Understanding Merchandising Effectiveness of Online Stores

JUHNYOUNG LEE, MARK PODLASECK, EDITH SCHONBERG, ROBERT HOCH AND STEPHEN GOMORY

A b s t r a c t

While techniques and tools for Web marketing are being actively developed, there is much less available for Web merchandising. This paper contributes to the understanding and analysis of merchandising effectiveness in online stores. First, we categorize the requirements of business analysis for online stores. Especially, we distinguish the analysis requirements of Web merchandising from those of Web marketing. Then we focus on Web merchandising by identifying its elements, such as merchandising cues, shopping metaphors, store design and layout, and product assortment. Second, we tackle the problem of measuring the effectiveness of these elements of Web merchandising. For tracking and analyzing the merchandising performance, we introduce a new set of metrics, which we call *micro-conversion rates*. We explain how these metrics provide detailed insight into the performance of different business strategies. Finally, we describe a set of system and data requirements for tracking micro-conversions in online stores. Based on the identified system requirements, we present an architecture for a data-mart system for electronic commerce analysis that incorporates the new metrics.

Keywords: e-commerce, marketing, merchandising, micro-conversion, business intelligence

A u t h o r s

Juhnyoung Lee (jyl@us.ibm.com) is a research staff member at IBM T.J. Watson Research Center. His current research focuses on understanding and solving problems in marketing and merchandising of online stores in the Web by using data management, warehousing, mining, and visualization technologies.

Mark Podlaseck (podlasec@us.ibm.com) works at IBM T.J. Watson Research Center in Hawthorne, New York. He is currently designing a Web-based catalog for the works of composer, Philip Glass.

Edith Schonberg (ediths@us.ibm.com) is manager of the E-Commerce Intelligence group at IBM T.J. Watson Research Center. Her current interests are in understanding Web site design effectiveness, mechanisms for measuring visitor behavior, and personalization.

Robert Hoch (rhoch@us.ibm.com) is a Software Developer at IBM Research.

Stephen Gomory (gomory@us.ibm.com) is a programmer at IBM T.J. Watson Research Center where he works on software for analyzing e-commerce Web sites.

INTRODUCTION

In just a few years, Web sites have evolved from electronic brochures into channels for sales, customer service and information gathering. In order to maximize their return on investment, Web merchants are finding it necessary to thoroughly understand the effectiveness of their sites and to take appropriate actions when and where the sites fall short (Greening 1999; Hoffman and Novak 1996; Mulvenna *et al.* 1998). Web merchants generally analyze their sites' effectiveness from two perspectives: marketing and merchandising.

Marketing on the Web is narrowly defined as the group of activities used to attract customers to an online store and retain them. Techniques for online marketing include the use of banner ads, email ads, and personalization. Examples of marketing-related business questions are: Which banner ads generate the most traffic and sales? Which portal sites are pulling in the most qualified traffic? Who are the buyers referred by a particular ad? Web usage metrics for answering these questions are the banner ad *click-through rate*, which is the percentage of viewers who click on a banner ad, and the *conversion rate*, which is the percentage of visitors who purchase from the store (Novak and Hoffman 1997). Recently, ad banner *return*

Volume 10 (1): 20–28. www.electronicmarkets.org

on investment (ROI) has become the significant metric for Web marketing. Marketers want to know not only the number of visitors who come to a site from a particular banner ad and purchase from the site, but also how much revenue and profit is generated by these visitors (Duvall 1999; Haar 1999; Schmitt 1999; Tower *et al.* 1998).

Merchandising consists of the activities involved in acquiring particular products and making them available in the quantities and at the places, times, and prices to enable a retailer to reach its goals (Berman and Evans 1998; Levy and Weitz 1998). Merchandising also involves how and where to present products, and which ones to highlight and promote. Online merchandisers are responsible for product assortment and product display, including promotions and cross-selling. In large online stores, merchandisers make adjustments to the Web site content weekly or even daily. To assist Web merchandisers with such tasks, data mining techniques that generate association rules to determine products suitable for cross-selling are available (Anand *et al.* 1997; Brijs *et al.* 1999). In general, however, while the needs of reporting and analysis for *marketing* is being addressed, useful *merchandising* metrics and tools lag behind. For example, Web page hit counts provide a broad indication of visitor interest, but keeping track of which products are shown on served Web pages can be difficult, in particular, when page content is highly dynamic and personalized. Furthermore, there is a need to know to what extent interest translates into sales.

This paper contributes to the understanding and analysis of Web merchandising effectiveness in online stores. From a general set of business analysis requirements for online stores, we distinguish Web merchandising from Web marketing. We then focus on the Web merchandising, and categorize and explain the areas of Web merchandising such as merchandising cues, shopping metaphors, store design and layout, and product assortment. We define a new set of metrics for measuring merchandising performance, which we call *micro-conversion rates*. We show how these metrics provide detailed insight into the success of different Web merchandising strategies and tactics. Finally, we present an architecture for a data mart system for electronic commerce analysis that incorporates the new metrics.

The concept of micro-conversion rates is based on banner ad marketing metrics. More specifically, we view an online store as a collection of ads for individual products for sale in the store. From this perspective, we measure click-through rates, conversion rates, and ROI for a broad range of internal Web site features such as cross-sell, recommendation, and promotion. These metrics provide information at the individual product level, as well as on the product attribute and aggregate level. We believe that this type of information is both actionable and necessary for merchandising success.

ANALYSIS OF WEB MARKETING AND MERCHANDISING

In this section, we summarize and categorize areas of business analysis for the effectiveness of online stores. We identify three areas based on the types of analysis: the overall store performance, marketing, and merchandising. Each area addresses a different audience. The first area, the overall store performance analysis, summarizes the general statistics of store performance, and is directed primarily towards upper management. The second area of analysis, used by marketers, can be further divided into three subareas: advertising, external referrals, and customer segmentation. Marketers need to quantify the effectiveness of advertising activities such as banner ads and email ads, as well as how portal sites and other external sites are used by customers to reach the online store. Customer segmentation is also of paramount importance to marketers because it helps better understand their shoppers and their needs. Different sets of variables, such as demographic characteristics and shopping behavior, can be used for clustering customers and selected based on the types of business questions of interest. It is important to note that the Web has a few variables for clustering customers which do not exist in physical stores. They include the domain names from where the customers come and the external portal sites or banner ads which refer the customers. Finally, customer behavior regarding the shopping basket may be interesting and also used for clustering customers as well. Table 1 summarizes the Web marketing analysis with some sample business questions.

Table 2 summarizes the analysis of Web merchandising with some sample business questions. There are four areas of Web merchandising analysis: product assortment, merchandising cues, shopping metaphors, and Web design features. The first area, *product assortment*, deals with whether the products in an online store appeal to the visitors. If the product assortment is not optimal, the merchants may adjust, for example, brand, quality, selection, inventory or price of the products. Examples of business questions related to product assortment are given in Table 2.

Merchandising cues are techniques for presenting and/or grouping products to motivate purchase in online stores. Examples of merchandising cues include cross-sells, up-sells, promotions, and recommendations, and these merchandising cues are associated with hyperlinks on Web pages. For example, a *cross-sell* in an online store is a hyperlink which refers the visitor to a Web page marketing an item complementary in function to the item marketed on the current Web page. An *up-sell* link refers the visitor to a Web page presenting a similar but higher-priced item. A *recommendation* link highlights product pages that are likely to be of interest to the shopper based on knowledge of the shopper and the behavior of a large population. A *promotion* hyperlink refers a visitor to a product page from a 'What's Hot' page or a high traffic area such as the Home page for informing, persuading and/or reminding

Table 1. Areas of Web marketing analysis and sample business questions

Areas of analysis	*Business questions*
Overall store performance	What is the sales value for a specific period of time, say, week or month? What is the number of customer visits for the day? What is the store conversion rate for the week? What is the sales value index for the week?
Advertising	Which banner ads are pulling in the most traffic? How many sales are driven by each banner ad? What products do shoppers from a particular banner ad purchase? What is the conversion rate for each banner ad?
External referrals	Which portal sites are pulling in the most traffic? Which portal sites are generating the most sales? How many sales are generated by each referral site/search engine site? What products do shoppers from a particular portal site purchase?
Customer segmentation	How many visitors are from a specific domain? What is the distribution of first-time vs. repeat shoppers? What characterizes the shoppers of a particular set of products? What characterizes the shoppers who abandon shopping baskets?

Table 2. Areas of Web merchandising analysis and sample business questions

Areas of analysis	*Business questions*
Merchandising cues	How much do cross-sells/up-sells contribute to gross revenue? What are the best performing cross-sell pairs? Worst? What is the overall conversion rate for cross-sells/up-sells? How much do promotions contribute to gross revenue? Which promotions are generating the most sales? Which levels in site hierarchy are the best promotions located at?
Shopping metaphors	What generates the more sales value: search or browsing? How much does search contribute to gross revenue? What is the conversion rate for search?
Design features	What are the features of links customers most frequently click? What are the features of links customers most frequently buy from? What are the parts of page customers most frequently buy from? Do products sell better in the upper left corner?
Product assortment	What are the top sellers of the week? What is the conversion rate for a particular department? How is a product purchased: purchase frequency and quantity? What characterizes the products that end up being abandoned? How much of the sales of each product are driven by search?

the shoppers about a product and/or other aspects of the store. Web merchants need to understand the performance of different merchandising cues used in their stores in terms of traffic and sales driven by them.

Shopping metaphors in an online store are the means that shoppers use to find products of interest. Examples include browsing through the product catalog hierarchy, various forms of searching, and configuration for 'build-to-order' type products. The performance of different shopping metaphors in the store is a concern for online merchants. Like merchandising cues in online stores, shopping metaphors are associated with hyperlinks on Web pages. This allows one to categorize and group together hyperlinks in an online store by their types of merchandising cue and shopping metaphor.

Web design features presents another area of analysis for

Web merchandising. The design features of hyperlinks include media types (e.g., image or text), font (if text), size, color, and location. Examples of business questions related to Web design features are also given in Table 2.

Just as Web marketing uses banner ads and/or referral sites to attract customers from external sites to an online store, Web merchandising uses hyperlinks and image links within the store to lead customers to click to Web pages selling products. Web merchants employ a variety of tactics for merchandising by using hyperlinks. From this perspective, the problem of tracking and measuring the effectiveness of different merchandising tactics in an online store can be partitioned into three sub-problems:

1. classifying hyperlinks by their merchandising purposes;
2. tracking and measuring traffic on hyperlinks and analyzing their effectiveness (e.g., profit); and
3. attributing the profit of hyperlinks to their merchandising cue type, shopping metaphor type, and design features.

Analysis of the effectiveness of marketing tactics may be conducted in a similar way by using the metrics such as click-through rates and ad banner ROI described earlier. The only difference is that the originating hyperlinks in marketing efforts are presented and controlled in external sites. Classifying hyperlinks by their merchandising aspects and analyzing their performance provides an effective and powerful means to measure the performance of merchandising tactics in an online store.

MICRO-CONVERSION RATES

Having identified the areas of Web merchandising analysis, we now introduce a set of new metrics, referred to as *micro-conversion rates*, which can be used for measuring the effectiveness of efforts in these merchandising areas. The metrics are based on the conversion rate which is used for measuring online marketing performance. Traditionally, the conversion rate of an online store indicates the percentage of visitors who purchase from the store. While this measure is useful for evaluating the overall effectiveness of the store, it does not help to understand the factors within the store which may have affected the sales performance. The notion of micro-conversion rate extends the traditional measure by considering the following four general shopping steps in online stores:

1. *product impression*: the view of the hyperlink to a Web page presenting a product;
2. *click-through*: the click on the hyperlink and view of the Web page of the product;
3. *basket insertion*: the placement of the item in the shopping basket;
4. *purchase*: the purchase of the item – completion of a transaction.

Basic micro-conversion rates are computed for each adjacent pair of measures resulting in the first three rates in the following list. In addition, the aggregation of these three is also interesting. By looking at this look-to-buy rate, online merchants can tell if a product is overexposed or underexposed and take action to change the presentation of the product:

1. *look-to-click rate*: how many product impressions are converted to click-throughs.
2. *click-to-basket rate*: how many click-throughs are converted to basket placement.
3. *basket-to-buy rate*: how many basket placements are converted to purchases.
4. *look-to-buy rate*: what percentage of product impressions are eventually converted to purchases.

Note that the first of these, look-to-click rate, is similar to the click-through rate used for measuring the amount of traffic on banner ads. The micro-conversion rates go further by relating the traffic-related measure to sales which happen later in the shopping process. By precisely tracking the shopping steps with these metrics, it is possible to spot exactly where the store loses the most customers and to generate hypotheses as to what might have caused the loss.

The micro-conversion rates extend the traditional measure by considering the merchandising purposes associated with hyperlinks viewed in the first shopping step, i.e., product impression. In this way, the micro-conversion rate is related to tactics of merchandising, and can be used for evaluating the effectiveness of different merchandising aspects of online stores. Unlike the traditional conversion rate that describes the entire site with a single number, the micro-conversion rates can be computed for individual products, individual shopping metaphors, individual merchandising cue types, individual design features. For example, the micro-conversion rates can be computed for individual products to measure the product performance in the site. The rates for individual products can be rolled up to give the rates for categories of the products and then all the way up to the entire site, if an appropriate taxonomy of products is given. Furthermore, the micro-conversion rates can be computed for Web marketing analysis elements such as banner ads and different customer segments by associating these elements with the shopping steps. As a result, these rates provide a rich set of information pertaining to the analysis areas of Web merchandising and marketing described in the previous section. Table 3 presents sample business questions on merchandising addressed by micro-conversion rates for different analysis areas of Web merchandising.

Figure 1 illustrates sample micro-conversion rates for three different merchandising cues, i.e., promotions, recommendations, and cross-sells. Visitors are seeing twice as many impressions of promotions (40,000) than cross-sell impressions (20,000), and twice as many cross-sell impressions than those of recommendations (10,000). However, the look-to-click rate for recommendations

Table 3. Micro-conversion rates and sample business questions

Areas of analysis	*Micro-conversion rates*	*Business questions*
Merchandising cues	Look-to-click	What cross-sells are working best? Worst?
	Click-to-basket	Is a cross-sell more likely to be placed in a shopping basket if the first item has already been placed there?
	Basket-to-buy	Are the consumers who responded to a cross-sell any more or less likely to abandon a product in a shopping cart?
Shopping metaphors	Look-to-click	Are consumers finding what they want from the search engine?
	Click-to-basket	Do consumers who found a product through the search engine want the same amount of product detail as those who found it by browsing?
	Basket-to-buy	Are the consumers who responded to a search result any more or less likely to abandon a product in a shopping cart?
Design features	Look-to-click	Are visitors clicking more image links than text links?
	Click-to-basket	Are there product links that are misleading?
	Basket-to-buy	Where are the problems in the check-out process?
Product assortment	Look-to-click	Is a product's exposure optimized for the current level of consumer interest?
	Click-to-basket	Is the detailed information about a product appropriate?
	Basket-to-buy	What kinds of products are abandoned in shopping carts?

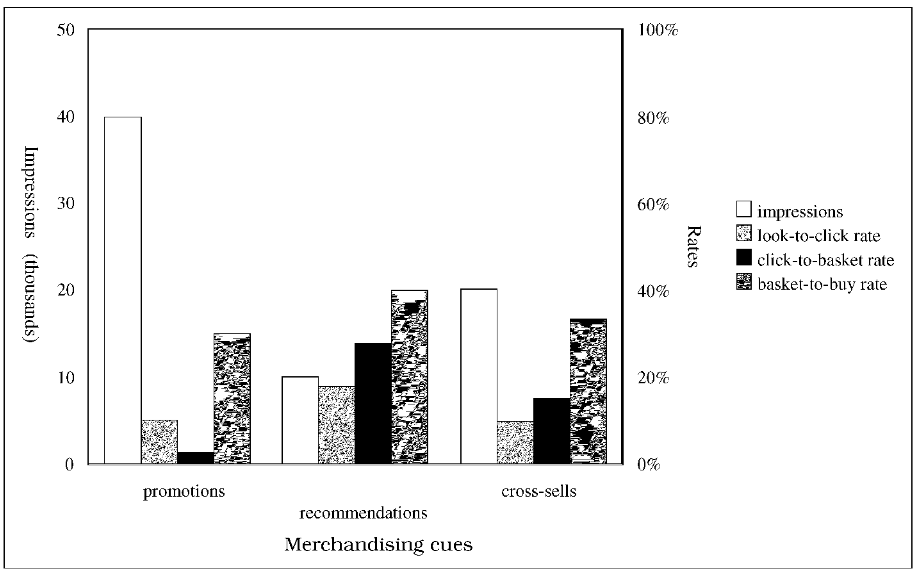

Figure 1. Examples of micro-conversion rates

(18%) is twice as high as those for promotions or cross-sells. Additionally, recommendations are resulting in a relatively high look-to-buy rate (2%). This means that the recommendation engine is relatively effective at personalization. On the other hand, this example shows that promotions in this store are not effective. Of the visitors who place a promoted item in a shopping basket, 30% of them purchase the product. However, the click-through rate for promotions is 10% and the click-to-basket rate is only 2.5%, so the look-to-buy rate is 0.075%, which shows poor overall performance, and an overexposure of the promoted items. Finally, the look-to-buy rate for cross-sells in this example is about 0.5%.

DATA REQUIREMENTS FOR WEB MERCHANDISING ANALYSIS

In this section, we briefly describe several data requirements for the analysis of Web merchandising explained in the previous section. While some source data are readily available from most online store sites, others are not and need to be collected with some special tools. Also, the collected data have to be integrated to show micro-conversions over shopping steps and to provide insight into the merchandising effectiveness of online stores.

First, the analysis of Web merchandising effectiveness based on micro-conversions requires the combination of the site traffic and sales data. In most online stores, the two types of data are typically stored in separate storage systems in different structures: the traffic data in Web server logs in a file format, and the sales data in the database of the associated commerce server. The commerce server database also contains information about customers and products (including product taxonomy) that may also be useful in computing a rich set of micro-conversions. It is important to combine data from the two different sources with a common key and to construct an integrated database system or a data mart system for business analysis.

Second, showing a complete set of micro-conversions requires product impression data. Capturing product impressions involves tracking the content of served Web pages, which is challenging because more and more Web pages are dynamically generated. (A simple example of a dynamically created Web page is a search result page commonly found in online stores.) Currently, the standard Web server logging mechanism does not capture the content of Web pages. One possible method is to enhance the Web server logging as a way to dynamically parse the content of served Web pages and extract useful data such as product impressions and information on hyperlink types. This ability of dynamically scanning Web pages as they are served is critical for tracking Web usage, because more and more Web pages are dynamically created from databases and contain personalized and adaptive content.

Finally, it is important to classify and identify hyperlinks by their merchandising purposes, to attribute the profit generated from the hyperlinks to their merchandising cue type, shopping metaphor type, and/or design features. For this purpose, Web pages and hyperlinks in an online store need to be tagged with semantic labels describing their merchandising features. Semantic labels of a hyperlink may include, for example, a product label, a cross-sell or promotion label, and a tag indicating where the product is being displayed. Such semantic labels for hyperlinks in a site may be explicitly provided in a form of meta-data during the site creation. If this is not the case, semantic labels need to be inferred from various sources such as the file name and/or path portion of URLs, types, values and orders of parameters in URLs of dynamic Web pages, and the location of a hyperlink in the page.

E-COMMERCE INTELLIGENCE FOR ANALYZING WEB MERCHANDISING

ECI ('E-Commerce Intelligence') is an ongoing project at IBM T.J. Watson Research Center. The architectural goal of the ECI project is to provide an analysis environment that is rich in data expressed in meaningful terms for analysts to answer the business questions described earlier. Data employed to achieve this goal come from Web servers, commerce servers, and the merchant's enterprise data systems. Combining these disparate data sources into a single analysis environment places e-commerce activities into their business context. In order to provide this environment, three distinct components are necessary. The first is an instrumentation of the Web server for recording detailed information about the user view of the Web site. The second is a set of descriptive meta-data that ties the Web site design to business terms and product information. The third is the data transformation necessary to tie this information together with additional merchant data in order to produce our multidimensional data model for analysis. Below we will briefly describe these components, and summarize the multidimensional data model for analysis.

Figure 2 illustrates the ECI system architecture. Sources of analysis data include various databases and files in the operational commerce site. Some data such as product categorization and cost of goods come from outside the operational system, and some such as advertising data and customer profile information come from outside the corporation. Target data is one that goes into the fields within the data mart system for e-commerce analysis. Figure 2 shows the types of source and target data, and the steps in transforming source to target data. During the data transformation process, data from various sources is cleansed, normalized, integrated and then loaded into a multidimensional data model for use by business analysts. Note that additional dimensions such as customer segmentation can be added or developed over time as patterns emerge from the accumulation of data.

Because existing Web server log files do not provide information on Web page content, we implemented a software module that runs in the commerce Web server environment, and parses the content of each Web page as it is served. The software module generates an enhanced Web server log that contains hyperlink information. In addition, it classifies the links by extracting relevant parameters such as product identifier or merchandising cue type. This information is combined with referral data and URL, and stored in the enhanced Web server log which we call *semantic log*.

The semantic log initially provides data about hits to the site. We use a software tool to process the log to identify requests from different visits and groups them together into sessions by using user identification data such as cookies. A session is a series of Web pages requested by a visitor in a single visit (Cooley *et al.* 1999; Pitkow 1996). This tool stores the created detailed information about the customer's interaction with the online store in a database for analysis.

The *site dictionary* is the data structure that stores information used to infer merchandising features of Web

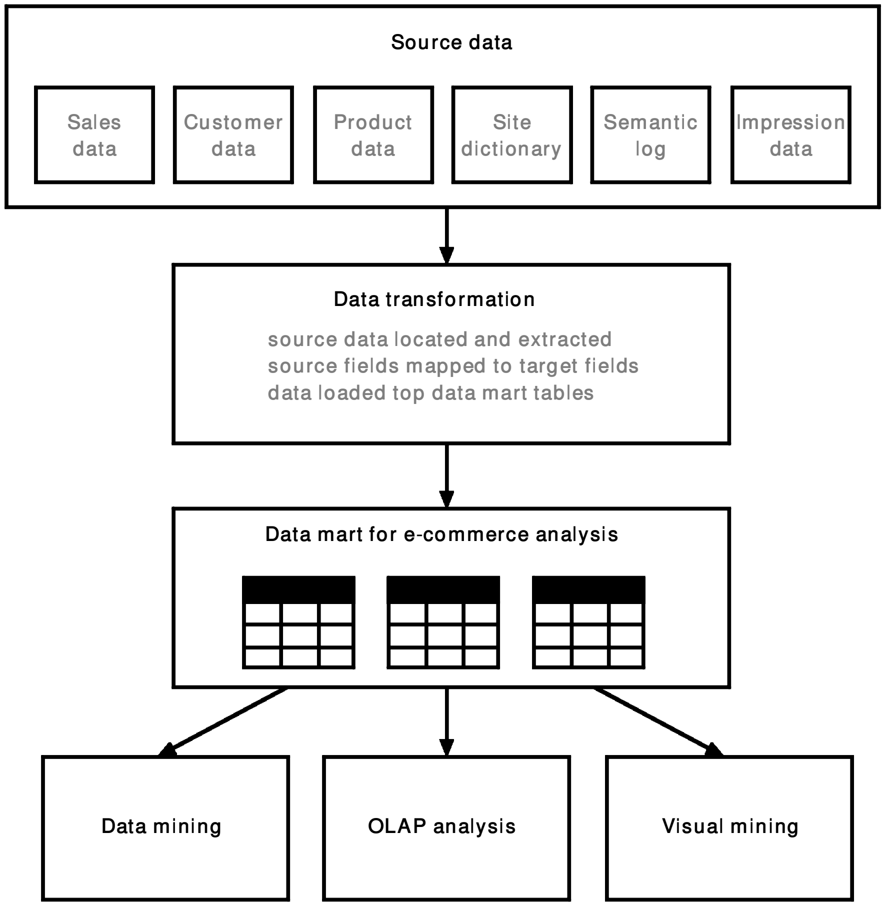

Figure 2. E-Commerce Intelligence system architecture

pages and hyperlinks when they are not given explicitly in a form of meta-data. The filename, path, and/or parameter portion of URLs, the location of the hyperlink in a page, and/or other attribute values of the hyperlink are used as keys to find the merchandising cue type, shopping metaphor type, and design features of a Web page or a hyperlink. The information in the site dictionary is used when the Web log is processed and when data from various sources is integrated. Note that the information stored in the site dictionary is site-specific: URLs, Web page types and layout, and site structure may vary from one site to another in many ways. However, most large-scale commerce sites rely on templates for page design and navigation. As a result, it is usually straightforward to classify URL patterns and parameters and map them to product displays, shopping metaphors, and merchandising cues. Since online stores will change frequently, it is important to automate this process of site meta-data collection and provide a versioning capability in order to maintain historical comparisons.

The end goal of the data preparation step is to produce a data mart consisting of data about Web traffic, product impressions and product purchases. In this environment, the business analysts can answer complex and useful questions such as those outlined earlier. The data mart is implemented as three related star schemas providing On-Line Analytical Processing (OLAP) capabilities. Figure 3 represents the three main areas of analysis and demonstrates the areas of overlap. Each oval contains the area name, its basic facts and the dimensions available. This provides a convenient way to diagram the common dimensions across these areas of analysis. The individual areas of analysis are given in italics in the figure and described in more detail below.

Traffic analysis is well defined, and forms the basis of existing analysis and reporting tools. Traffic analysis data comes directly out of server log information via a session construction tool. In this context, it is interesting because it provides the half of the data necessary to compute metrics such as conversion rate. In addition, it provides sufficient information about the origin of the visitor to be useful for measuring effectiveness of Web-based ad campaigns.

The order dimension has at its center the basic facts (item price and units ordered) about customer purchases through the Web site. In addition, through Web log analysis, using information provided by the semantic labeling, it is possible to associate a dimension describing the product impression associated with the sale (creative) and to identify the path that brought the user to the sale

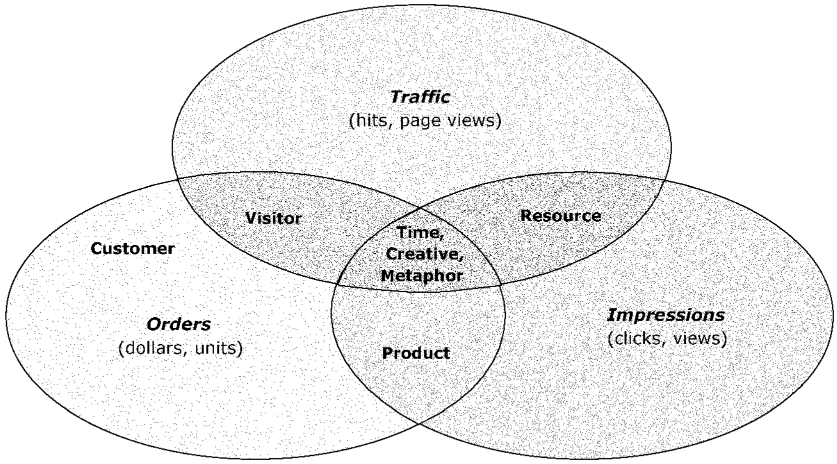

Figure 3. Multidimensional data model

(metaphor). The traffic analysis also contributes information about the Web-demographics of the customer, i.e., the sub-domain, platform, browser, referral-source for this session. Lastly, using enterprise data from the merchant, the product identifier can be tied to a product taxonomy as well as other product characteristics such as color, size, and brand.

The third area of analysis is that of product impressions, which are recorded in the semantic log as the frequency that a product was seen and/or clicked on by a visitor. The impression is characterized by its placement within the site, its role as a merchandising cue, and the type of visual presentation on a page. For example, the product view may occur on a promotion page as a featured element with a large graphic. Knowing the number of visitors that see, and consequently click on this graphic gives an indication of its effectiveness and is the first step in calculating the micro-conversion rate.

CONCLUDING REMARKS

Measuring the effectiveness of business efforts in online stores is relatively new. There has not been much work done on the subject, although it poses an important and imminent challenge for e-commerce across almost every industry. This paper has presented a collection of Web analysis requirements in the areas of marketing and merchandising in the form of business questions. It then outlined the process of tracking and analyzing the effectiveness of Web merchandising in online stores by using the notion of micro-conversion rates. We are currently implementing the mechanisms for tracking Web merchandising performance presented in this paper as part of the E-Commerce Intelligence project at IBM T.J. Watson Research Center. We also experiment with data extracted from an online store to demonstrate how different Web merchandising efforts are tracked and measured.

The work presented in this paper can be extended in several ways to address different questions on both business and system effectiveness of online stores. Some of the possibilities of extending this work for future study are as follows: First, the classification of hyperlinks and labeling them with semantic vocabularies can be generalized and applied to other types of business questions in other areas such as marketing and operation. One example is clustering customers by their shopping behavior with the types of hyperlinks they click on. In this case, hyperlinks in an online store need to be categorized and labeled to distinguish characteristics of shoppers' behavior. Second, the approach of OLAP-based exploratory analysis used in this work may be combined with a different, but complementary approach, i.e., data mining, that can find useful information on navigation paths, association, and clustering, which cannot be directly obtained by a pure OLAP approach (Büchner and Mulvenna 1998). Finally, the metrics for merchandising effectiveness presented in this work need to be adjusted and extended for the use in different shopping paradigms in the Internet such as online auction.

References

Anand, S., Hughes, J., Bell, D. and Patrick, A. (1997) 'Tackling the Cross-Sales Problem Using Data Mining,' *Proceedings of the 1st Pacific-Asia Conference on Knowledge Discovery and Data Mining*, 331–43.

Berman, B. and Evans, J.R. (1998) *Retail Management: A Strategic Approach*, 7th edn, New Jersey: Prentice-Hall.

Brijs, T., Swinnen, G., Vanhoof, K. and Wets, G. (1999) 'Using Association Rules for Product Assortment Decisions: A Case Study,' *Proceedings of the 5th ACM SIGMOD International Conference on Knowledge Discovery and Data Mining*, 254–60.

Büchner, A.G. and Mulvenna, M. (1998) 'Discovering Internet Marketing Intelligence through Online Analytical Web Usage Mining,' *SIGMOD Record* 27(4), 54–61.

Cooley, R., Mobasher, B. and Srivastava, J. (1999) 'Data

Preparation for Mining World Wide Web Browsing Patterns,' *Journal of Knowledge and Information Systems* 1(1).
Duvall, M. (1999) 'Exploiting Electronic Intelligence: Decoding the Customer,' *InteractiveWeek* 15 March.
Greening, D.R. (1999) 'Tracking Users: What Marketers Really Want to Know,' *Web Techniques* July.
Haar, S.V. (1999) 'Web Metrics: Go Figure,' *Business 2.0* June.
Hoffman, D.L. and Novak, T.P. (1996) 'Marketing in Hypermedia Computer-Mediated Environments: Conceptual Foundations,' *Journal of Marketing* July.
Levy, M. and Weitz, B. (1998) *Retailing Management*, 3rd edn, Irwin/McGraw-Hill Series in Marketing.
Mulvenna, M., Norwood, M. and Büchner, A.G. (1998) 'Data-Driven Marketing,' *International Journal of Electronic Markets* 8(3).
Novak, T.P. and Hoffman, D.L. (1997) 'New Metrics for New Media: Toward the Development of Web Measurement Standards,' *World Wide Web Journal* 2(1), Winter, 213–46.
Pitkow, J. (1996) 'In Search of Reliable Usage Data on the WWW,' *Technical Report*, College of Computing, Graphics, Visualization, and Usability Center, Georgia Institute of Technology.
Schmitt, E. (1999) 'Measuring Web Success' *The Forrester Report*, November.
Tower, D.L., McCarthy, J.C., Armstrong, A. and MacKenzie, M. (1998) 'Measuring Internet Return,' *The Forrester Report* 2(11).
Underhill, P. (1999), *Why We Buy: The Science of Shopping*, Simon & Schuster.