SHOPSMART: Product Recommendations through Technical Specifications and User Reviews

Alexander Yates
Temple University
Computer and Information Sci.
303A Wachman Hall
1805 N. Broad St.
Philadelphia, PA 19122
yates@temple.edu

James Joseph
Temple University
Computer and Information Sci.
324 Wachman Hall
1805 N. Broad St.
Philadelphia, PA 19122
james.joseph@temple.edu

Ana-Maria Popescu
Yahoo, Inc.
2821 Mission College
Santa Clara, CA
amp@yahoo-inc.com

Alexander D. Cohn
Temple University
Computer and Information Sci.
324 Wachman Hall
1805 N. Broad St.
Philadelphia, PA 19122
tua58997@temple.edu

Nick Sillick
Temple University
Computer and Information Sci.
324 Wachman Hall
1805 N. Broad St.
Philadelphia, PA 19122
nick@sillik.org

ABSTRACT

This paper describes a new method for providing recommendations tailored to a user's preferences using text mining techniques and online technical specifications of products. We first learn a model that can predict the price of a product given automatically-determined features describing technical specifications and users' opinions. We then use this model to rank a list of products based on individual users' preferences about various features. On a data set collected from Amazon reviews and online technical specifications, rankings produced by this model rank the best product for a user in the 87th percentile of products in its category, on average. Our approach outperforms several comparison systems by 21 percentiles or more.

Categories and Subject Descriptors

I.2.1 [**Artificial Intelligence**]: Applications and Expert Systems

General Terms

Experimentation

Keywords

Recommendation, e-Commerce, personalization.

1. INTRODUCTION

Online shoppers can read user reviews, technical specifications, editorial reviews, and a variety of other resources to help them shop for the best product. Yet the sheer volume of the information available, as well as the technical complexity of the products themselves (*e.g.*, *plasma vs. LCD TV? shorter focal length or faster shutter speed?*), can create a daunting problem for even the most intrepid shopper. Recommender systems, including collaborative filtering techniques [2] and content-based methods [1], address this problem by automatically providing recommendations to users for which product to buy, or at least to research more closely. Existing approaches fall short of shoppers' needs, however, in cases where users have little or no transaction histories, or where shoppers' current preferences are significantly different from their preferences when they purchased previous items. Since many kinds of products purchased online (*e.g.*TVs, digital cameras, cellphones) are relatively big-ticket items, they are purchased infrequently, and a shopper's transaction history is a poor indicator of her or his current preferences. For other frequent online transactions, such as DVD rentals, recommender systems face the problem that a shopper's preferences may change with each new transaction (as in, "No crime drama tonight; I feel like a sci-fi movie.").

To address this challenge, we present the SHOPSMART system for making product recommendations using mined user reviews, technical specifications, and explicit shopper's preferences. Specifically, SHOPSMART addresses the following key questions: 1) Is it possible to provide helpful recommendations to shoppers in the absence of information on their transaction history? 2) How can we elicit a user's preferences without requiring the user to rank previous purchases?

2. PRODUCT VALUE MODEL

SHOPSMART proceeds in two steps. It first uses a learned model, called a *product value model* to predict the intrinsic value of a product to the *average* user. It then uses a shopper's preferences to personalize the predictions and present a ranked list of products in decreasing order of suitability for that shopper. We describe the personalization step in the next section.

CIKM'08, October 26–30, 2008, Napa Valley, California, USA.
ACM 978-1-59593-991-3/08/10.

ShopSmart's product value model learns from a set of automatically extracted product features. In training, we use product price as an indicator of value, since in a competitive market the price marks the amount a buyer and seller are willing to trade for. This works well for consumer electronics markets, where products are highly differentiated and each shopper purchases only a small number of items, but other economic indicators can be used (such as revenue for movies) for other markets. ShopSmart uses an SVM regression model and subset-evaluation feature selection for its product value model. Experiments on a set of digital cameras, flatscreen TVs, and LCD monitors showed that the product value model achieves an accuracy of over 90% for two categories, and 81% on LCD monitors.

3. PERSONALIZED VALUE MODEL

While the product value model can provide accurate predictions for a product's value to the average user, ShopSmart's ultimate goal is to personalize these predictions for each shopper. Its *personalized value model* incorporates a shopper's preferences into the predictions.

ShopSmart elicits preferences from the shopper of the form, "How much do you care about feature X?" It stores the resulting preferences in a vector, with a value between 0 and 10 for every feature in the product category's feature set. For instance, for digital cameras, a shopper might indicate a preference level of 3 for weight and 8 for battery life if they were particularly interested in long-lasting cameras but did not care especially about how heavy the camera is. Importantly, shopper's preferences are elicited at the time they are seeking a purchase, so that they reflect the shopper's current tastes and interests. In contrast, systems based on transaction histories cannot anticipate or cope with changes in shoppers' preferences.

ShopSmart's personalized value model combines a shopper's preferences with the product value model's predictions. Let $V(\hat{x})$ represent the value of a product $\hat{x} =< x_1, \ldots, x_n >$ according to the product value model, and let $\hat{y} =< y_1, \ldots, y_n >$ represent a shopper's preferences (on a scale of 0 to 10) for each feature of a product category. The personalized value model adjusts each product's feature vector $\hat{x}$ according to the shopper's preferences as follows. Let $\hat{F} =< f_1, \ldots, f_n >$ be the adjusted set of features for a product. Each f_i is given by:

$$f_i\left(x_i, y_i\right) = \frac{1}{2} + \frac{y_i}{5} * \left(x_i - \frac{1}{2}\right) \qquad (1)$$

The personalized value model provides recommendations to a shopper by combining preferences and product features using the above equation, and then measuring the *change in predicted value* according to the product value model:

$$\text{Change in Value}\left(\hat{x}, \hat{y}\right) = \frac{V\left(\hat{F}\right) - V\left(\hat{x}\right)}{V\left(\hat{x}\right)} \qquad (2)$$

It uses this change in value to score each product $\hat{x}$, and then sorts them in decreasing order by these scores. The highest-scoring product is most recommended.

We tested ShopSmart's ranking performance on a data set of automatically collected product feature vectors and a set of user preference vectors. We followed Liu *et al.* [3] in mining online customer reviews to automatically identify product features and the average user opinion about each feature (*e.g.*, on average, users rated the "picture quality" of camera X as a 7 out of 10). We augmented these product features with technical features (*e.g.*, physical dimensions, megapixels, screen size, *etc.*) collected from online product manuals. In total, we collected 55 examples of cameras, 105 flatscreen TVs, and 78 LCD monitors. We also manually collected example user preference vectors by examining a separate set of reviews for phrases like "really cared about the battery life". We adjusted user preferences up or down according to such phrases, and labeled the user preferences with the products that the user bought. In total, we collected 30 user preference vectors for digital cameras, 36 for TVs, and 30 for LCD monitors.

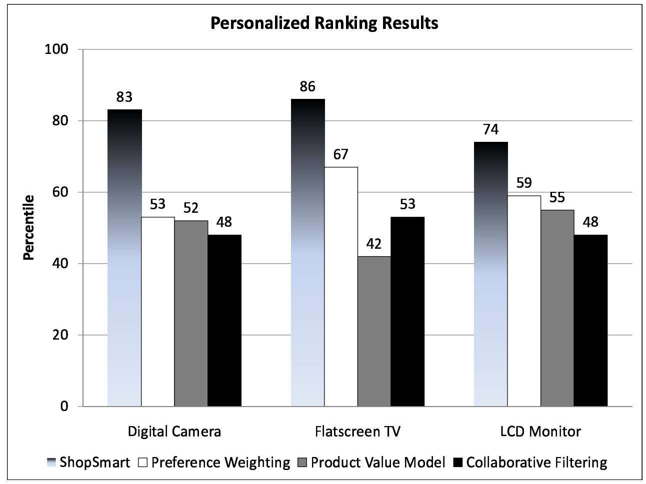

Figure 1: Average percentile score for four product recommendation techniques. ShopSmart outperforms the next-best system by 21 percent on average across the 3 product categories. The percentile score represents the percent of all products that were ranked below the best product.

We measured the accuracy of our model by ranking products for each preference vector according to Equation 2, and measuring how far down the list the target product appeared. We averaged this percentile score over all preference vectors. We compare ShopSmart against three other ranking techniques: 1) a baseline (called "Preference Based" ranking) that uses the dot product of a preference vector and a product vector as the function to rank products by; 2) the Product Value Model without any preference vector information; and 3) a collaborative filtering technique.

Figure 1 shows the performance of these four ranking techniques. ShopSmart significantly outperforms the next-best system, the Preference Weighting baseline, by an impressive 21 percentiles on average across the three product categories. It is almost always able to rank the best product near the top of the list.

4. REFERENCES

[1] M. Balabanovi and Y. Shoham. Fab: content-based, collaborative recommendation. *Commun. ACM*, 40(3):66–72, 1997.

[2] G. Linden, B. Smith, and J. York. Amazon Recommendations: Item-to-Item Collaborative Filtering. *Internet Computing*, 7(1):76–80, 2003.

[3] B. Liu, M. Hu, and J. Cheng. Opinion observer: Analyzing and comparing opinions on the web. In *WWW*, 2005.