

Recommender Systems in E-Commerce

J. Ben Schafer, Joseph Konstan, John Riedl

GroupLens Research Project
Department of Computer Science and Engineering
University of Minnesota

Minneapolis, MN 55455
1-612-625-4002

{schafer, konstan, riedl}@cs.umn.edu

ABSTRACT

Recommender systems are changing from novelties used by a few E-commerce sites, to serious business tools that are re-shaping the world of E-commerce. Many of the largest commerce Web sites are already using recommender systems to help their customers find products to purchase. A recommender system learns from a customer and recommends products that she will find most valuable from among the available products. In this paper we present an explanation of how recommender systems help E-commerce sites increase sales, and analyze six sites that use recommender systems including several sites that use more than one recommender system. Based on the examples, we create a taxonomy of recommender systems, including the interfaces they present to customers, the technologies used to create the recommendations, and the inputs they need from customers. We conclude with ideas for new applications of recommender systems to E-commerce.

Keywords

Electronic commerce, recommender systems, interface, customer loyalty, cross-sell, up-sell, mass customization.

1. INTRODUCTION

In his book Mass Customization (Pine, 1993), Joe Pine argues that companies need to shift from the old world of mass production where “standardized products, homogeneous markets, and long product life and development cycles were the rule” to the new world where “variety and customization supplant standardized products.” Pine argues that building one product is simply not adequate anymore. Companies need to be able to, at a minimum, develop *multiple* products that meet the *multiple* needs of *multiple* customers. The movement toward E-commerce has allowed companies to provide customers with more options. However, in expanding to this new level of customization, businesses increase the amount of information that customers must process before they are able to select which items meet

their needs. One solution to this information overload problem is the use of *recommender systems*.

Recommender systems are used by E-commerce sites to suggest products to their customers. The products can be recommended based on the top overall sellers on a site, based on the demographics of the customer, or based on an analysis of the past buying behavior of the customer as a prediction for future buying behavior. Broadly, these techniques are part of personalization on a site, because they help the site adapt itself to each customer. Recommender systems automate personalization on the Web, enabling individual personalization for each customer. Personalization to this extent is one way to realize Pine’s ideas on the Web. Thus, Pine would probably agree with Jeff Bezos, CEO of Amazon.com™, when he said “If I have 2 million customers on the Web, I should have 2 million stores on the Web.”

Recommender systems enhance E-commerce sales in three ways:

Browsers into buyers: Visitors to a Web site often look over the site without ever purchasing anything. Recommender systems can help customers find products they wish to purchase.

Cross-sell: Recommender systems improve cross-sell by suggesting additional products for the customer to purchase. If the recommendations are good, the average order size should increase. For instance, a site might recommend additional products in the checkout process, based on those products already in the shopping cart.

Loyalty: In a world where a site’s competitors are only a click or two away, gaining customer loyalty is an essential business strategy (Reichheld and Sesser, 1990) (Reichheld, 1993). Recommender systems improve loyalty by creating a value-added relationship between the site and the customer. Sites invest in learning about their users, use recommender systems to operationalize that learning, and present custom interfaces that match customer needs. Customers repay these sites by returning to the ones that best match their needs. The more a customer uses the recommendation system – teaching it what they want – the more loyal they are to the site. “Even if a competitor were to build the exact same capabilities, a customer ... would have to spend an inordinate amount of time and energy teaching the competitor what the company already knows.” (Pine, et al. 1995) Finally, creating relationships between customers can also increase loyalty. Customers will return to the site that recommends people with whom they will like to interact.

This paper makes five contributions to the understanding of recommender systems in E-commerce. First, we provide a set of recommender system examples that span the range of different applications of recommender systems in E-commerce. Second, we analyze the way in which each of the examples uses the recommender system to enhance revenue on the site. Third, we describe a mapping from applications of recommender systems to a taxonomy of ways of implementing the applications. Fourth, we examine the effort required from users to find recommendations. Fifth, we describe a set of suggestions for new recommender system applications based on parts of our taxonomy that have not been explored by the existing applications.

The paper is useful to two groups: academics studying recommender systems in E-commerce, and implementers considering applying recommender systems in their site. For academics, the examples and taxonomies provide a useful initial framework within which their research can be placed. The framework will undoubtedly be expanded to include future applications of recommender systems. For implementers, the paper provides a way of making choices among the available applications and technologies. An implementer can choose a moneymaking goal, select the interfaces that will help achieve that goal, and select an implementation technique that supports the goal within the interface.

2. Recommender System Examples

In the following section we present six e-commerce businesses that utilize one or more variations of recommender system technology in their web sites. For each site, and each variation, we give a brief description of the features of the system. In later sections we refer to these examples as we explain the types of recommendations provided, the type of technology used, and the types of information gathered. For organizational purposes these sites have been alphabetized. The examples listed were correct as of May 31, 1999. Due to the rapid changes in the Internet they may no longer be valid.

2.1 Amazon.com

We focus on recommender systems in the *book* section of Amazon.com.

Customers who Bought: Like many E-commerce sites, Amazon.com™ (www.amazon.com) is structured with an information page for each book, giving details of the text and purchase information. The Customers who Bought feature is found on the information page for each book in their catalog. It is in fact two separate recommendation lists. The first recommends books frequently purchased by customers who purchased the selected book. The second recommends authors whose books are frequently purchased by customers who purchased works by the author of the selected book.

Eyes: The Eyes feature allows customers to be notified via email of new items added to the Amazon.com catalog. Customers enter requests based upon author, title, subject, ISBN, or publication date information. Customers can use both simple and more complex Boolean-based criteria (AND/OR) for notification queries. Requests can also be directly entered from any search results screen, creating a persistent request based on the search.

Amazon.com Delivers: Amazon.com Delivers is a variation on the Eyes feature. Customers select checkboxes to choose from a list of specific categories/genres (Oprah books, biographies, cooking). Periodically the editors at Amazon.com send email announcements to notify subscribers of their latest recommendations in the subscribed categories.

Book Matcher: The Book Matcher feature allows customers to give direct feedback about books they have read. Customers rate books they have read on a 5-point scale from “hated it” to “loved it.” After rating a sample of books, customers may request recommendations for books they might like. At that point a half dozen non-rated texts are presented which correlate with the user’s indicated tastes. Feedback to these recommendations is provided by a “rate these books” feature where customers can indicate a rating for one or more of the recommended books.

Customer Comments: The Customer Comments feature allows customers to receive text recommendations based on the opinions of other customers. Located on the information page for each book is a list of 1-5 star ratings and written comments provided by customers who have read the book in question and submitted a review. Customers have the option of incorporating these recommendations into their purchase decision.

2.2 CDNOW

Album Advisor: The Album Advisor feature of CDNOW™ (www.cdnw.com) works in two different modes. In the single album mode customers locate the information page for a given album. The system recommends 10 other albums related to the album in question. In the multiple artist mode customers enter up to three artists. In turn, the system recommends 10 albums related to the artists in question.

My CDNOW: My CDNOW enables customers to set up their own music store, based on albums and artists they like. Customers indicate which albums they own, and which artists are their favorites. Purchases from CDNOW are entered automatically into the “own it” list. Although “own it” ratings are initially treated as an indication of positive likes, customers can go back and distinguish between “own it and like it” and “own it but dislike it.” When customers request recommendations the system will predict 6 albums the customer might like based on what is already owned. A feedback option is available by customers providing a “own it,” “move to wish list” or “not for me” comment for any of the albums in this prediction list. The albums recommended change based on the feedback.

2.3 eBay

Feedback Profile: The Feedback Profile feature at eBay.com™ (www.ebay.com) allows both buyers and sellers to contribute to feedback profiles of other customers with whom they have done business. The feedback consists of a satisfaction rating (satisfied/neutral/dissatisfied) as well as a specific comment about the other customer. Feedback is used to provide a recommender system for purchasers, who are able to view the profile of sellers. This profile consists of a table of the number of each rating in the past 7 days, past month, and past 6 months, as well as an overall summary (e.g., 867 positives from 776 unique customers). Upon further request, customers can browse the individual ratings and comments for the sellers.

2.4 Levis

Style Finder: Style Finder allows customers of the Levi Straus™ (www.levis.com) website to receive recommendations on articles of Levi's clothing. Customers indicate whether they are male or female, then view three categories -- Music, Looks, Fun -- and rate a minimum of 4 "terms" or "sub-categories" within each. They do this by providing a rating on a 7-point scale ranging from "leave it" to "love it." They may also choose the rating of "no opinion." Once the minimum number of ratings are entered customers may select "get recommendations." Here, they are provided with thumbnails of 6 items of recommended clothing. Customers may provide feedback by use of the "tell us what you think feature" which allows them to enter an opinion rating for the recommended article of clothing. Feedback may change one or all of the six items recommended.

2.5 Moviefinder.com

Match Maker: Moviefinder.com's Match Maker (www.moviefinder.com) allows customers to locate movies with a similar "mood, theme, genre or cast" to a given movie. From the information page of the movie in question, customers click on the Match Maker icon and are provided with the list of recommended movies, as well as links to other films by the original film's director and key actors.

We Predict: We Predict recommends movies to customers based on their previously indicated interests. Customers enter a rating on a 5-point scale -- from A to F -- for movies they have viewed. These ratings are used in two different ways. Most simply, as they continue, the information page for non-rated movies contains a personalized textual prediction (go see it -- forget it). In a variation of this, customers can use Powerfind to search for top picks based on syntactic criteria such as Genre, directors, or actors and choose to have these sorted by their personalized prediction or by the all customer average.

2.6 Reel.com

Movie Matches: Similar to Amazon.com's Customers who Bought, Reel.com's Movie Matches (www.reel.com) provides recommendations on the information page for each movie. These recommendations consist of "close matches" and/or "creative matches." Each set consists of up to a dozen hyperlinks to the information pages for each of these "matched" films. The hyperlinks are annotated with one sentence descriptions of how the new movie is similar to the original movie in question ("Darker thriller raises similarly disturbing questions...").

Movie Map: The Movie Map feature of Reel.com recommends movies to customers based on syntactic features. Customers enter queries based on Genre, movie types, viewing format and/or prices, and request results be constrained to "sleepers" or "best of this genre." The recommendations are editor's recommendations for movies that fit the specified criteria.

2.7 Summary

In Table 1 we have summarized the applications, interfaces, recommendation technology, and how users find recommendations for all of the example applications. The first column just names each application, under the E-commerce site that houses it. The second column describes the interface that is used for delivering the recommendations. The third column

describes the recommendation technology used by the site, and the inputs required by that technology. The fourth column describes how users find recommendations using the application. Each of the columns of Table 1 is the subject of one of the sections of this paper, describing the meaning of the entries in the table, and their role in supporting recommender systems for E-commerce.

3. Recommendation Interfaces and Ways to Make Money

An old proverb states that there is "more than one way to skin a cat." One would assume that the method selected depends on the desired outcome. Similarly, there is more than one way to display recommendations to a customer. The method selected may well depend on how the e-commerce site wants the customer to use the recommendation. In the following we will examine seven recommendation interfaces, and how each assists the site in making money. While some of these methods have their roots in traditional commerce, each of them draws upon the strengths of the electronic medium to provide more powerful recommendations.

Browsing: In traditional commerce a customer might walk into a video store and ask the clerk to recommend "a comedy from the 50s." Ideally, the clerk would recommend several movies, and the customer could go off to locate the recommended movies, browse the box covers, and see which ones appealed to them. However, the quality of the recommendations provided was dependent on the particular clerk's knowledge of an enormous range of movies. Reel.com has several advantages when implementing *browsing* into their Movie Map feature. First, the recommendations of several clerks/editors can be combined so that higher quality recommendations can be provided no matter what the query parameters. Furthermore, recommendations are returned with immediate links to the items being recommended -- no more searching the store for the obscure videos recommended. Recommended browsing helps the E-commerce site by converting browsers into buyers. It does so by helping the users narrow down their choices and feel more confident in their decision to buy by providing organized access to the recommendations.

Similar Item: Another modification of traditional commerce techniques is the *similar item* recommendation. Systems such as Reel.com's Movie Matcher, Amazon.com's Customer's who Bought and one variation of CDNOW's Album Advisor attempt to expose customers to items they may have forgotten about, or of which they may have simply been unaware. Their implementation in E-commerce sites allows for more specific and personalized recommendations. The items displayed can be entirely selected based on the item(s) in which a customer has shown interest. In doing so, sites increase customer's exposure to their product line, and ideally are able to sell more items per order.

Email: Recommendations can also be delivered directly to customers through *email*, in a extension of traditional direct mail techniques. Amazon.com's Eyes feature allows them to notify customers the minute an item becomes commercially available. Eyes enables Amazon.com to attract customers into their store before other stores with the same product can reach those

customers.. Furthermore, both Eyes and Amazon.com Delivers allows the site to keep a customer aware of the site and of items the customer may have missed. Customers appreciate the email recommendations because they help them watch out for new items they are interested in purchasing. These features assist the

site in making money by increasing both loyalty, and the number of return visits.

Business/Applications	Recommendation Interface	Recommendation Technology	Finding Recommendations
Amazon.com			
Customers who Bought	Similar Item	Item to Item Correlation <i>Purchase data</i>	Organic Navigation
Eyes	Email	Attribute Based	Keywords/freeform
Amazon.com Delivers	Email	Attribute Based	Selection options
Book Matcher	Top N List	People to People Correlation <i>Likert</i>	Request List
Customer Comments	Average Rating Text Comments	Aggregated Rating <i>Likert</i> <i>Text</i>	Organic Navigation
CDNOW			
Album Advisor	Similar Item Top N List	Item to Item Correlation <i>Purchase data</i>	Organic Navigation Keywords/freeform
My CDNOW	Top N List	People to People Correlation <i>Likert</i>	Organic Navigation Request List
eBay			
Feedback Profile	Average Rating Text Comments	Aggregated Rating <i>Likert</i> <i>Text</i>	Organic Navigation
Levis			
Style Finder	Top N List	People to People Correlation <i>Likert</i>	Request List
Moviefinder.com			
Match Maker	Similar Item	Item to Item Correlation <i>Editor's choice</i>	Navigate to an item
We Predict	Top N List Ordered Search Results Average Rating	People to People Correlation <i>Aggregated Rating</i> <i>Likert</i>	Keywords/freeform Selection options Organic Navigation
Reel.com			
Movie Matches	Similar Item	Item to Item Correlation <i>Editor's choice</i>	Organic Navigation
Movie Map	Browsing	Attribute Based <i>Editor's choice</i>	Keywords/freeform

Table 1: Recommender System Examples

Text Comments: More and more frequently, sites are providing customers with recommendations based directly on the *text comments* of other customers. Amazon.com's Customer Comments and eBay's Feedback Profile streamlines the gathering of "the word on the street" by allowing customers to locate an item of interest and browse the comments of other customers. This helps sites make money by providing impartial information on the goods/services being sold – the thought being, if enough people claim that a book is good, or a seller is credible, than it is likely to be true. This not only helps convert browsers into buyers, but should increase loyalty to a site. If customers learn they can trust these third party recommendations, than they are more likely to return the next time they are faced with a questionable decision..

Average Rating: Even simpler access to "the word on the street" is the *average rating* feature. Rather than asking customers to browse a list of text based opinions, other customers can provide numerical ranking opinions. By aggregating these rankings into an average rating, Customer Comments and Feedback Profile both provide users with a "one stop" check on the quality of an item. Similar to text comments, average ratings should facilitate in converting browsers into buyers, and increasing customer loyalty to the site.

Top-N: Amazon.com's Book Matcher, Levi's Style Finder and My CDNOW, among others, take advantage of recommendations through a *top-N* list. Once each site has learned details about a customer's likes and dislikes, each is able to provide the

customer with a personalized list of the top number of unrated items for that customer. It is as though one could gather all of the clothes that might interest a given client onto a single rack without distracting them with items they will not be interested in. This helps the vendor in several ways. First, it is another example of converting browsers into buyers – it provides increased exposure to the vendor’s wares, but only to those items that should truly interest the user. Second, it may help the customer in making a decision about items that they originally held in doubt – the suggestion from the site may be another point in favor of the item.

Ordered Search Results: Finally, a less restrictive variation of the top-N list are *Ordered Search Results* recommendations. While top-N limits the predictions to some predefined number, ordered search results allow the customer to continue to look at items highly likely to be of interest to them. Moviefinder.com’s We Predict feature allows customers to have query returns sorted by the predicted likelihood that the customer will enjoy the item. Once again, this helps convert browsers into buyers.




Figure 1: Recommendation Taxonomy

4. A Taxonomy for Mapping Applications to Recommendation Techniques

In this section we describe the Recommendation Technology column of Table 1 in detail. We first lay out the different recommender system applications in a taxonomy of recommendation types. We then describe the different user inputs, which are the italicized entries in the table. The goal of the taxonomy is to present a completely user-focused analysis of the different recommender systems, so the taxonomy is based on the features most important to customers of the E-commerce sites. The two key dimensions in the taxonomy are the degree of automation, and the degree of persistence in the recommendations (Figure 1).

The automation axis ranges from completely Automatic recommendations to completely Manual recommendations. From the perspective of the customer, Automatic means that the

recommendation is generated without explicit effort by the customer. The customer just interacts with the site as he or she wishes, and suddenly a recommendation appears that is appropriate for the customer’s interests. Manual means that the customer takes explicit effort to seek out recommendations that will fit her interests. Note that recommendations that are Manual from the perspective of the user may be generated by the site using a computer program. We consider these Manual, since we are taking the customer’s perspective. Likewise, recommendations that appear automatically for the customer, but that are generated by hand by the site are considered Automatic. Whether the site uses a computer or a human to implement its recommendation algorithms is unimportant to the customer.

The persistence axis ranges from completely Ephemeral recommendations to Persistent recommendations. Ephemeral recommendations are made based entirely on a single customer session, and are not based on any information from previous sessions of this customer. Persistent recommendations are based on the site recognizing the customer, and suggesting products to the customer based on the customer’s likes and dislikes in previous sessions.

This section is structured at the high level around the four recommendation techniques: non-personalized, attribute based, item-to-item correlations, and people-to-people correlations. For each technique, we briefly introduce it, explain its place in the taxonomy, and give examples of it from our recommender system examples.

4.1 Non-Personalized Recommendations

Non-personalized recommender systems recommend products to customers based on what other customers have said about the products on average. The recommendations are independent of the customer, so each customer gets the same recommendations. Non-personalized recommender systems are Automatic, because they require little customer effort to generate the recommendation, and are Ephemeral, because the system does not recognize the customer from one session to the next since the recommendations are not based on the customer. Non-personalized recommender systems are common in physical stores, since they can be set up on a display that is viewed without change by every customer.

For instance, the average customer ratings displayed by Amazon.com and Moviefinder.com are non-personalized recommendations. These recommendations are completely independent of the particular customer targeted by the recommender system. eBay has a slightly different form of non-personalized recommendation in its feedback profile. Customers give feedback on each other, rather than on products! The average and individual feedback is then available for consideration by buyers to decide whether a particular seller is a good risk, and by sellers to decide whether a particular buyer is a good risk. All three of these systems are nearly completely on the Automatic and Ephemeral end of the axes. Another type of non-personalized recommendation is the text comments supported in Amazon’s Customer Comments and eBay’s Feedback Profile. Both of these are still Ephemeral, but move closer to the Manual end of the other axis. In a sense, the system is merely providing raw data to the user, who must then collate the data and make sense out of it manually.

4.2 Attribute-Based Recommendations

Attribute based recommender systems recommend products to customers based on syntactic properties of the products. For instance, if the customer does a search for a historical romance book, and the E-commerce site responds with a list of three recommended books, that is an example of an attribute-based recommendation. Attribute-based recommendations are often Manual, since the customer must directly request the recommendation by entering his desired syntactic product properties. Attribute-based recommendations can be either Ephemeral or Personal, depending on whether the E-commerce site remembers the attribute preferences for a customer.

Reel.com's Movie Map is an example of a attribute-based recommendation. The recommendations are entirely based on the category of movie the customer selects. Since customers must explicitly go to Movie Map and navigate to a category to obtain a recommendation, Movie Map is Manual. Since Movie Map does not remember a customer's interest from one visit to the next, it is Ephemeral. Amazon.com Delivers is also Manual, since customers must explicitly sign up and provide a set of interest categories. However, Amazon.com Delivers is Persistent, since Amazon.com continues to send out recommendations in selected categories until the customer turns off the request.

4.3 Item-to-Item Correlation

Item-to-item correlation recommender systems recommend products to customers based on a small set of products the customers have expressed interest in. For instance, if a customer has placed a few products in her shopping basket, the recommender system may recommend complementary products to increase the order size. Item-to-item correlation recommender systems can be Automatic, if they are based on observations of the customer's unchanged behavior. They can also require some Manual effort, if the customer must explicitly type in several items of interest in order to generate a recommendation. Item-to-item correlation recommender systems are usually Ephemeral, since they do not need to know any history about the customer to generate a recommendation based on the products she has selected.

Reel.com's Movie Matches, Moviefinder's Match Maker, and Amazon.com's Customers who Bought are similar from the perspective of customer experience. All three suggest other products a customer might be interested in based on a single other product that customer has expressed interest in. These systems are Automatic and Ephemeral, since they require neither action from nor identification of the customer. CDNOW's Album Advisor is different, since it is triggered by the user asking for recommendations by typing in a set of artists. This application is still Ephemeral, but is closer to Manual because it requires some customer effort.

4.4 People-to-People Correlation

People-to-people correlation recommender systems recommend products to a customer based on the correlation between that customer and other customers who have purchased products from the E-commerce site. This technology is often called "collaborative filtering", because it originated as an information filtering technique that used group opinions to recommend

information items to individuals (Resnick et al. 1994, Hill et al. 1995, Shardanand & Maes 1995, Konstan et al. 1997). Though we use the word correlation in the name of this technique, hinting at nearest-neighbor techniques based on linear correlation, the technique can be implemented with many other technologies as well (Breese et al. 1998). Since we are focused on the effect of the technique on users, we differentiate according to user experience rather than implementation details. People-to-people correlation recommender systems are close to Automatic, since the recommendations themselves are generated automatically by the system. The system does not have to learn over time from customers. In some systems this is done by having customers explicitly rate products, in which case the system is moved part of the way towards Manual. In other systems, the learning is implicit from the buying patterns or click-stream behavior of the users, in which case the system is pure Automatic. These systems are most often Persistent, since learning about patterns of agreement between users requires substantial data which is most easily collected over time. In principle, such a system could be Ephemeral if user sessions are long enough.

Amazon.com's Book Matcher, Moviefinder's We Predict, and Style Finder are all examples of Persistent but not quite Automatic people-to-people correlation recommender systems. Users explicitly rate products and other products are recommended based on the ratings. Since the ratings are entered only to get the recommendations, these systems are not considered fully Automatic.

My CDNOW is a fully Automatic example of this technique, since customer opinions are inferred by the actions a customer takes in setting up his personal music site within the CDNOW E-commerce site. Recommendations are provided organically within the context of the personal music site.

4.5 User Inputs

Each of the previous four recommendation technologies requires some form of input upon which to base the recommendations. Typically this input is provided by the customer(s). However, it is possible that the input may also be provided by the business as well. The systems in our examples utilize one or more of the following inputs.

Purchase data: Which products a customer has purchased. Systems such as Amazon.com's Customers who Bought and My CDNOW make recommendations based entirely patterns of "co-purchase" between multiple customers. In principle, this may be augmented with how many of each product the customer has purchased.

Likert: What a customer says he thinks of a product, typically on a 1-5 or 1-7 scale. The scale may be numeric or textual, but must be totally ordered. Systems such as eBay's Feedback Profile and Levi's Style Finder utilize Likert inputs.

Text: Written comments intended for other customers to read. Usually not interpreted by the computer system. Currently included in systems such as Amazon.com's Customer Comments.

Editor's choice: Selections within a category made by human editors, usually employed by the E-commerce site, though independent editors are possible in principle. Editor's choice is

important in both Reel.com's Movie Matches/Map and Moviefinder.com's Match Maker.

5. Finding Recommendations

Just as sites can utilize different methods for calculating or displaying recommendations, so can they utilize different methods for allowing customers to access the recommendations. Through our recommender system examples we have identified four different methods for finding recommendations each of which may provide access to more than one recommendation interface and/or technology. These four methods are ordered in the amount of customer effort required to find the recommendations.

Organic Navigation: Requiring the least amount of work to actually access recommendations is the *organic navigation* process. In applications such as Album Advisor, Movie Matches, and Feedback Profile, customers do nothing extra in order to receive recommendations. In each of these applications, recommendations appear as part of the item information page. These recommendations can consist of additional items to consider, average ratings, or a list of other customer comments. However, the underlying similarity is that through the course of normal navigation of the site, customers are provided with a recommendations.

Request Recommendation List: Requiring not much more work from the customer is the *request recommendation list* process. Customers using applications such as Book Matcher and Style Finder can access recommendations based on their previously recorded likes/dislikes. To do so, they simply have to request these recommendations from the system.

Selection Options: In the *selection options* process customers must truly interact with the system in order to receive recommendations. Typically, customers choose from a set of predefined criterion/options upon which to base their recommendations. For example, users of Amazon.com Delivers have a choice from nearly 50 pre-defined categories in which to receive periodic recommendations. Even more involved, users of Moviefinder.com's We Predict system can select from a finite list of title, format, length and genre options to define a search, as well as customizing options such as ranking method and display features.

Keyword/Freeform: Arguably, the *keyword/freeform* option requires the most interaction from the customer. In applications such as Eyes, customers provide a set of textual keywords upon which to retrieve future recommendations. A version of Album Advisor takes the freeform input of multiple artists upon which to make recommendation matches. The We Predict and Movie Map applications produce recommendations from the results of a query conducted using the keywords provided. While each uses the keywords in very different manners, each requires the user to know specifically what types of things they are interested in.

6. E-Commerce Opportunities

Many varieties of recommender systems are already in use. We have already explored multiple interfaces, technologies, and information needs for these types of systems. However, there remain many opportunities for the expansion of recommender

systems in E-commerce sites. These range from simple variations on existing systems, to entirely new types of systems.

As discussed above, many sites currently use purchase data as an implicit, positive rating. CDNOW has realized in My CDNOW that owning something cannot always be interpreted as a positive. Recall that CDNOW allows customers to later go back and indicate "own it but dislike it". However, few sites are attempting to extract *implicit negative* ratings from purchase data. One way to do this would be through the analysis of data on returned products. While customers may return an item for a variety of reasons, in general any return could be considered as a negative rating on the item in question. Another model of implicit negative rating can be derived from detail views. If the site presents a few products in low detail and the customer chooses to view some products at higher detail, but ignores others, a mild negative rating can be inferred for the unselected items. Many recommender system algorithms perform better with both negative and positive ratings, so the negative data can be valuable.

Another creative use of a recommender system would be to use it in reverse to explain to a user what type of thing a product is. For instance, a recommender system might be used to tell the user "this product you're looking at is similar to these other products that you have liked in the past". Recommender system algorithms that correlate items can be used in this way. For best results they should be modified to return items that the user has purchased in the past, rather than the usual set of items the user has not purchased in the past.

Current recommender systems only use a small subset of the available information about the customer in making their recommendations. Some systems use demographic information, some use purchase data information, some use explicit ratings, some use ownership data, but no system effectively uses all this data simultaneously for real-time recommendations. How should these diverse types of data be combined? Should individual recommender systems running on each type of data produce independent recommendations? Or can better recommendations be produced by using all of the available data simultaneously?

Recommender system algorithms that use many different types of data create the possibility for "subtle personalization", in which the site provides a completely organic personalized experience to the customer. The customer interacts with the site just as she would have before personalization. She does not need to take any explicit actions to inform the site of her interests or desires. The site subtly changes the interface in nearly invisible ways to create a more personal experience for her, without her even noticing that anything has changed! (Balabanovic & Shoham, 1997) (Basu, Hirsh, & Cohen, 1998) (Sarwar et al, 1998)

Recommender systems are currently used as virtual salespeople, rather than as marketing tools. The difference is that many recommender systems target each individual customer differently, making it difficult to produce the reports that marketing professionals are used to. These reports usually partition the population into a manageable number of segments. One way to bring these two worlds together would be to use the people to people correlations used by some recommender system algorithms to create segments for the reports. Open questions include "how can names be assigned to the automatically

generated segments?” and “are automatically generated segments more useful for managing marketing campaigns than traditional segments?”

Recommender systems can be made more useful as marketing systems in other ways, too. Current recommender systems are mainly “buy-side” systems. That is, they are designed to work on behalf of the customer in deciding what products they should purchase. However, modern marketing is designed not just to maximize utility to the customer, but to maximize value to the business at the same time. The recommender system could produce an indication of the price sensitivity of the customer for a given product, so the E-commerce site could offer each product at the price that maximizes the lifetime value of the customer to the site. For instance, one customer might be willing to purchase the product at a price that would earn the site ten cents of profit, while another customer might purchase the same product at a one dollar profit. There are challenging ethical issue in implementing systems like these that use information from studying the customer in determining how to get more money from the customer. One economic study suggests that sites may need to pay customers for their information (Avery et al. 1999).

In a related concept, “sell-side” recommender systems could allow businesses to decide which clients to make special offers towards. In traditional commerce a company could offer a coupon for a free pound of bananas with the purchase of a box of cereal and a gallon of milk in order to increase sales of milk. The success of this depends on the customer viewing the coupon and remembering to bring it to the store. A recommender system could be designed which notices that a customer already has bananas and milk in his shopping cart and rarely purchases cereal. This customer might be a good choice for the above offer.

One limitation to recommender systems is collecting enough data to make effective recommendations for new users. One way to speed the transition is for sites to share information about their users. Shared information benefits users, because they get more accurate recommendations in less time, but decreases the benefit to individual sites because users are not as loyal to them. Since sites own the information they collect, they have little incentive to share with competitors. However, it seems quite possible that consortia of non-competing sites may form with the goal of sharing data to increase the value to companies within the consortia. Customers of these consortia will need assurances that their privacy will be carefully protected, even as their data are shared beyond the boundary of a single site.

7. Conclusion

Joe Pine’s book *Mass Customization* lists the five fundamental methods for achieving the goal of mass customization. Each of the first four of these goals can be realized through recommender systems:

- ◆ “Customize services around standardized products and services”: Recommender systems provide a customized service that enables E-commerce sites to sell their largely commodity products more efficiently.
- ◆ “Create customizable products and services”: Recommender systems are a customizable product of the E-commerce site.

- ◆ “Provide point of delivery customization”: The recommender system directly customizes the point of delivery for the E-commerce site.
- ◆ “Provide quick response throughout the value chain”: We predict that recommender systems will be used in the future to predict demand for products, enabling earlier communication back the supply chain.

Recommender systems are a key way to automate mass customization for E-commerce sites. They will become increasingly important in the future, as modern businesses are increasingly focused on the long-term value of customers to the business (Peppers & Rogers 1997). E-commerce sites will be working hard to maximize the value of the customer to their site, providing exactly the pricing and service they judge will create the most valuable relationship with the customer. Since customer retention will be very important to the sites, this relationship will often be to the benefit of the customer as well as the site – but not always. Important ethical challenges will arise in balancing the value of recommendations to the site and to the customer.

There are many different techniques for implementing recommender systems, and the different techniques can be used nearly independently of how the recommender system is intended to increase revenues for the site. E-commerce sites can first choose a way of increasing revenue, then choose the degree of persistence and automation they desire, and finally choose an recommender system technique that fits that profile.

Technologists often assume that the holy grail of recommender systems is fully Automatic, completely Ephemeral recommendations. Our study does not bear this assumption out at all. Many E-commerce sites use Persistent systems that require Manual effort from the user. One reason for this preference is that systems that are more Persistent create a relationship with the customers. If creating the relationship requires some degree of Manual effort from the customers, they will prefer to return to the site in which they have invested the effort, increasing the degree of “stickiness” of the relationship between the site and its customers. On the other hand, purely Manual recommender systems are entirely portable: the customer can freely go to another site with the same Manual features to obtain the desired recommendations. The optimal technology for the customers may well be fully Automatic, completely Ephemeral recommendations, since these recommendations leave them free to visit any similar E-commerce site. However, most recommender systems are deployed by the people running E-commerce sites. The optimal technology for them will be Persistent, and is likely to be only partially Automatic, requiring some input from customers to increase “stickiness”, but rewarding the customers with valuable recommendations based on their input. Our prediction is that most recommender systems will be run by E-commerce sites, and will be Persistent and partially Automatic. A few recommender systems will be run by groups whose goal is to support customers, and their recommendations will tend to be Ephemeral and fully Automatic to minimize customer effort, or Ephemeral and fully Manual, to maximize customer control and portability (Schneiderman 1997).

Recommender systems are creating value for both E-commerce sites and their customers. We hope our taxonomy of the ways in which recommender systems make money for sites, the way they

can be implemented, and our analysis of future directions for recommender systems in E-commerce helps to stimulate the creativity that is needed to produce the recommender systems of the future.

8. ACKNOWLEDGMENTS

We gratefully acknowledge the support of the National Science Foundation under grants IIS 9613960, IIS 9734442, and DGE 9554517. Support was also provided by Net Perceptions Inc, a company that Konstan and Riedl co-founded, and that sells a recommender system for E-Commerce.

9. REFERENCES

Christopher Avery, Paul Resnick, and Richard Zeckhauser 1999. The Market for Evaluations. *American Economic Review* 89(3): pp 564-584.

Marko Balabanovic and Yoav Shoham 1997. Fab: Content-based, collaborative recommendation. *Communications of the ACM*, 40(3): pp. 66-72.

Chumki Basu, Haym Hirsh, and William Cohen 1998. Recommendation as classification: using social and content-based information in recommendation. In *Proceedings of the 1998 Workshop on Recommender Systems*, pages 11-15.

John S. Breese, David Heckerman, and Carl Kadie 1998. Empirical analysis of predictive algorithms for collaborative filtering. In *Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence (UAI-98)*, pp 43-52.

Will Hill, Larry Stead, Mark Rosenstein, and George Furnas 1995. Recommending and evaluating choices in a virtual community of use. In *Proceedings of ACM CHI'95 Conference on Human Factors in Computing Systems*, pages 194-201.

Joseph A. Konstan, Bradley N. Miller, David Maltz, Jonathan L. Herlocker, Lee R. Gordon, and John Riedl 1997. GroupLens: Applying collaborative filtering to

Usenet news. *Communications of the ACM*, 40(3): pp 77-87.

Don Peppers and Martha Rogers 1997. The One to One Future : Building Relationships One Customer at a Time. Bantam Doubleday Dell Publishing.

B. Joseph Pine II 1993. Mass Customization. Harvard Business School Press. Boston, Massachusetts

B. Joseph Pine II, Don Peppers, and Martha Rogers 1995. Do you want to keep your customers forever? *Harvard Business School Review*, 1995(2): pp. 103-114.

Frederick F. Reichheld and W. Earl Sasser, Jr 1990. Zero Defections: Quality Comes to Services. *Harvard Business School Review*, 1990(5): pp. 105-111.

Frederick F. Reichheld 1993. Loyalty-Based Management. *Harvard Business School Review*, 1993(2): pp. 64-73.

Paul Resnick, Neophytos Iacovou, Mitesh Suchak, Peter Bergstrom, and John Riedl 1994. GroupLens: An open architecture for collaborative filtering of netnews. In *Proceedings of ACM CSCW'94 Conference on Computer-Supported Cooperative Work*, pp 175-186.

Badrul M. Sarwar, Joseph A. Konstan, Al Borchers, Jon Herlocker, Brad Miller, and John Riedl 1998. Using filtering agents to improve prediction quality in the groupLens research collaborative filtering system. In *Proceedings of 1998 Conference on Computer Supported Collaborative Work*.

Ben Shneiderman 1997. Direct Manipulation for Comprehensible, Predictable, and Controllable User Interfaces. *Proceedings of IUI97, 1997 International Conference on Intelligent User Interfaces*, Orlando, FL, January 6-9, 1997, 33-39.

Upendra Shardanand and Patti Maes 1995. Social information filtering: Algorithms for automating "word of mouth". In *Proceedings of ACM CHI'95 Conference on Human Factors in Computing Systems*, pages 210-217.