown in Figure 1. Relevant steps have been explained in this section. Specific details about our dataset and experimental setup will be provided in the next section.

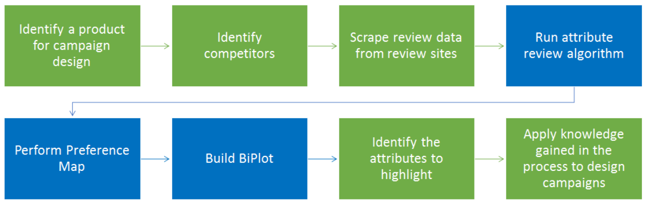

Figure 1: A schematic of the steps in our use case: The steps in green are part of the workflow, while those in blue are part of the proposed algorithm.

3. DATASET

We test our approach on a dataset consisting of 1309 reviews related to four digital camera models (Canon G3, Canon Powershot SD500, Canon S100, and Nikon Coolpix 4300), having a total of 13 distinct attributes. These attributes (or features) that we analyzed are: flash, zoom, battery, auto (quality of automatic mode), photo quality, view (quality of view through the viewfinder), delay (delay between photos), look, start (startup speed), color, night (quality of night photos), lens and resolution. The reviews are pre-processed to identify mentions of camera attributes within their texts. The 13 attributes are mentioned a total of 583 times in the product reviews that we collected.

Expert ratings. It is an interesting exercise to compare our attribute recommendation system to expert opinion. To this end, we went through popular digital camera review sites `dcresource`[1] and `imaging-resource`[2] for extracting expert ratings on the thirteen attributes for our

[1] `http://www.dcresource.com`, Accessed 11 July '14.
[2] `http://www.imaging-resource.com`, Accessed 11 July '14.

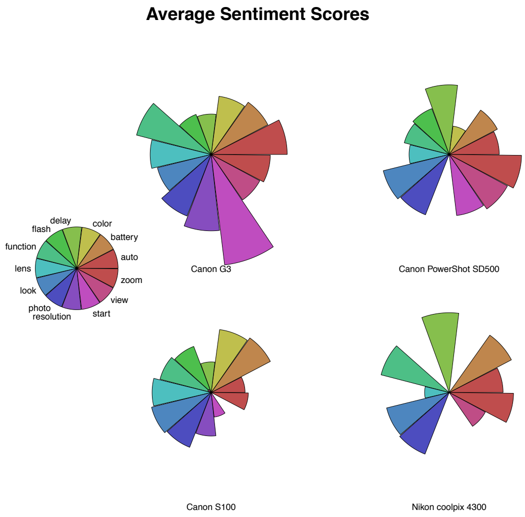

Figure 2: (Color online) Reviewer-averaged sentiment scores of attributes for our camera models.

four camera models. Since none of the popular camera review sites provide direct numeric ratings for attributes, we mapped expert opinion to a score of 1 or 2 depending upon the comments provided. For example, comments containing words like *exceptional*, *excellent* and *good* about an attribute were mapped to two, and *weak* and *worst* were assumed to be a one rating. The data that we collected has been made publicly available at `http://goo.gl/v8BGj4`.

4. EXPERIMENTS AND RESULTS

We assign a sentiment score to each sentence in each review in our dataset with the Alchemy API[3] and transfer the score to the attribute mentioned in the sentence. The higher the magnitude of the score, the stronger is the strength of the associated sentiment. Following this, the positive and negative sentiment scores of all the 52 ($= 13 \times 4$) camera-attribute pairs were averaged together over all the reviewers who mentioned the pair in his/her reviews, the neutral sentiments contributing zero to the sum. The missing observations are assumed to be neutral sentiments and hence the scores in such cases are assumed to be zero. These average sentiments for each camera over all attributes are shown in a radial chart in Figure 2. As a specific example, the battery of the Canon S100 was mentioned in 13 reviews, with seven, one, and five review(s) showing positive, negative and neutral scores respectively. While the numbers of positive and negative mentions seem comparable, the average positive and negative sentiment scores were found to be 1.3461 and 0.3569 respectively, indicating that the strength of the negative sentiment was not as strong as the positive sentiment. In our experiments, the two values were averaged to obtain 0.8515.

We now have a matrix with four rows (corresponding to each camera model) and thirteen columns (corresponding to each model attribute). The cells of this matrix are the reviewer-averaged sentiment scores associated with each camera and attribute pair. A principal component analysis (PCA) is then performed on this matrix of camera-attribute pairs. The PC1 and PC2 for this example, cumulatively explain 85% of the variability in the data. We then produce the biplot of the weighted scores of the products and the eigenvectors of each of the attributes, as shown in Figure 3.

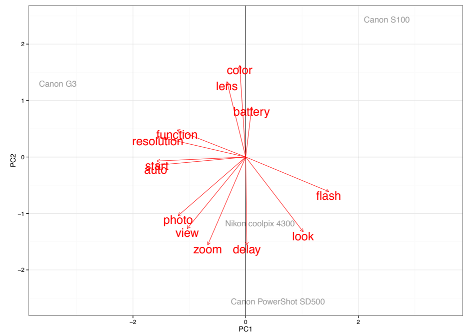

Figure 3: (Color online) A biplot of the weighted scores of products and eigenvector attributes. Attributes are in red and product names are in gray.

This graph provides a lot of information for design of marketing campaigns. First, in the graph, two attributes (in red) that are pointing towards the same direction, are attributes that tend to be highly positively correlated. A product that is in the same direction as an attribute, has a high value for this attribute. Thus, from the graph, we can conclude that attributes, which are closer and in the same direction as a product, are the ones that should be recommended for highlighting in marketing content for that particular model. For example, Canon G3 and Canon S100 received high sentiment scores on attributes like lens and color, while Nikon Coolpix 4300 and Canon PowerShot SD500 received high positive sentiments on low shutter delay and zoom quality. Thus, for example, lens and color should be recommended for designing marketing content in the campaign for Canon G3, rather than the zoom.

Second, this methodology also helps to contrast competing products simultaneously and provides *competitive intelligence* to the marketer. Thus, based on the given set of consumers' reviews, one can deduce that Nikon Coolpix 4300 and Canon PowerShot SD500 are similar with respect to the attributes studied, as compared to Canon G3 and Canon S100. For example, if Nikon Coolpix 4300 and Canon PowerShot SD500 are competing products, then it is meaninful to recommend only discriminatory features that add value to a particular product for its campaign. It is more sensible to recommend flash for Nikon Coolpix 4300 (more closer to the model than Canon 500) than the zoom, which is approximately equidistant from the both the products.

Analysis of expert opinion. From the data collected on expert comments (Sec. 3), we find that many of the discussed attributes are rated as 2, which implies that these attributes are "excellent" or "good" (Table 1). We assume that high expert score is analogous to high positive sentiment.

Table 2 shows the Kendall-Tau rank correlation coefficients between the preference mapping technique and the plain average sentiment scores (which is the unweighted sum of the attributes as opposed to the weighted sum for each camera). For three cameras we have statistically significant (at 0.05 level) correlation between the methods and a moder-

[3] `http://www.alchemyapi.com`

Table 1: Proportion of Attributes Rated as Excellent/Good and Poor.

Camera	Excellent/Good	Poor
Canon G3	0.385	0.538
Canon S100	0.615	0.231
Canon Powershot SD500	0.385	0.538
Nikon Coolpix 4300	0.615	0.385

Expert ratings were not available for all the attributes. So the sum of the values in a row may not add up to one

Table 2: Correlation between ranks of the attributes based on average sentiment scores and preference mapping scores.

Camera	Kendall-Tau	p-Value
Canon G3	0.564	0.007
Canon S100	0.615	0.003
Canon Powershot SD500	0.641	0.002
Nikon Coolpix 4300	0.294	0.172

ate correlation for Nikon Coolpix 4300. This shows that our method has high correlation with the intuitive understanding of the importance of the attributes and helps in further refinement. We could not observe any direct relation between the predictions based on the preference mapping and the attributes highly rated by experts.

5. CONCLUSIONS AND FUTURE WORK

The *preference mapping* technique, as described by us in this research, recommends *potentially "valuable" attributes* of products to marketers for *highlighting in a marketing campaign*. Our method provides the marketer the ability to design marketing content that can potentially increase response rates. We have used *sentiment scores* for product attributes, extracted from *review texts* to identify product features to be highlighted in campaigns. By focusing on attributes that are *known to have received positive sentiments of customers*, the risk in the campaign is minimized. Moreover, the comparison with the experts' comments suggests that sometimes, what customers value more about a product may be different from attributes that experts consider of high quality. So, designing marketing content taking into account what a large section of consumers show positive sentiments towards may help in engaging more effectively with a larger section of the consumers. The sentiment score in our research is a continuous variable and PCA has been used to identify appropriate attributes that have high scores. If some or all the scores are categorical in nature, multi-factor analysis [9] is preferable over PCA. The proposed technology does not require large amounts of customer preference data to be available internally with the advertiser (for example, customers who have viewed the same product or customers who have bought the same product), from their own sales and browsing patterns. Rather, we use reviews that *directly reflect customer preferences*. The reviews can be collected from any external source with consumers' opinion. The other major strength of our approach is that it is more likely to be positively viewed by the future customer. Such an approach enables having an informed conversation with the potential customer and is likely to improve customer satisfaction.

As future work, we would like to cluster products using attribute sentiment scores as features and observe the correlation of the clustering output to the representation produced by our preference mapping technique. Also, the quality of the reviews can be improved by choosing relevant users by mapping them to specific customer segments. This can lead to better insights on the data and finer levels of control in the design of marketing content.

Acknowledgements

We thank Ritwik Sinha from Adobe Research Labs India for valuable inputs at various stages of this work.

6. REFERENCES

[1] M. I. Alpert. Identification of determinant attributes: A comparison of methods. *Journal of Marketing Research*, pages 184–191, 1971.

[2] Y. H. Cho, J. K. Kim, and S. H. Kim. A personalized recommender system based on web usage mining and decision tree induction. *Expert Systems with Applications*, 23(3):329–342, 2002.

[3] M. L. Cropper, L. Deck, N. Kishor, and K. E. McConnell. Valuing product attributes using single market data: a comparison of hedonic and discrete choice approaches. *The Review of economics and Statistics*, pages 225–232, 1993.

[4] L. Dooley, Y. S. Lee, and J. F. Meullenet. The application of check-all-that-apply (CATA) consumer profiling to preference mapping of vanilla ice cream and its comparison to classical external preference mapping. *Food quality and preference*, 21(4):394–401, 2010.

[5] K. Greenhoff and H. MacFie. Preference mapping in practice. In H. MacFie and D. Thomson, editors, *Measurement of Food Preferences*, pages 137–166. Springer US, 1994.

[6] J. X. Guinard, B. Uotani, and P. Schlich. Internal and external mapping of preferences for commercial lager beers: comparison of hedonic ratings by consumers blind versus with knowledge of brand and price. *Food Quality and Preference*, 12(4):243–255, 2001.

[7] H. Helgesen, R. Solheim, and T. NÃęs. Consumer preference mapping of dry fermented lamb sausages. *Food Quality and Preference*, 8(2):97–109, 1997.

[8] S. Lappin and H. J. Leass. An algorithm for pronominal anaphora resolution. *Comput. Linguist.*, 20(4):535–561, Dec. 1994.

[9] S. Lê, J. Josse, F. Husson, et al. Factominer: an r package for multivariate analysis. *Journal of statistical software*, 25(1):1–18, 2008.

[10] V. Lehdonvirta. Virtual item sales as a revenue model: identifying attributes that drive purchase decisions. *Electronic Commerce Research*, pages 97–113, 2009.

[11] G. Linden, B. Smith, and J. York. Amazon. com recommendations: Item-to-item collaborative filtering. *Internet Computing, IEEE*, 7(1):76–80, 2003.

[12] L. Zhang and B. Liu. Identifying noun product features that imply opinions. In *HLT '11*, pages 575–580, 2011.