

Prospects for Personalization on the Internet

Alan L. Montgomery and Michael D. Smith

October 2007

Alan L. Montgomery (alan.montgomery@cmu.edu) is an Associate Professor of Marketing at the Tepper School of Business and Michael D. Smith (mds@cmu.edu) is an Associate Professor of Information Technology at the Heinz School of Public Policy, Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, PA, 15213.

Copyright © 2007 by Alan L. Montgomery and Michael D. Smith
All rights reserved

Prospects for Personalization on the Internet

Abstract:

Personalization is a key component of an interactive marketing strategy. Its purpose is to adapt a standardized product or service to an individual customer's needs. The goal is to create profit for the producer and increased value for the consumer. Applications of personalization have advanced greatly in conjunction with the Internet, since it provides an environment that is information rich and well suited to interactivity. This article reviews past research on personalization and considers some examples of personalization in practice. We discuss what we believe are key problems and directions for personalization in the future.

Keywords: Interactive Marketing, Personalization

1. Introduction

Personalization is intimately connected with the idea of interactive marketing and refers to the customization of some or all the elements of the marketing mix to an individual level. Personalization is more refined than customization, in that the customization is automated by the marketer on the behalf of the customer, as opposed to customization that a customer requests on their own behalf. For example, a customer ordering specific options for a micro-computer from Dell would be an example of customization, while Dell preloading a software bundle that it expects that the user would like to purchase is an example of personalization. This distinction is quite important because it puts the burden on the marketer to determine the appropriate degree of customization. As a consequence the marketer must anticipate what the customer wants. Generally this is done by leveraging customer-level information using analytical tools. Hence, personalization is paired closely to the technology and applications for which it is employed.

The promise of personalization has always been that it can deliver on the premise of a customer-oriented marketing strategy. Blattberg and Deighton (1991) prominently mention the notion of personalization in developing their notion of interactive marketing. They held out that promise that interactive marketing would allow the customer to be provided more pertinent information with less effort. This embodies the notion of what we refer to as personalization. Personalization is meant to eliminate tedious tasks for the customer and allow the marketer to better identify the user's needs and habits from past behavior.

It is important to understand that personalization predates the Internet and as a consequence is clearly broader than this single technology. However, it is clear that the Internet

has cultivated much of the interest in personalization and has continued to advance it more than any other technology. Therefore in this paper we limit our focus of personalization to the Internet and present our perspectives on current applications of personalization and what the future holds for personalization. Our goal is not to present a systematic survey of personalization, but to identify examples of excellence which in our opinion represent the state of the art for the practice of personalization on the Internet. A secondary goal is to consider future challenges for personalization.

2. Survey of Research on Personalization

The interest in personalization has spawned a wide variety of research. In this section we provide a survey of research into personalization. To better organize the work that has been done we classify this work into several areas: definitions of personalization, conceptual framework for personalization models, methodology for implementing personalization, and the effectiveness of personalization. We leave a discussion of privacy and its related research to section 4 when we discuss future challenges for personalization.

Definitions: Personalization is often invoked as an important element in an interactive marketing strategy. There are a plethora of versions of interactive marketing (Blattberg and Deighton 1991) including: customerization (Wind and Rangaswamy 2001), customer relationship management (Imhoff et al. 2001), one-to-one marketing (Peppers and Rogers 1997, Allen et al. 2001), permission marketing (Godin 1999), customer intimacy (Wiersema 1998), real-time marketing (McKenna 1997), McKinsey's continuous relationship management, Gartner Group's technology-enabled marketing, enterprise relationship management, internet marketing (Hanson 2000, Roberts 2003), database marketing, and e-marketing.

All of these forms of marketing refer to the concept of using technology to help identify and satisfy the customer’s needs—which is at the core of any marketing strategy. However, with each of these invocations of interactive marketing comes differences in what personalization means and how it should be used. In table 1 we provide a list of definitions for personalization compiled by Vankalo (2003). The common element in all these definitions is that personalization is an adaptation of the marketing mix to an individual customer based upon the marketer’s information about the customer.

<i>Source</i>	<i>Definition</i>
Peppers & Rogers (1999)	“Customizing some feature of a product or service so that the customer enjoys more convenience, lower cost, or some other benefit.”
Hanson (2000)	“A specialized form of product differentiation, in which a solution is tailored for a specific individual.” (p. 450)
Allen (2001)	Company-driven individualization of customer web experience. (p. 32-33).
Imhoff, Loftis & Geiger (2001)	“Personalization is the ability of a company to recognize and treat its customers as individuals through personal messaging, targeted banner ads, special offers on bills, or other personal transactions.” (p. 467)
Wind & Rangaswamy (2001)	Personalization can be initiated by the customer (e.g., customizing the look and contents of a web page) or by the firm (e.g., individualized offering, greeting customer by name, etc.). (p. 15)
Cöner (2003)	Personalization is performed by the company and is based on a match of categorized content to profiled users.
Roberts (2003)	“The process of preparing an individualized communication for a specific person based on stated or implied preferences.” (p. 462)

Table 1. Definitions of personalization. Entries listed chronologically. Adapted from Vesanen (2004) and Vankalo (2004).

Conceptual Framework for Personalization: Given the plethora of interactive marketing strategies it is not surprising that we lack a unified framework for personalization and customization; however two authors have recently proposed personalization frameworks that attempt to integrate the various viewpoints. Vesanen (2007) proposes that personalization is a broad concept that encompasses execution, marketing outputs in the form of product/service, promotion/communication, price and delivery, and the creation of value for both the customer and the marketer (see Figure 2 of Vesanen (2007)). Vesanen then decomposes each of these concepts further into a series of subtopics to develop a taxonomy which relates the common components of interactive marketing strategies.

A second framework has been proposed by Micelli, Ricotta, and Costabile (2007). They identify four dimensions along which to characterize personalization: value, knowledge, orientation, and relationship quality. Value refers to customer expectations for personalization features. Knowledge pertains to the customer's expertise. Orientation depends upon the consumer's behavior, whether the consumer is utilitarian or hedonic in their motivation. Relationship quality refers to tie that customer has with the firm and encompasses trust. The authors argue that these dimensions are useful for gauging the value of personalization for a marketer. Both of these frameworks are helpful in reconciling and integrating previous work on personalization, and point to our need to better understand both how, when, and what type of personalization is valued.

Methodology for Implementation of Personalization: Given that a key facet of personalization is adapting to a customer's interests, it is imperative to learn about the customer. Montgomery and Srinivasan (2003) define active and passive learning forms of learning which provides a useful dichotomy for understanding personalization methodology.

Active learning refers to posing direct questions to the customer, while passive learning requires the marketer to make inferences about the customer's interests generally using information from past transactions or clickstream data.

The most obvious example of active learning on the Internet is to pose questions directly to the user. However, customers do not always know the answer and they may not always be truthful. Therefore, a popular technique for trying to decompose a customer's preferences into the worth of each attribute is conjoint analysis. One important application of this technique is product design, which is quite relevant for personalization since we may want to employ a mass customization framework, but need to understand the value that customers attach to these customizations. Dahan and Hauser (2002) review six web-based methods for actively learning about customers: web-based conjoint design, user design, fast polyhedral adaptive conjoint estimation, virtual concept testing, information pump, and securities of trading of concepts. These approaches demonstrate that virtual design is practical in real-time environment.

Passive learning is more complex because it is necessary to make inferences from actions that may be indirectly related to the personalization that needs to be performed. For example, the marketer may wish to infer about price sensitivity from the path taken through a website. Fortunately, there are several major sources of information on the Internet that are available from users from which to passively learn: transaction data, clickstream data, and email.

There is an extensive literature on the use of transaction data that includes information about product, price, and quantity that is identified at a consumer level. The information can be used to make inferences about price sensitivity, and the best predictor of price response is past purchase behavior and not demographics. Rossi, McCulloch, and Allenby (1996) consider how purchase histories can be used to determine personalized couponing policies. They estimate

a statistical model known as a multinomial probit model to determine the price sensitivity of a supermarket shopper using information from previous purchases in a single product category, canned tuna. They used the model to determine what value (if any) of coupon to offer a shopper at the checkout stand and the expected profits for different coupons based on various information sets (Table 2). The blanket coupon drop assumed that all consumers received a coupon for 10¢ off on their next purchase. This blanket coupon drop is used as a benchmark to compare other strategies. Using demographic information to choose who received a coupon, they increased expected profits by 10 percent over a blanket coupon drop. In comparison, by using only data on the last purchase, they increased profits by 60 percent. Clearly purchase information is more valuable than demographic information in this context. In the full information case, they used information on both previous purchases and demographics and obtained a 160 percent increase in expected profitability over the blanket coupon drop. Purchase information used properly is an incredible resource. The value of transaction information for personalization is apparent.

Information Set	Profit gain
Blanket coupon drop	1.0
Coupon drop using information from consumer demographics	1.1
Coupon drop using information from last purchase	1.6
Coupon drop using entire purchase history	1.9
Coupon drop using entire purchase history and demographics	2.6

Table 2. The relative value of information sets for the distribution of a coupon as estimated by Rossi, MuCulloch, and Allenby (1996). Expected profits increase as more information is used. The most dramatic profit increased come from using purchase histories and not just demographic information.

Just as we can collect purchase data in physical stores, we can collect purchase data in virtual stores. However, the data set is much richer. In the brick and mortar retail store, we know what the customer purchased and what items and prices were available within the store. But imagine that we also record shopper movements through the store, what items they looked at and considered and those they ignored, how long shoppers considered their decisions, and whether they bought one product and then put a complementary product into their carts. This is exactly the kind of information we can collect in online shopping environments. Generally, we refer to this data as clickstream data. More precisely clickstream data refers to the sequence of pages viewed at a website and the time at which each page was viewed.

Several models for analyzing clickstream data have been proposed in the marketing literature. Montgomery et al (2004) propose a dynamic multivariate probit model that uses path information to predict where a user is likely to go in the remainder of the session. They discuss an application in which the pages and offers could be adaptively designed to personalize the features based upon a user's implied interests and goals. Sismeiro and Bucklin (2004) develop a sequential probit model, in which movements through the website is decomposed into a series of tasks that must be performed before moving to the next step. For example, after a user completes a product configuration they must then input personal information. They find that users likely to purchase can be identified early in their session; suggesting that greater attention to these individuals may increase the profitability of the website. Moe and Fader (2004) develop a model that measures evolving visit behavior across sessions, and find that users that visit more frequently are more likely to purchase, but also that changes in visit frequency over time can also be valuable in predicting conversion.

An alternative approach is to use data mining techniques to extract usage patterns (Mobasher et al 2000, Mobasher et al 2001, Nasraoui et al 2003, Pierrakos et al 2003, Ting et al 2006). These models hint at the potential for clickstream data in personalization, however we need to also understand that clickstream is underutilized (Sen et al. 2006) and it is likely to take years before its potential is fully leveraged.

Another rich stream of information for learning about customers is email. Ansari and Mela (2003) propose customization (or in our terminology personalization) of email communications with the customer. Specifically, they consider a problem in which the marketer sends targeted messages to customers in order to solicit new business. They estimate the value of different types and ordering of messages to increase the likelihood of purchase. Their methodology could be employed to create an adaptive web design as well. More generally, one can think of email as a textual database. Great strides have been made over the past ten years in analyzing the textual data that is encoded in user emails.

Effectiveness of Personalization: Another stream of research aims at identifying when personalization is and is not effective. Tam and Ho (2006) develop a model of personalization that posits that effectiveness is moderated by content relevance and self reference. Additionally, Tam and Ho (2005) consider the application of the elaboration likelihood model from the consumer behavior literature and finds that a consumer's need for cognition may play a role in the effectiveness of web personalization. More critically Nunes and Kambil (2001) find that many consumers do not like personalization but prefer the control that a user-customizable interface provides. They suggest that users prefer to provide information in a gradual process, and it may be a violation of trust to jump to a highly personalized interface.

3. The State of the Art in Personalization on the Internet

In this section we present a series of areas that illustrate how personalization has been applied to the Internet in order to illustrate the state of the art in personalization. Again we remind the reader that these examples are not meant to be exhaustive but are selected to be illustrative of elements of personalization technologies that we see implemented currently.

Personalized Price and Promotions for Airlines – Prices are simple to customize at an individual level, however keeping a consistent marketing strategy is difficult since personalized pricing strategies can degenerate into a confusing, Byzantine system. Perhaps the airline industry is one of the best examples of how a fractured pricing system can lead to confusion amongst consumers. Two examples of personalization in this marketplace are Southwest's Ding and Farecast, but each takes much different approaches to the problem.

Southwest's Ding system works by providing customers with special offers that are only available through Ding. Typically consumers report 20-25% reductions in prices over standard fares. Ding works by asking customers to register at Southwest and then download a small applet which is constant communication with Southwest to receive special, up-to-the-minute offers. The customer specifies markets that they are interested in flying to, but in order to secure the fare they must act quickly using the Ding applet. This type of system allows Southwest to make unique, personalized offers based upon the customer's history and stated preferences.

An alternative approach is exhibited by Farecast, which serves as an intermediary with the airline for the consumer. What makes Farecast unique is that it has an extensive database to track previous prices and make predictions using data mining techniques about what future prices are likely to be available. Farecast then can make personalized recommendations about

whether to purchase at the current price or wait. Additionally, Farecast can sell guarantees for a specified charge to the consumer (e.g., \$9.95) so that if the price drops below what is forecasted then it will reimburse the difference. Farecompare will also look for fares that match a consumer's prespecified range. These are examples of services that offer personalization by filtering another supplier's products.

Personalized Search – Search engines work by matching keywords or phrases that describe what a user is searching for against an index of web pages. One difficulty is that two users could enter the same keyword, such as “travel”, and the search engine would yield the same set of matches. One user may be preparing for a recurring business trip, and another for a one-time vacation. Each user has different expertise related to the keyword. They may already be familiar many popular sites and be interested in identifying web pages that are new or different. The goal of personalization for search is to adapt the search results based upon the user's previous searches and potentially their knowledge base and goals. For example, in the previous example the key question is whether we identify that a user travels on business frequently to San Francisco while the other has not.

We speculate that it is possible that user profiling can make these types of inferences if a user's history is known (Montgomery 2001). Contextualization is the process through which a user's task is inferred from their behavior, again path analysis demonstrates that this is possible (Montgomery et al. 2004). The next problem is whether a search engine can use this information to filter or augment search requests. This is precisely what personalized search engines are meant to provide. Google, Yahoo, Ask Jeeves, and A9 all provide personalized search services. Minekey follows a more hybrid strategy in that it is an applet that exists on the

user's computer and makes recommendations of new web sites based upon a user's browsing history.

Personalized Recommenders – Collaborative filtering was one of the first examples of personalization technologies widely available on the web. Currently collaborative filtering systems have been employed by Amazon, Barnes & Noble, LibraryThing, and Storycode to recommend books; Blockbuster, Eachmovie, Hollywood Video, MovieLens, and Netflix to recommend movies; Audioscrobbler, CDNow, iLike, iTunes, Last.fm, MusicMatch, MSN Music, MyStrands, RealPlayer MusicStore, Rhapsody, and Napster to suggest music; TiVo for television shows; Findory for news items; and StumbleUpon for website recommendations. This is not an exhaustive list, but clearly it demonstrates the pervasive use of recommendation systems.

A recent movement is towards recommenders that operate independently of a retailer and allow users to update their own information. Music provides a nice illustration of the promise of recommender systems and this new approach. The traditional problem in music is that it is extremely expensive to promote new music, and given the small number of songs that terrestrial radio stations play the mass market is a difficult through which to introduce new music. Personalization would seem to represent a potential solution by dramatically reducing promotion costs. Pandora is a personalized music recommendation system. It works by having users select a song or artist that they enjoy listening to. Pandora then makes recommendations and can adapt as users specify approval, disapproval, or that the user is tired of the song and do not play for one month. The site works by using 400 different musical attributes defined and rated by the Music Genome Project (e.g., Accordion Playing, Disco Influences, Midwest Rap Roots, etc.). The key to the system is that songs are selected based upon their similarity to

those attributes a user enjoys, which stands in contrast to collaborative filtering which makes suggestions based upon your similarity to other user's preferences.

Some of the challenges identified by Wedel et al. (2007) for recommender systems are the large number of choice alternatives, missing data, scale usage heterogeneity, customer choices depend upon the recommendations which introduces endogeneity, and scalability. Perhaps the best assessment of the current state of the art in recommender systems is Netflix. Netflix has sponsored a Prize to the first collaborative filtering system that improves predictions by more than 10% in terms of RMSE over its own Cinematch system. The range of predictions by consumers on the Netflix system is 1 through 5. A 10% improvement in RMSE means that a new recommender system must reduce RMSE from .9525 to .8572. This means that Cinematch can generally predict within a point whether a user will like a movie, but clearly this leaves a good deal of room on a five-point scale to improve recommendations. The promise of recommender systems is clearly demonstrated by the widespread adoption, but clearly improved techniques must be developed if we are to truly realize their potential.

Personalization Strategies – Perhaps no company has more aggressively pursued personalization technology than Amazon. Amazon's CEO Jeff Bezos stated that "If we want to have 20 million customers, then we want to have 20 million 'stores'. ... Our mission is to be the earth's most customer-centric company" (Ferranti 2000). Amazon provides dozens of forms of personalization features, some of the more notable ones are Your Amazon, Today's Deals, Gifts & Wish Lists, Recommendations by Category, Your Browsing History, Your Lists, and Your Profile. Amazon has made a clear and aggressive foray into personalization, although it is neither the largest nor the most profitable e-retailer. Personalization is not a panacea, but clearly it has demonstrated its promise in the marketplace.

4. New directions for Personalization

Our goal to this point has been to identify the promise of personalization made by research and examine the implementation of personalization in online environments. Our next step is to consider new directions for research and practice if personalization is to truly deliver on its promise. In this section we consider three problems that we believe must be addressed as personalization advances into the future: privacy, adaptive web design, and computation.

4.1. Privacy

From the outset a primary criticism leveled at personalization is that it constitutes an invasion of the consumer's privacy. There is a tradeoff between information to implement personalization and the potential violation of privacy that comes with this information. Just as the definition is diverse, the definition of privacy is perhaps more problematic (Mascarenhas et al. 2003). Adding to this tension is a varied legal framework that sees privacy vary on a continuum as a contractual agreement to a basic human right (Smith 2001). The crucial issue in the future to address is how can we understand and balance these competing objectives.

To motivate our discussion consider the idea of a sufficiency in statistics. Sufficiency means that all the information the analyst needs to make inferences about a parameter is contained within a sufficient statistic. For example, if an analyst is interested in making an inference about the mean of a normal population with a known standard deviation, then the analyst would only need to know the mean of the sample. All the information from the sample about making inferences on the unknown population mean is encapsulated in this single

parameter. The practical importance of the sufficiency reduction is that it allows one to discard the raw observations. Unfortunately, for some models without closed form solutions the sufficient statistic is the vector of all information. For example, the sufficient information in a multinomial probit model using the method of Rossi, McCulloch, and Allenby (1996) is the vector of all information.

The question for personalization models is whether a sufficient reduction can be made and whether this sufficient statistic can be disguised to protect the anonymity of an individual. From a privacy standpoint sufficiency reductions are helpful because the analyst no longer needs the raw information but can work with a reduced, potentially more anonymized form of data without a loss of information. The added benefit of sufficiency reductions is that they reduce the information that must be stored and transferred. The privacy of the sufficient statistic for a personalization model must still be addressed.

One solution insuring privacy is finding a technological solution. Sackmann et al. (2006) suggest that a combination of controlled disclosure of data, agreement on data collection, transparent processing and usage, and enforcing policy-compliance. For example, the Platform for Privacy Preferences (P3P) is an XML specification that states what kind of data is stored, how the data is used, how long it will be kept, and how it can be accessed. This allows the user some input over what information will be shared. Although to disguise clickstream data it still may be necessary to make requests through an intermediary server or proxy. However, P3P presents one step towards enabling consumers to be strategic in how their information will be shared. Along with the strategic action is a need to understand when is it in their interest to share this information and at what cost.

Xu et al. (2007) consider the effects of enabling the user to have some control over the type of information shared. They considered a personalized web search in which the user was given control over how much information was revealed. They found that search results could be significantly improved by only sharing some higher-level user profile information. Smyth and Balfe (2006) have also suggested that similar gains could be made by using social networks to improve search. Although some information about what segment you belong to still must be shared with the search engine.

One solution that would better insure the privacy of the individual is to have the individual perform the computation necessary to make the personalizations. This would require smart agents to act on the users behalf to collect information and integrate with their history. However, this requires greater computational resources to be devoted by the consumer and a large feed of information from the producer. Currently, there are no successful commercial bots that reverse this flow of information that have been developed. But this does provide an alternative computational model.

4.2. Personalization Using Adaptive Web Design

The traditional model of web navigation is that an expert creates a site by anticipating what objective a user has. Usually this means that there is a hierarchical structure for the site which allows users to traverse this structure in order to find the information that interests them. Perkowski and Etzioni (2000) portray these websites as “fossils cast in HTML”. Weld et al. (2003) provides both a survey of progress made in adaptive interfaces at the University of Washington as well as a summary of the lessons learned. Specifically, they find that users often are unable to specify their goals and making inferences about these goals based upon their

actions are better. Additionally, they believe adaptive interfaces must balance the costs of user errors in adaptive environments along with the gains. Also, defaults and automatic responses can help improve adaptive schemes.

The problem of adaptive design is not unique to personalization. This problem occurs frequently within many online contexts, such as web design, auctions, advertising, search, recommender systems, user navigation support, and tutoring models. Research on adaptive choice in online environments is scattered across many different fields of human computer interaction, machine learning, marketing, and statistics. To continue to move personalization forward we must work on integrating these diverse research streams.

4.3. Computational Issues

Research into the use of data intensive marketing strategies has been intense over the past decade. For example, researchers have proposed methods for the analysis of clickstream data for adaptive web design (Montgomery et al. 2004), and analyzing consumer choice for shopbot design (Brynjolfson, Smith, and Montgomery 2007). Much of this progress has been due to Monte Carlo Markov Chain (MCMC) estimation procedures. From a research standpoint these methods have been highly successful, allowing the estimation of previously intractable models. Practically though these simulation techniques can be very time consuming: the estimation of a hierarchical choice model could take hours to complete (Allenby, Rossi, and McCulloch 1996). In practice, Amazon typically requires its system to respond in 2 seconds or less. If the current stream of Bayesian research in marketing is to have an important impact in practice computational methods that can be implemented in real-time are necessary.

A promising computational strategy that has been employed by Google, Amazon, Sun, and IBM (as well as a host of others) is to construct grid computing environments (Bryant 2007). This allows companies to harness the power of up to hundreds of thousands of low-cost personal computers by splitting up the computational tasks so that they can be handled in parallel. Grid computing provides a promising direction for being able to analyze the quantitative models employed in marketing problems. Unfortunately the simulation techniques like MCMC used currently to estimate marketing models are by their nature sequential. Brockwell and Kadane (2005), however, proposed a promising method of parallelizing these algorithms. If a general parallelization scheme could be implemented then this would help bridge the gap between academic models that are computationally expensive and the need for fast response in practice.

5. Conclusions

The promise of personalization is the ability to deliver more relevant products and services to consumers. The current marketplace for most consumer goods with hundreds of products may seem overwhelming, just consider the example of LCD TVs. But if this number is compared with the large degree of heterogeneity between individual preferences and the hundreds of millions of consumers, clearly there are many magnitudes of growth that can occur, potentially there could be more products than consumers in the marketplace. Many managers may cringe at these ideas since traditional product and brand management models are inadequate to support such a system, however interactive marketing strategies can be hugely successful in this environment. The reward of personalization within an interactive marketing strategy is increased profits for the firm and increased value for the consumer.

The technology to deliver personalization systems is available. Currently the World Wide Web (WWW) is the most popular tool for delivery of personalization systems. However, the WWW is a component of the Internet. In the coming years we are likely to see an increase in mobile devices, embedded computers, smart appliances, and RFID tags and scanners that are networked into the Internet. This ubiquitous computing environment will likely yield richer environments for personalization that are contextually aware (Kenny and Marshall 2000). The technological and strategic impediments to personalization are substantial. Although these impediments do not appear to provide a barrier to growth in the application of personalization anytime in the near future.

References

- Allen, C., D. Kania, and B. Yaeckel (2001). *Internet world guide to one-to-one web marketing*. John Wiley & Sons, Inc.: New York.
- Ansari, A. and C.F. Mela (2003), “E-Customization”, *Journal of Marketing Research*, 40 (2), 131-145.
- Blattberg, R. and J. Deighton (1991). “Interactive Marketing: Exploiting the age of addressability.” *Sloan Management Review*, 33 (1): 5-14.
- E. Brynjolfsson, M.D. Smith, and A. L. Montgomery (2007), “The Great Equalizer? An Empirical Analysis of Consumer Choice Behavior at an Internet Shopbot”, Working Paper, Carnegie Mellon University, Tepper School of Business.
- Brockwell, A. and J.B. Kadane (2005), “Identification of Regeneration Times in MCMC Simulation, with Application to Adaptive Schemes”, *Journal of Computational and Graphical Statistics*, 14 (2), 436-458.
- Cöner, A. (2003). “Personalization and Customization in Financial Portals”, *Journal of the American Academy of Business*, 2(2), 498-504.
- Dahan, E. and J.R. Hauser (2002), “The virtual customer”, *Journal of Product Innovation Management*, 19, 332-353.
- Ferranti, Marc (2000), “Amazon’s Bezos touts personalization”, *InfoWorld*, June 20, 2000.
- Godin, S. (1999), *Permission Marketing*, Simon & Schuster.
- Hanson, W. (2000). *Principals of Internet Marketing*. South-western College Publishing: Ohio.
- Imhoff, C., L. Loftis, and J. Geiger (2001). *Building the Customer-Centric Enterprise, Data Warehousing Techniques for Supporting Customer Relationship Management*. John Wiley & Sons, Inc: New York.
- McKenna, R. (1997). *Real Time: Preparing for the Age of the Never Satisfied Customer*. Harvard Press: Cambridge, Massachusetts.
- Mascarenhas, O.A.J., R. Kesavan, and M.D. Bernacchi (2003), “Co-managing online privacy: a call for joint ownership”, *Journal of Consumer Marketing*, 20 (7), 686-702.

- Micelli, G., F. Ricotta, and M. Costabile (2007). "Customizing Customization: A Conceptual Framework for Interactive Personalization", *Journal of Interactive Marketing*, 21 (2), 6-25.
- Mobasher, B., Cooley, R. and Srivastava, J. (2000) "Automatic Personalization based on web usage Mining" *Communications of the ACM*, Vol. 43, No.8, pp. 142-151
- Mobasher, B., Dai, H., Kuo, T. and Nakagawa, M. (2001) "Effective Personalization Based on Association Rule Discover from Web Usage Data" In *Proceedings of WIDM 2001*, Atlanta, GA, USA, pp. 9-15
- Moe, W. and P. Fader (2004), "Dynamic Conversion Behavior at E-Commerce Sites", *Management Science*, 50 (3), 326-335.
- Montgomery, A.L. (2001), "Applying Quantitative Marketing Techniques to the Internet", *Interfaces*, Vol. 30, No. 2 (March-April), 90-108.
- Montgomery, A.L. and K. Srinivasan (2003), "Learning About Customers Without Asking", N. Pal and A. Rangaswamy (eds.), *The Power of One - Leverage Value from Personalization Technologies*, eBRC Press, Penn State University, pp. 122-143.
- Montgomery, A.L., S. Li, K. Srinivasan, and J. Liechty (2004), "Modeling Online Browsing and Path Analysis Using Clickstream Data", *Marketing Science*, 23 (4), 579-595.
- Wedel, M., P. Naik, L. Bacon, A. Bodapati, W. Kamakura, J. Kreulen, P. Lenk, D. Madigan, and A. Montgomery (2007), "Challenges and Opportunities in High Dimensional Choice Data Analyses", Working Paper, University of Maryland.
- Nasraoui O., Petenes C. (2003), "Combining Web Usage Mining and Fuzzy Inference for Website Personalization", in *Proc. of WebKDD 2003 – KDD Workshop on Web mining as a Premise to Effective and Intelligent Web Applications*, Washington DC, August 2003, p. 37.
- Nunes, P.F. and A. Kambil, "Personalization? No thanks.", *Harvard Business Review*, 79 (4): 32-.
- Peppers, D. and M. Rogers (1997) *Enterprise One-to-One: Tools for Competing in the Interactive Age*. Currency Double Day: New York.
- Perkowitz, M. and O. Etzioni (2000), "Towards adaptive web sites: Conceptual framework and case study", *Artificial Intelligence*, 118 (1-2): 245-275.

- Pierrakos, D., G. Paliouras, C. Papatheodorou, and C.D. Spyropoulos (2003), "Web Usage Mining as a Tool for Personalization: A Survey", *User Modeling and User-Adapted Interaction*, 13 (4), 311-372.
- Roberts, M.L. (2003). *Internet Marketing: Integrating Online and Offline Strategies*. McGraw-Hill/Irwin.
- Rossi, P.E., R.E. McCulloch, G.M. Allenby (1996), "The value of purchase history data in target marketing", *Marketing Science*, 15(4), 321-340.
- Sackmann, S., J. Strüker, and R. Accorsi (2006), "Personalization in Privacy-Aware Highly Dynamic Systems", *Communications of the ACM*, 49 (9), 32-38.
- Sen, A., P.A. Dacin, and C. Pattichis (2006), "Current Trends in Web Data Analysis", *Communications of the ACM*, 49 (11), 85-90.
- Sismeiro, C. and R.E. Bucklin (2004), "Modeling Purchase Behavior at an E-Commerce Web Site: A Task Completion Approach", *Journal of Marketing Research*, 306-323.
- Smith, Jeff H. (2001), "Information Privacy and Marketing: What the U.S. Should (and Shouldn't) Learn from Europe", *California Management Review*, 43(2), 8-31.
- Smyth, B. and E. Balfe (2006), "Anonymous personalization in collaborative web search", *Information Retrieval*, 9: 165-190.
- Tam, K.Y. and S.Y. Ho (2006), "Understanding the Impact of Web Personalization of User Information Processing and Decision Outcomes", *MIS Quarterly*, 30 (4), 865-.
- Tam, K.Y. and S.Y. Ho (2005), "Web Personalization as a Persuasion Strategy: An Elaboration Likelihood Model Perspective", 16 (3), 271-293.
- Ting, I-Hsien, Chris Kimble, Daniel Kudenko (2006), "UBB Mining: Finding Unexpected Browsing Behaviour in Clickstream Data to improve a Web Site's Design",
- Vesnan, J. (2007). "What is personalization? A conceptual framework", *European Journal of Marketing*, 41 (5-6): 409-418.
- Vankalo, M. (2003). "Internet-Enabled Techniques for Personalizing the Marketing Program", unpublished Master thesis, Swedish School of Economics and Business Administration, Helsinki.

- Weld, D.S., C. Anderson, P. Domingos, O. Etzioni, K. Gajor, T. Lau, and S. Wolfman (2003), “Automatically personalizing user interfaces” in International Joint Conferences on Artificial Intelligence (IJCAI03), Acapulco, Mexico.
- Wiersema, F. (1998). *Customer Intimacy*, Knowledge Exchange.
- Wind, J. and A. Rangaswamy (2001). “Customerization: The Next Revolution in Mass Customization”, *Journal of Interactive Marketing*, 15 (1), 13-.
- Xu, Y., B. Zhang, Z. Chen, and K. Wang (2007), “Privacy-Enhancing Personalized Web Search”, *International World Wide Web Conference*, Proceedings of the 16th international conference on the World Wide Web, Bann, Alberta, Canada. Publisher ACM Press.