Personalizing the user experience on ibm.com

by C. M. Karat
C. Brodie
J. Karat
J. Vergo
S. R. Alpert

In this paper, we describe the results of an effort to first understand the value of personalizing a Web site, as perceived by the visitors to the site as well as by the stakeholder organization that owns it, and then to develop a strategy for introducing personalization to the ibm.com Web site. We started our investigation by conducting literature reviews, holding brainstorming sessions with colleagues around the world, and performing heuristic usability evaluations of several relevant Web sites. We adopted a User-Centered Design approach and conducted a number of usability studies applied to the subset of the ibm.com Web site that business customers use for all aspects of purchase, service, and support of computer equipment. These studies employed a number of low- and medium-fidelity prototypes that we developed for this purpose. Our proposal for personalizing ibm.com consists of a set of 12 personalization features, selected for the value they offer to customers and to the business.

We use the phrase "personalizing a Web site" to mean using personal information about an individual to tailor the experience of that individual on the site. The information we include varies widely, from basic information, such as age and income, to information that is just entering the realm of interactive applications, such as the individual's intention or emotional state. We refer to a personalization technique as *adaptive* if it has the ability to automatically adapt content, navigation support, and presentation style to the individual visiting the site. The site is said to be *adaptable* if it allows the user to specify how the site should look and respond.[1] A *personalization policy* is the policy of a Web site about the use of personal data about an individual to adapt that Web site for the individual. A personalization policy might include the degree of control given to visitors to the site over their personal data. When personalizing a Web site, a *personalization feature* is a method of using some personal information about an individual to tailor a Web site visitor's experience on the Web site. Examples of personalization features include collaborative filtering and adaptive navigation.[2–4]

We take a broad view of personalization in the context of e-commerce according to which there is an exchange between two parties, the *customer* and the *provider* of a product or a service, with the essential goal of providing increased value to both parties.[5,6] Much of the research work on personalization to date has focused on how various personalization features affect customers. We go beyond this approach by considering the value and costs of personalization to both businesses and customers. Specifically, the value of personalization for the customer can be seen as a function of two variables: the cost of divulging personal information and the perceived resulting benefits. Similarly, the value of personalization for the provider can be seen as a function of two variables:

0018-8670/03/$5.00 © 2003 IBM

the cost of gathering information and the perceived benefits. Whereas the provider's costs and benefits can generally be expressed in monetary units, the customer's value proposition is more complex and can involve factors such as security, privacy, trust, and the value of business relationships.

Previous work on personalization of interactive applications focused on specific techniques. These include click-stream analysis,[7] collaborative filtering,[2–4] and data mining of Web user logs.[8–11] More recent developments include pattern classification and the development of recommender systems,[12,13] combining historical profile data and online visitation patterns,[14] and online heuristic decision-making based on flowchart and rule-based constructs.[15] In general, these methods attempt to predict user interests or goals and automatically personalize and adapt the presentation of information.

Although the previous work mentioned above is substantial, it usually focuses on a specific technique, a rather narrow context. Our research differs in that it is not aimed at identifying a single "best" personalization technique or technology. We currently view such an approach as less attractive because, as available data suggest, (1) the value of a specific technique to a customer varies with the customer and over time, (2) the value of a technique to a business depends on the business objectives, and (3) the benefits to be obtained from a combination of techniques are likely to surpass the benefits of any one technique.

In this paper we describe an approach to personalizing an e-commerce Web site and then we validate this approach through a number of usability studies involving the subset of the ibm.com Web site that business customers use for all aspects of purchase, service, and support of computer equipment. The ibm.com site includes four million pages of content on 2200 sub-sites (a sub-site is a collection of pages within a larger site with a common style and a shared navigation mechanism). In consultation with business executives, we decided to limit the scope of our research on personalization of the ibm.com site to the content areas related to servers and personal computer information, purchase, service, and support. We coordinated our activities with related human-computer interaction (HCI) efforts in the IBM Personal Computer Division and IBM Systems Group. As a multidisciplinary team of five researchers we had eight months to complete the work, and we collaborated with several additional groups across IBM in order to accomplish our goals.

We followed a User-Centered Design approach.[16] We first developed a master list of 75 possible personalization features and a tentative strategy for personalization that consisted of a set of hypotheses to be validated. Next, we undertook a set of three successive usability studies whose outcome was a list of the 12 highest-value features and policies for personalization from a customer point of view, a personalization strategy for the site, and a business case for the personalization of the site based on customer cooperation. We also developed the Personalization Value Model, which illustrates the value of Web site personalization to customers and the business.

The rest of the paper is organized as follows. In the next section, "Approach and first phase results," we describe our initial steps: literature review, brainstorming, and heuristic evaluation of IBM and competitive sites. In addition, we list the personalization feature categories we developed, as well as a number of hypotheses for a personalization strategy. In the section "Methodology for user studies," we describe the methodology we used in the three successive user studies that followed, including the participants, the user tasks, and the procedures used. Then, in the section "Results," we describe the results of our studies, including details of the business case for personalization, the Personalization Value Model, and the personalization strategy and ancillary recommendations. In "Discussion and future work," we discuss our results and some ideas for further study. We summarize our results in the last section.

Approach and first phase results

In planning and executing the ibm.com personalization project, we followed a User-Centered Design approach.[16] Figure 1 illustrates the major steps of the project and the resulting deliverables. As the figure shows, we performed the work in two phases. In the first phase we reviewed the literature, held brainstorming sessions with colleagues within the company who were able and willing to contribute to the effort, performed heuristic evaluations of a number of Web sites of interest, and identified the business requirements. The results of this phase included a master list of possible personalization features and a set of hypotheses for a personalization strategy. The box labeled V11/V12 Application Code stands for two versions of the code supporting the ibm.com site (as illustrated in Figure 1, we have used infor-

Figure 1 ibm.com personalization project: major activities and deliverables

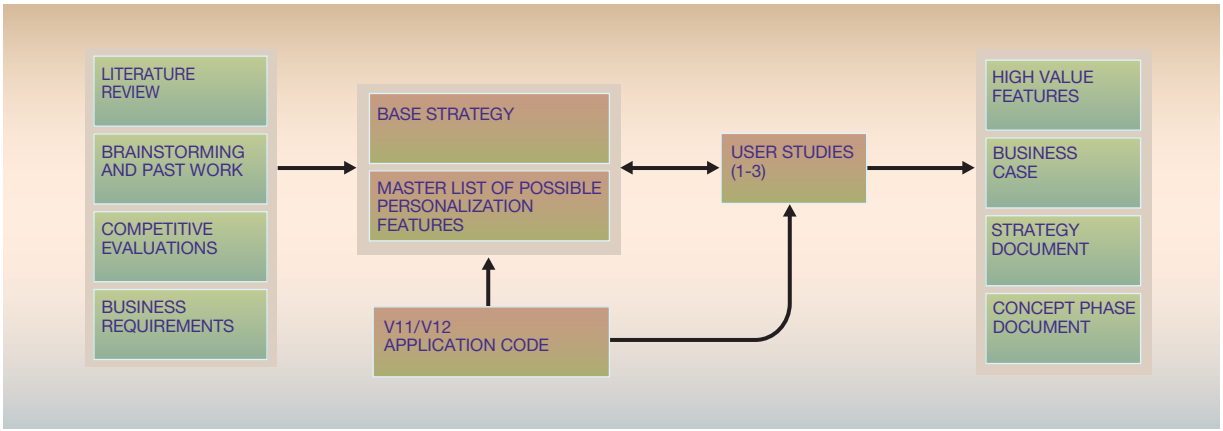

mation derived from this code as input to the usability studies). In the second phase of the effort we performed three successive usability studies that involved sessions with selected users and a number of interactive prototypes of personalized user interfaces for the site. These studies enabled us to define a personalization strategy, a list of high-value personalization features for ibm.com, a business case for personalization of ibm.com, and the Personalization Value Model. In the rest of this section we describe the first phase activities and results.

Literature review and brainstorming sessions. We conducted a review of published research in the areas of Web site personalization, e-commerce involving one-to-one marketing and permission marketing, and adaptive hypermedia. In the material surveyed, we included company confidential reports on personalization, e-commerce and HCI research papers, and market intelligence reports. In addition to the literature review, we conducted brainstorming sessions with IBM researchers around the world who are working in areas related to personalization. We started these sessions by stating our goals and requesting ideas for personalization features to be considered. In the discussions that followed we entertained a wide range of ideas, having made a conscious effort not to exclude any on the basis of questionable feasibility or possible implementation difficulties. As might be expected, the brainstorming sessions produced a variety of approaches, from those that were technically feasible to those that were futuristic (one such idea involves inferring the emotional state of the user and adapting the system responses accordingly).

Heuristic evaluation of Web sites. We performed a heuristic evaluation of ibm.com and a number of competitive sites (Dell Computer Corporation, Hewlett-Packard Company, Compaq Computer Corporation, Sun Microsystem, Inc., and Amazon.com, Inc.), six sites in all. Our goal was to gain an understanding of current best practices in the personalization of sites that market computer equipment and, in the process, expand our list of personalization features and determine the opportunities for IBM leadership in this area. We evaluated the sites using a set of six task scenarios covering the purchase and support of computer hardware and accessories. Each one of us was randomly assigned to perform heuristic reviews involving a subset of six user task scenarios on a subset of the six sites. The heuristic reviews of the six sites showed that, at the time, most of these sites had only limited personalization capabilities (amazon.com was the exception). The reviews produced a list of 18 design recommendations that were incorporated in the strategy as working hypotheses, and also in the master list of possible personalization features.

Identifying business requirements. For identifying the business requirements we employed an adaptation of the contextual inquiry method.[17] Contextual inquiry is an HCI method that enables practitioners to identify usability issues through observation of users in context, use of probing questions, and collection and analysis of key data points. Inductive reasoning is employed to identify issues through the "voice of the customer" (e.g., through actual statements by customers). Hierarchies are built from the bottom up based on data instances, to a larger view

of patterns and affinity diagrams of the associations that highlight common customer issues and requirements.[17] The ibm.com stakeholders (people from marketing, sales, development, finance, industry solutions, support, and hardware and software brands) were the primary user group to identify business requirements. To gain a deeper understanding of stakeholder goals, we adapted contextual inquiry methods by the combined use of probing questions with user observation, through analysis of key data points and the construction of affinity diagrams and the modeling of personalization features. We met with 12 representatives of various stakeholder groups. A typical meeting involved two members of our team and a stakeholder, sometimes accompanied by an associate. We asked questions about the site-related business goals for which they were responsible and then discussed the potential impact of personalization. We probed for specifics in order to keep the discussion grounded and relevant. We observed each user work environment for a maximum of 120 minutes. The resulting affinity diagrams document the stakeholders' business requirements regarding customer experience goals for the site, the quality of customer relationships, business financial goals, and infrastructure goals for the site. The analysis also identified the target population for usability studies, themes regarding the personalization pilot on the site, and identification of obstacles and limitations in achieving the identified goals. These results provided the business grounding for the empirical user studies.

Master list of possible personalization features. We gathered information from a variety of sources and initially cataloged 75 personalization features and policies. We view this list as evolving and changing over time. Although space limitation prevents us from providing the complete list here, we present in Table 1 a list of categories (or clusters) that we use for grouping personalization policies and features.

We hypothesized that Web site visitors would appreciate a centralized mechanism for accessing and modifying their personal information. To test this hypothesis we used the *Personal Book* mechanism, created by one of the authors (C. M. Karat).[18] The Personal Book, listed in the first entry in Table 1, is a personal space created when a visitor registers with the site. It is reachable from any page within the site and gives the visitor access to his or her profile, as well as easy and direct access through tabs on the Personal Book to all other personalization features (e.g., a list of previous purchases, compatible accessories for such purchases, and replacements for discontinued currently owned items). Figure 2 shows a Personal Book used in the Study 3 prototype. The Personal Book also provides the ability to filter products based on user needs and to store a "wish list" and get advance notice of special offerings involving items on the list.

Our personalization strategy is based on three personalization policies that, we hypothesized, are crucial to the success of our effort: (1) Web site visitors have control over their personal data; (2) visitors are asked to supply personal information only when necessary, and immediate value is provided in return (permission marketing); (3) visitors may select a level of identity appropriate for the task to be performed. We now discuss each of these in more detail.

Control over personal data. In the past many companies viewed the data collected about visitors to their Web site as data the company owned and that could be used in any way deemed useful to the company. In Europe, social and legal pressures have forced companies to concede ownership of data collected on their Web sites to the subject of the data.[19] In spite of similar pressures, the United States has been slower to adopt such laws, relying instead on businesses to self-regulate.[20,21] The policy we adopted entails that the customers own their personal data, which means customers can, at any point in time, view, edit, and delete information about their person, their purchases, and their actions on the Web site. Moreover, the use of such data by the e-commerce company requires permission, and such permission is given for specific uses only.

Permission marketing. Permission marketing is based on the idea that a customer's profile is built gradually over time as the individual develops trust in the e-commerce company.[5] The customer is asked to provide only the information needed to enable specific services and receives immediate value in return for the information disclosed. To collect data for personalization features, many Web sites require that the customer first register by filling out a lengthy questionnaire. Hagen finds that people often defeat the purpose of these forms by entering incorrect information.[22] As part of our effort we will determine whether permission marketing increases the customers' willingness to share data.

Levels of identity. The personal information available to a Web site about a visitor may vary over a wide range, from no information to extensive information.

Table 1 Personalization feature clusters

Personalization Feature	Description
Personal Book	An area on the Web site that is available from any page on the site. This area provides one place where a registered, logged-in user can access and modify all personal data about himself or herself and access all personalized functions on the site.
Universal profile	A profile that is active across the entire site.
Subscription-based services	All services in which the user must indicate an interest before they are active for that user. For example, a user may subscribe to a security alert system to receive alerts either in his or her Personal Book or through e-mail whenever a security patch becomes available.
Service and support	Features that assist users in obtaining support for the systems they own. An example is remote diagnostics.
Recommendations based on profile data	Features that include collaborative filtering technologies and recommender systems.
Adaptive presentation tailored to user characteristics	Features that adapt the contents of the Web page based on either the user's known characteristics (e.g. adapt information content to the user's job type) or context (e.g. because the current line speed is low, do not send unnecessary graphics).
Personal preferences in page layout or format (customization)	This includes features that allow users to specify aspects of the Web page that they would like.
Adaptive navigation	Features that adapt the Web page presentation and/or content based on user profile, other pages visited during this session, current emotional state, or current task or context.
Live chat-like or phone-based help or sales support (personal shopper)	Live help in the form of instant messaging, e-mail, or phone calls in which the company representative has access to both the user's profile and the pages visited during the session.
Feedback that the system recognizes a "repeat" visitor	An indication on the Web page that the system knows who is logged on and what their current role is (e.g. IT professional, personal, etc.)
Transaction history	Features that provide the Web site visitors easy access to a list of their prior transactions on the Web site, including purchases, support provided, etc.
Loyalty programs, incentives	Features that reward repeat visitors and/or customers by offering them special incentives (e.g. special prices or gifts).
Future purchase considerations	Features that make it easier for Web site visitors to make purchases in the future. Examples of these features include the creation of an individualized catalog of items a customer often buys, recommendations based on past purchases, or wish lists in which a customer can indicate an interest in items should they go on sale.
Your store, built by an expert	A customized Web page or catalog that is created by an industry expert based on a group profile.

Schaffer defines *levels of identity* as a discrete measure of the degree of personal information that the user decides to provide to a Web site through the particular role in which the user decides to enter the site. The higher the level of disclosure, the higher the trust the customer has in the e-commerce company.[23] Table 2 shows an example. At the lowest level of trust, the user has the cookies turned off, and the site has no information (visitor is *invisible*). At the highest level of trust the user has created multiple profiles for a number of different purposes (visitor has *differentiated roles*). A user might have, for example, a home profile and a number of work profiles.

Methodology for user studies

We performed a series of three successive studies, carried out both in laboratory and field settings (customer locations). The studies included group sessions (Studies 1 and 2) similar to design walk-throughs,[16] and individual sessions (Study 3). The group sessions involved from two to six participants. Two researchers conducted each user session—a facilitator ran

the session, and a colleague recorded the session on videotape and collected verbal comments from participants. Some user sessions were held in the Usability Lab at the IBM Watson Research Center, in Hawthorne, New York; others were held at customer locations in New York City; Raleigh, North Carolina; and Austin, Texas. The user sessions in the Usability Lab studio were observed from the control room through a one-way mirror and closed circuit television, and the proceedings were recorded by video cameras. In the field settings, the facilitator, the colleague, and the participant were located in a conference room equipped with a laptop computer for the participant's use and a video camera for recording the session.

Participants. All the participants were people comfortable with the Web (at least three hours of use per week) and were at least moderately "technology savvy" in their purchasing behavior. They made purchases on the Web themselves but may have enlisted expert help when selecting items to be purchased. They had been involved in the purchase of a server in the past year. About half of them had also been involved in the purchase of desktop and notebook systems. The participants for Study 1 were IBM employees who matched the profile characteristics outlined above. The participants for Studies 2 and 3 were a mixture of current ibm.com customers and target customers. The participants included IT decision makers and business unit executives for large traditional companies, medium and small business companies, and small dot com companies. The participants for Studies 2 and 3 were recruited by an outside agency using a participant recruitment profile based on the characteristics described above.

Twenty participants were recruited for Study 1, 23 participants for Study 2, and 22 participants for Study 3, for a total of 65 participants. Participants recruited by the agency received $150 for taking part in a two-hour user session. New groups of target users were recruited for each study. The three studies were conducted from August to December of 2001.

User tasks. We created a set of task scenarios that covered purchasing IT equipment, maintaining and upgrading the equipment, and obtaining support for products. The scenarios contained personalization features appropriate for the tasks and covered all the personalization feature clusters listed in Table 1. The task scenarios featured Pat User (whose gender was randomly assigned in each session), who was required to complete a set of tasks that might normally occur over an 18-month period. Thus we observed initial visits to a site, as well as return visits to the same site. Following Study 1, we made several small adjustments in the personalization features (added four features), the questionnaires, and the scenario and storyboard presentations for Study 2 (see Figure 3). Feedback from Studies 1 and 2 was incorporated into the design of Study 3.

In the group sessions in Studies 1 and 2, the scenarios were read aloud and at the same time a story-

Figure 2 The Personal Book

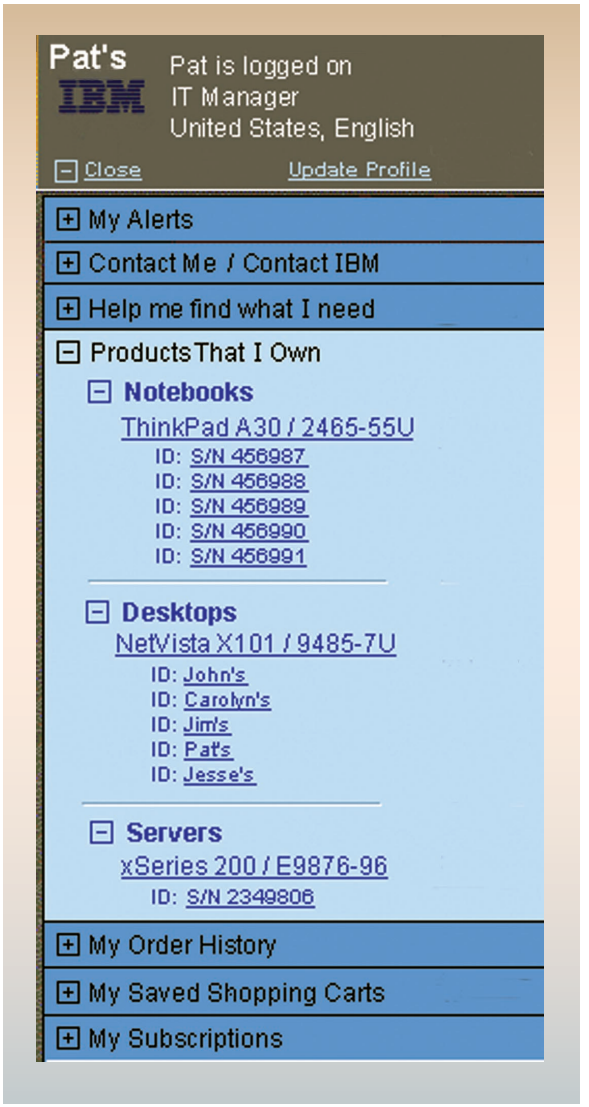

Figure 3 Development of the high-value personalization feature and policy list

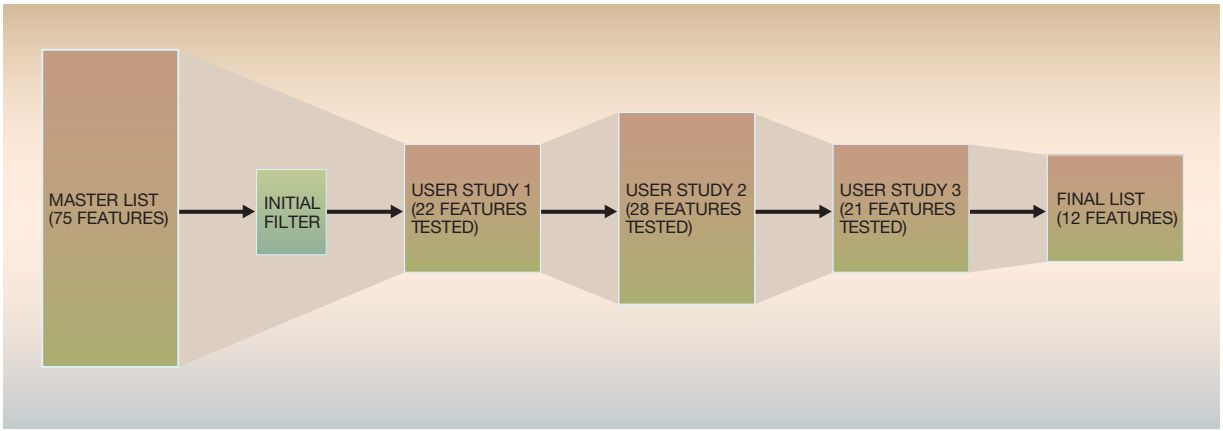

Table 2 Levels of identity

Level of Identity	Description
Invisible	An individual who not only has not registered with the site, but has cookies turned off so that the Web site cannot detect whether he or she has ever visited before. Lowest level of trust in the site.
Anonymous	An individual who has cookies enabled but has not registered on the site. Shows slightly more trust of the site.
Identified	An individual who has registered with the site, providing personal information in exchange for the use of personalization features. Shows a high degree of trust in the site.
Associated	An individual who has both registered with the site and indicated that he or she is associated with a particular team or organization. Shows a very high degree of trust that the site will provide value to him or her and the team.
Differentiated	An individual who has created multiple profiles on the Web site for different purposes (e.g. home and business, different business roles). Shows a very high degree of trust in the site.

board prototype was presented on a large screen. In Study 3, which was based on individual user sessions, the participants interacted with a medium-fidelity prototype and carried out six typical task scenarios.

A first-visit scenario reads as follows:

> You are Pat User and you have just become the manager of an IT department that develops and hosts Web applications for client companies. You know that you will need to travel in your new job, so you need to purchase a laptop computer. You want to spend less than $2000. As part of your shopping for this laptop, you look at ibm.com. You decide to buy an A Series ThinkPad*.

A return-visit scenario reads as follows:

> Three months have passed since you made your first purchases on the ibm.com site. You are still an IT manager for a small company that provides Web-hosting services for client companies. You are about to purchase a server in order to provide Web-hosting services for a new client. The client requires that their data be hosted on a separate server for security reasons. You have $8000 to spend on the server. The server is required to support 150 users at a time.

Procedure. Both group and individual sessions began with the participants filling out a questionnaire on demographic and job-related information. In the group walk-throughs (Study 1 and 2), the facilitator presented to participants three task scenarios accompanied by a storyboard prototype projected on a large

screen. The first scenario concerned buying a server and a mix of desktop and notebook systems for a department in which 10 employees were beginning a new project. The second scenario involved upgrading a server in order to handle the workload of 10 additional employees, and purchasing desktop and notebook systems for these employees. The third scenario focused on buying accessories—in this case zip drives—for Pat's entire department. Each scenario involved between 5 and 13 personalization features and policies (e.g., presentation of accessories constrained to those compatible with a selected, previously purchased machine; presentation of servers compatible with previously determined business characteristics when searching for servers; user control of data). Presenting a scenario took about 20 minutes and involved about 10 screen shots. Following the scenario presentation, the facilitator conducted a 5-minute discussion with the participants covering the features presented. Comments were recorded on flip charts mounted on the walls. Each participant then completed a post-scenario questionnaire that rated each personalization technique on a 7-point scale ranging from "highly valuable" to "not at all valuable," and provided additional comments in writing. Following the three scenario presentations, participants filled out a post-session questionnaire in which they were asked to identify the most and the least valuable personalization features (in relationship to their jobs). The entire session lasted about 2 hours.

In the individual sessions (Study 3), each participant was asked to read a description of six task scenarios involving Pat User's department, and then perform these tasks using interactive prototypes. The application prototypes were implemented using Microsoft PowerPoint** and ran on an IBM ThinkPad. The order of the presentation of tasks was varied among participants using a Latin squares design (a Latin square of order n is an n by n array of n symbols in which every symbol occurs exactly once in each row and column of the array). As an example, one scenario involved the task "purchase additional memory for the laptop computers you bought last month." Participants were encouraged to "think aloud" as they performed their tasks. Each scenario involved three or four personalization features and policies. After each scenario, the participants filled out a questionnaire asking them about their reactions to the features presented in the scenarios. Following a discussion period led by the facilitator, participants filled out a post-scenario questionnaire about the features in the scenario. The participants were asked for ratings and design comments on the personalization features. At the end of the session, participants completed a post-session questionnaire in which they were asked to rank-order their least and most favored personalization features based on their experiences over the six scenarios. They were also asked about their expectations regarding future interactions with a personalized Web site for completing these types of typical tasks.

Results

This research, which is exploratory in nature, was intended to investigate a broad range of personalization techniques in an experiential setting and identify the more valuable areas on which to focus in follow-up work (an anticipated live pilot study preceding deployment). The group walk-throughs in Study 1 and 2 gave us critical feedback that allowed us to filter out some personalization features and concentrate on the high-value features. The results of Study 3 produced more detailed feedback concerning the highly rated personalization features and the possible impact of these features on the volume of transactions on the site. We started our usability studies with a master list of 75 features and policies. We identified a final list of 12 features and policies that, according to the study participants, will provide a coherent and valuable personalized user experience on the site.

The average ratings for features over the three studies ranged from 4.4 to 6.4, with 7 being the highest and 1 the lowest possible score. In this exploratory research, we employed natural and large breaks in the data as determination points for highly valued features. In the follow-on pilot study of personalization on the site, there will be tests for statistically significant differences in the design alternatives for specific personalization features.

Table 3 contains the personalization features and policies that were rated highest in Study 1. The entries in the table reproduce the wording used on the questionnaire that the participants filled out. User ratings of the brief descriptions of features and policies shown in Table 3 were tied to their experiences of the demonstrations of the personalization features and policies illustrated in the context of the task scenarios and storyboards.

Table 4 shows mean ratings for the top 17 personalization features and policies in Study 2 and Study 3. The results are in general agreement with the find-

Table 3 Highest-rated personalization policies and features from Study 1

You are asked to provide only the information needed to allow you to access a particular feature.
A personal "myX" site is created for you when you provide information about yourself.
You are asked to provide information for your profile that will be active across the site.
You can choose to be called by, or chat with, a human representative who has access to your profile.
You can save shopping carts with price quotes, availability dates, and contact information in them.
You can view your order history.
You can create a wish list that contains items you may be planning to buy.
You can track transactions on the site.
You have the choice of having a search constrained based on current activities and profile data. The pages displayed are adapted based on your profile.
The pages displayed are adapted based on transient implicit information, such as your connection speed.

Table 4 Mean ratings for top 17 personalization policies and features in Study 2 and Study 3

Personalization Feature	Study 2	Study 3
User control of personal data	6.4*	6.4*
Automatic support alerts	5.9	6.2*
Order history provided	6.1*	6.1*
Help me find what I need	5.6	6.1*
Suggest alternate products	5.6	6.0*
List of products you own	5.5	5.8
Log-in feedback	5.6	5.8
Wish list	5.1	5.8
Personal Book	5.4	5.6
Saved shopping carts	5.6	5.6
Transaction tracking	6.2*	5.6
Only information needed is asked for	6.0*	5.5
Constrained search	5.6	5.5
Adapt presentation, transient data	5.5	5.5
Adaptive navigation	5.6	5.2
Information valid across the site	5.6	5.0
Contact company in context	5.5	5.0

* = Top-rated items in Study 2 and Study 3

ings from Study 1. Participants attached greatest value to user control over personal information as well as access to a transaction history.

Figure 4 illustrates the features available through "inventory-based personalization." The Products That I Own feature allows customers to maintain a list of products they own, and to find compatible accessories and upgrades easily and accurately by filtering the information presented on the site for these products. It also enables customers to get recommendations for replacing inventory items that are no longer in production. In addition, customers can track and review current and past purchases.

The "help me find what I need" feature (see Figure 5), which is an instance of context-sensitive help combined with restricted search capabilities, was rated by participants as "very valuable." This innovative feature is based on user-supplied information in response to application prompts. As shown in Figure 5, the application helps users narrow down the server products of interest (a patent on this user interface design feature has been filed).

The three personalization policies, included for evaluation in all three studies, received strong support from participants. Quantitative data was collected on user control of personal data and permission marketing policies, and they were in the top ratings by users in each study. The levels-of-identity policy was evaluated in several scenarios through user opt-in actions and through comments. This policy also received strong customer support.

Participants unanimously stated they would visit the site more often if personalization features were implemented, and the vast majority stated they would make more purchases. While there are differences in the ratings of features across the three studies, we do not find these differences to be significant. For example, Table 4 shows that the second ranked feature in Study 2 (transaction tracking) was tied for ninth in Study 3. The feature was still within the highly rated category for both studies (remember that Table 4 does not contain all personalization features tested—only those that were highly rated). We believe that variations in the individual and group tasks contributed to such minor differences.

The business case for personalization. All three studies included post-session questionnaires that elicited participants' attitudes toward future visits to, and

Figure 4 The Products That I Own inventory-based personalization feature

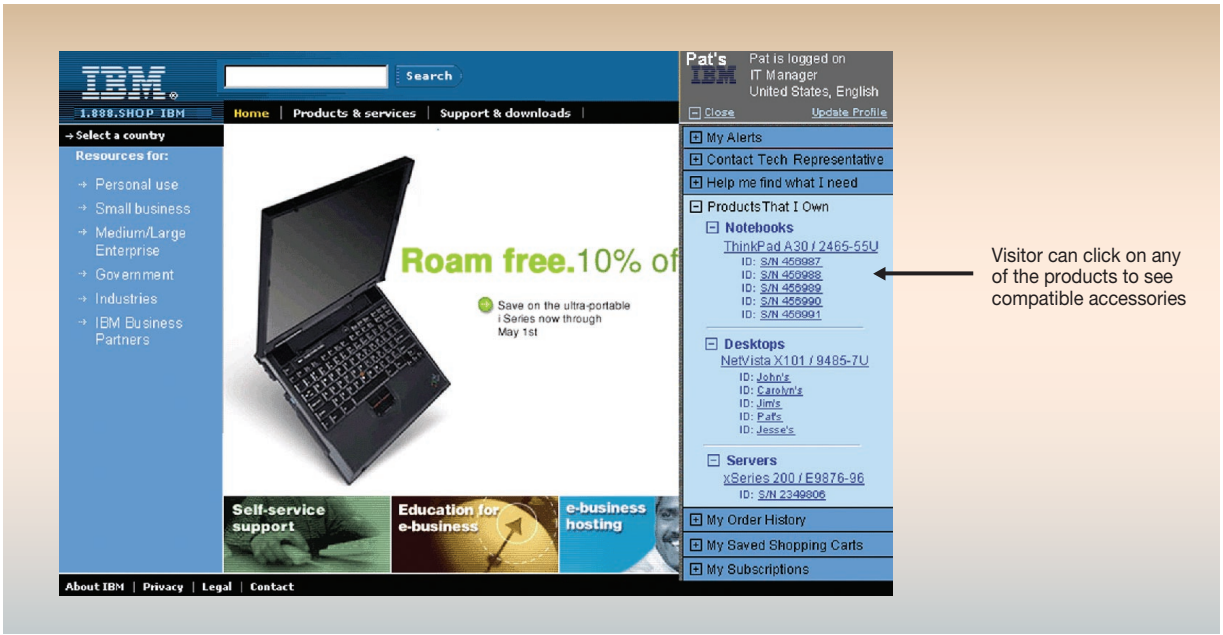

repeat purchases on, a site enhanced with the personalization features they identified as valuable. While we realize that such data might not accurately reflect actual future behavior, we believe they offer valuable input for making decisions about funding of future development efforts. The participants were asked, "If the features you rated highest were implemented on the X site, would you be more likely to use the site?" All 23 subjects in Study 2 and all 22 subjects in Study 3 responded, "Yes." The modal (most frequent) response to the follow-up question, "How often do you think you will visit the site over the next 12 months?" was 10+ (10 or more times) for both Study 2 and Study 3. The average response across the participants was 4.3 for Study 2 and 4.0 for Study 3, where 4 represents 5–10 visits, and 5 represents 10+ visits. When asked if they would be more likely to purchase from the site if the features of highest value were implemented, 22 out of 23 participants in Study 2 and all 22 participants in Study 3 responded in the affirmative (one participant responded, "Maybe"). The phrasing of the question was, "If these features were implemented on the X site, the company had a product that met your needs, and the product was available and was within your budget, would you be more likely to purchase from the site?" The participants' modal and mean response in terms of the number of additional purchases in a year was the range of 3–4 purchases. The modal response on the amount of the purchase was in the $5001–$10000 range. The mean across the participants was 5.8 in Study 2 and 5.7 in Study 3, with 5 representing $2501–$5000 and 6 representing $5001–$10000.

We used these results to build a business case for adding personalization features to the Web site. The projections were based on the following assumptions: (1) call center and Web-based revenue limited to the Americas (United States, Canada, Latin America), (2) the Web site involved is the subset of ibm.com dealing with personal computers and servers, (3) the targeted customer population is the one described in the user studies, and (4) a one-year horizon. Projections of increased site traffic and increased purchase transactions were added to existing mathematical and financial models relating site visits to revenue, and a significant business case for the value of personalization to the e-commerce site resulted. The costs of implementing the top-ranked personalization features were included in the business case as well. The projections were based on user statements on what they thought they would do, rather than on measures of actual performance. However,

Figure 5 The Help Me Find What I Need feature that constrains options based on user needs

because the participants responded to these questions after immersive personalized experiences and after having completed tasks relevant to their daily job responsibilities, such projections are viewed as better founded than most business case estimates for application software development.

The Personalization Value Model. Based on the affinity diagrams used to identify the business requirements (see the section "Identifying business requirements") and the comments made by participants during the immersive personalization experience in Study 3, we developed a Personalization Value Model (see Figure 6). The trigger for the model is the customer's first time opt-in (on the left in Figure 6). The system asks the customer for permission to use the minimal amount of personal information necessary to more efficiently and effectively carry out the task the customer is attempting to complete. The customer agrees and provides the information, and in exchange, the customer receives immediate value in the high quality completion of his task. The customer experiences improved ease of use in accomplishing tasks on the site through personalization. Customers highly value being in control of their personal data.

One customer, echoing many, told us "It makes me feel comfortable to be in control of my information. It makes me feel I can trust a company that is not looking to control or sell my information." Customers trust IBM because they are in control of their data and because IBM asks permission to use their data to provide them better service that they value. Customers told the researchers that they value the ability to complete tasks successfully and quickly on the site with personalization, and that the personalization functionality simplifies their jobs in small ways. As one customer expressed it: "If you can save me time or do some of the steps in my job for me so that I don't have to do them, that's real value to me, and you've got my business." Customers also thought that the personalization functionality such as the Products I Own and the Help Me Find What I Need tabs in the Personal Book allowed them to solve their own problems and made their decision-making simpler.

Customers expressed increased satisfaction with their personalized user experience on the site and stated that it would have financial and organizational benefits for them as well. The customer experience with personalization provides a feedback loop to a progressive opt-in to personalization, as they receive more value as more personal information is disclosed. The relationship between the customer and the organization is thus established and grown over time.

The top loop of the model shows the business value of personalization to ibm.com. With customer permission, personalization enables ibm.com to serve customers more effectively and efficiently as they inquire about information and make purchases or come to the site for after-sales support. Personalization allows ibm.com to improve marketing effectiveness and develop stronger customer relationships as customers self-select to receive promotional or other marketing information on specific products. These combined benefits produce financial and organizational benefits for ibm.com.

Personalization strategy and recommendations. Based on the results of our studies we recommended to ibm.com that the organization implement personalization on the site, including the three personalization policies—user control of personal data, permission marketing, and levels of identity. We also proposed nine specific personalization features that together with these three policies provide the basis for an effective personalized user experience for visitors to the site. Two personalization features, wish lists and saved shopping carts, were combined into one simply because participants saw it as one. Similarly, four personalization features that together formed a cohesive inventory-based personalization function through the Products I Own tab in the Personal Book were grouped together as one comprehensive feature. The twelve personalization features and policies are listed in Table 5 in the order of importance as suggested by the results of our user studies.

As can be seen in the table, user control of personal data was the most valued personalization function. Participants saw this policy as essential for their willingness to opt-in to a personalized user experience on the site. They welcomed being able to access and use their data through the Personal Book from any page on the site. Participants agreed that automatic updates on products they own, delivered to their Personal Book, to their e-mail address, or as a phone

Figure 6 Personalization Value Model

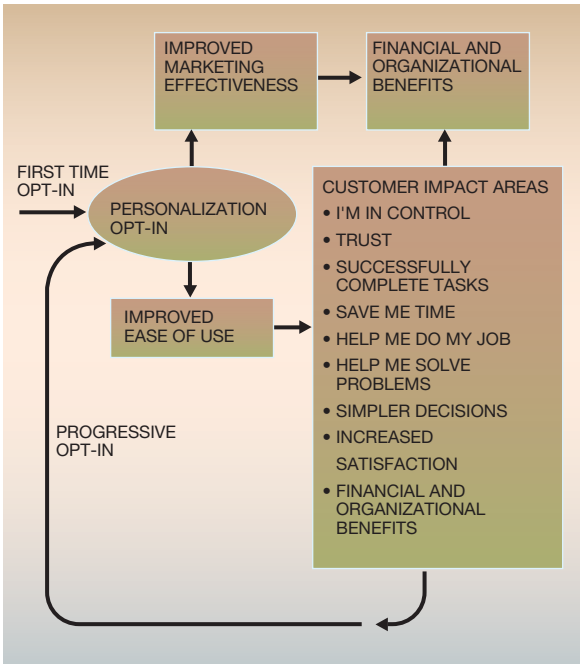

message, as preferred, would provide value and save them time. The Products That I Own inventory-based personalization function was viewed as providing ongoing value to customers by helping users find compatible accessories quickly and easily, recommending products to replace those no longer in production, and tracking inventory and current purchases.

The context-aware "help me find what I need" feature provides customers a way of quickly and effectively narrowing down options when searching for information. Permission marketing is valued as a means of earning a customer's trust and thus developing a business relationship with the customer. The log-in feedback, universal profile, and levels of identity features together assure customers that their identity is reliably handled by the site, regardless of where they are on the site, and that the information in their profiles will be used securely to provide enhanced service to them. Participants liked the idea of alerts informing them of special promotions related to the items saved in their shopping carts (future purchase considerations). They rated highly the ability to specify the expiration time for saved items in the shopping cart as well as the ability to share

Table 5 The 12 personalization policies and features recommended to ibm.com

Personalization Recommendations	Policy	Feature
1. User control of personal data: I control all the data in my profile and can review and edit it at any time.	x	
2. Automatic support: I can get automatic updates for the products that I own		x
3. Products that I own: I can view "products that I own" and get alternative recommendations for items that are no longer available, find compatible accessories and upgrades, and track/review current and past transactions.		x
4. Help me find what I need: I can use "help me find what I need" to help me filter through product choices and make purchase decisions. This enables me to easily conduct context-sensitive searches.		x
5. Permission marketing: I am asked to provide only the information needed to allow me access to the feature that helps me complete a task.	x	
6. Log-in feedback: After I have logged in, it is clear that the system knows who I am.		x
7. Universal profile: The information I provide is active across the entire site.		x
8. Future purchase considerations: This feature combines saved shopping carts with wish lists tailored for business buyers so that I can save shopping carts and indicate that I want to hear about special promotions on items in that cart.		x
9. Personal Book: A personal "my IBM" site is created for me when I provide information about myself.		x
10. Levels of identity: I can adopt the appropriate level of identity for the particular task on the site.	x	
11. Adaptive presentation: The pages displayed are adapted based on my recent navigation path (implicit data with session life span).		x
12. Contact IBM in context: I can communicate with IBM in the context of my profile and my current task. For example, I can give IBM technical support permission to see both my profile and my current navigation path while I am talking with IBM about my current goal on the site.		x

the shopping cart information with others in their organizations. They indicated that the use of implicit navigation data to adapt the presentation of information could be valuable if use of the data was limited to a single session. Although they liked this idea in concept, they thought that the technology had not matured sufficiently to ensure its usefulness. Participants also valued the ability to contact IBM in the context of their current task (Contact IBM in Context). They indicated that the ability to obtain an answer to a specific question in a "chat-like" session with a site support technician would be helpful. For more involved questions, participants preferred speaking on the phone with a site representative with whom they had previous interactions and therefore with whom they were comfortable sharing their profile data.

We suggested to ibm.com that the organization extend the current effort on personalization to other key areas of the site, such as software products, customer support, and services. Such an effort would first identify the extent to which the strategy we propose is applicable and where changes are needed. We also recommended that ibm.com develop guidelines to inform the development work underway on personalization, and we have participated in the creation of these guidelines.

The ibm.com organization is currently developing the information technology infrastructure to support the personalization features recommended above. Planning is underway for pilot studies and for rolling out the application across the 4 million pages and 2200 sub-sites of the ibm.com Web site.

Discussion and future work

In the work presented in this paper, we (1) identified a number of personalization features of interest, (2) explored these features in a set of usability studies, and (3) prioritized the list of features on the basis of costs and benefits to customers and providers. The results apply to information searching, sales, service, and support for servers, desktops, notebooks, and related accessories. The target customers are comfortable with the Web and at least moderately sophisticated in Web purchasing. A number of personalization features were highly rated in all three usability studies. Participants indicated they prefer

to provide personal information only to the extent necessary for the task at hand (permission marketing), they value the ability to access their transaction history, they value the ability to contact customer service personnel within the task context, and they would like more efficient search capabilities. They rated highly the context-based guidance provided by the "help me find what I need" feature. Participants also rated highly the Personal Book feature, as well as shopping carts with wish-list functionality. They unanimously stated that they would visit the site more often, and would make more purchases, if the most-valued personalization features were implemented.

We investigated personalization requirements only for a portion of the ibm.com site. There is a wealth of information contained in the comments generated during the group sessions and the written comments on the questionnaires. It is important to note that, according to participants, personalization features are valuable only if well-designed.

It is also interesting that the findings of Study 2 and Study 3 are consistent, given the differences in the User-Centered Design methods employed. In Study 2, the participants were involved in a design walk-through as a group and did not directly interact with the system—they reacted to a presentation by a facilitator. Although this presentation is designed to be engaging, it is still a passive experience. In Study 3, participants actively performed tasks similar to those discussed in the previous studies—they interacted with a prototype of the personalized Web site. The similarity in the results of the two studies is noteworthy. On one level, the cross-validation strengthens our confidence in the general findings. On another level, it provides some evidence about the value of the group methodology. Such group studies are often easier to conduct at early design stages. Our experience from this project has encouraged us to use this technique in future research. The similarities between the results from group and individual sessions provide additional evidence for the value of group design walk-throughs, as recently reported in the scientific literature.[24]

Based on what we have learned from these user studies, and with an eye toward a framework for future research, we propose the concept of a Personalization Value Space (PVS). We believe that the value of personalization features and policies to both users and providers must be appraised for a variety of business environments. Many factors within a given context can affect the value of any personalization feature or policy within that context. For example, on a Web site selling personal computers and servers, IT professionals making purchases for their organizations (transactions usually involving multiple units of the same item) rate highly the ability to view the transaction history over the past 10 years, whereas individual consumers buying home computers may not find this feature nearly as valuable. We hypothesize that once the value of personalization features and policies are understood for a subspace of PVS, it may be possible to predict the value at other points in this space.

We plan to explore PVS through a systematic examination of personalization policies (e.g., permission marketing, user control of personal data), feature categories (e.g., collaborative filtering, click stream analysis), user context (e.g., predisposition to trust, interaction goal), and business context (e.g., product offering, business goals). We believe that the effectiveness of personalization efforts is a function of the four above-mentioned components: policy, feature, user context, and business context. The determination of the exact functional relationship is a rich area for future research.

Conclusions

When starting the research work described in this paper our goals were to (1) understand the value of personalization to customers of ibm.com and to IBM, and (2) develop a strategy for introducing personalization to the Web site that provides the most benefit for customers and the business. We investigated a broad range of personalization techniques in the context of typical user tasks involving searching for information, purchase, service and support for servers, desktop and notebook computers, and related options and accessories. We conducted literature reviews, held brainstorming sessions, interviewed IBM stakeholders about the value of personalization to the business, performed heuristic usability evaluations of ibm.com and competitive sites, and conducted a series of three usability studies by using low- and medium- fidelity prototypes that we built for this purpose.

Our results indicate that personalization should not be considered in isolation, but rather as a space in which personalization features may take different values depending on user and business contexts. Thus personalizing interactions for e-business requires more than implementing a single function; it involves the development of a collection of functions that to-

gether achieve the desired goal. Our research identified a set of 12 personalization features with high value to customers and the business. This set of features is a key component in the personalization strategy planned for ibm.com. In the future, we propose to explore a Personalization Value Space concept that will extend the work described here by evaluating personalization features for a range of customers involved in e-commerce across a variety of domains. This concept may hold the promise of predictive capability in the design of personalized user experiences for e-commerce.

*Trademark or registered trademark of International Business Machines Corporation.

**Trademark or registered trademark of Microsoft Corp.

Cited references

1. P. Brusilovsky and M. T. Maybury, "From Adaptive Hypermedia to the Adaptive Web," *Communications of the ACM* **45**, No. 5, 31–33 (2002).
2. R. Burke, "Integrating Knowledge-Based and Collaborative Filtering," *Proceedings of AAAI 1999 Workshop on AI for Electronic Commerce* (*AIEC99*), (1999), pp. 14–20.
3. B. J. Schafer, J. Konstan, and J. Riedl, "Recommender Systems in E-commerce," *Proceedings of the First ACM Conference on Electronic Commerce*, ACM, New York (1999), pp. 158–166.
4. E. Schonberg, T. Cofino, R. Hoch, M. Podlaseck, and S. Spraragen, "Measuring Success," *Communications of the ACM* **43**, No. 8, 53–57 (2000).
5. S. Godin, *Permission Marketing: Turning Strangers into Friends, and Friends into Customers*, Simon and Schuster, New York (1999).
6. D. Peppers and M. Rogers, *Enterprise One to One: Tools for Competing in the Interactive Age*, Doubleday, New York (1997).
7. R. Barrett, P. Maglio, and D. Kellem, "How to Personalize the Web," *Proceedings of the Conference on Human Factors in Computing Systems* (*CHI'97*), Atlanta, GA, March, 1997, ACM, New York (1997), pp. 75–82.
8. J. Heer and E. H. Chi, "Separating the Swarm: Categorization Methods for User Sessions on the Web," *Proceedings of the Conference on Human Factors in Computing Systems* (*CHI 2002*), Minneapolis, MN, April, 2002, ACM, New York (2002), pp. 243–250.
9. B. Mobasher, R. Cooley, and J. Srinivastava, "Automatic Personalization Based on Web Usage Mining," *Communications of the ACM Special Issue on Personalization* **43**, No. 8, 142–151 (2000).
10. M. Perkowitz and O. Etzioni, "Adaptive Web Sites," *Communications of the ACM Special Issue on Personalization* **43**, No. 8, 152–158 (2000).
11. M. Spiliopoulou, "Web Usage Mining for Web Site Evaluation," *Communications of the ACM* **43**, No. 8, 127–135 (2000).
12. R. O. Duda, P. Hart, and D. G. Stork, *Pattern Classification*, John Wiley and Sons, New York (2001).
13. P. Resnick and H. R. Varian, "Recommender Systems," *Communications of the ACM* **40**, No. 3, 56–58 (1997).
14. D. VanderMeer, K. Dutta, A. Datta, K. Ramamritham, and S. B. Navanthe, "Enabling Scalable Online Personalization on the Web," *Proceedings of the Second ACM Conference on Electronic Commerce*, ACM, New York (2000), pp. 185–196.
15. V. Anupam, R. Hull, and B. Kumar, "Personalizing E-commerce Applications with Online Heuristic Decision Making," *Proceedings of the Tenth International World Wide Web Conference on World Wide Web*, ACM, New York (2001), pp. 296–307.
16. K. Vredenburg, S. Isensee, and C. Righi, *User-Centered Design: An Integrated Approach*, Prentice Hall, New York (2002).
17. H. Beyer and K. Holtzblatt, *Contextual Design*, Morgan Kaufmann Publishers, San Francisco, CA (1998).
18. D. Berreby, "Getting to Know You," *Think Research*, No. 1, 1–4 (1999), http://www.research.ibm.com/thinkresearch/index.shtml.
19. A. Kobsa, "Personalized Hypermedia and International Privacy," *Communications of the ACM* **5**, No. 5, 64–67 (2002).
20. Arent Fox—Attorneys at Law, Online Privacy Law, http://www.arentfox.com/home.html.
21. E. Volokh, "Personalization and Privacy," *Communications of the ACM* **43**, No. 8, 84–88 (2000).
22. P. R. Hagen, "Personalization versus Privacy," *Forrester Report*, 1–19 (November 2000).
23. J. S. Schaffer, personal communication (2001).
24. C. M. Karat, J. Karat, J. Vergo, C. Pinhanez, D. Reicken, and T. Cofino, "That's Entertainment! Designing Streaming, Multimedia Web Experiences," *International Journal of Human-Computer Interaction* **14**, Nos. 3/4, 369–384 (2002).

Accepted for publication May 21, 2003.

Clare-Marie Karat *Research Division, IBM Thomas J. Watson Research Center, 19 Skyline Drive, Hawthorne, NY 10532 (ckarat@us.ibm.com).* Dr. Karat is a research staff member at the Watson Research Center. She received her B.A. degree in psychology with honors from Stanford University and a Ph.D. in social psychology from the University of Colorado at Boulder. She conducts human-computer interaction (HCI) research in the areas of privacy, personalization, and conversational interface technologies. She is an international expert in cost-justifying human factors work and develops innovative user interface designs and methodologies. Dr. Karat has published numerous articles in professional and technical journals and has authored several chapters in recently published books. As an editorial board member of the ACM magazine *interactions* and the *Interacting with Computers Journal* of the British Computer Society and as a technical program committee member of the CHI, HFES, and INTERACT conferences, she provides leadership and maintains an active network of communication with HCI professionals in the field.

Carolyn Brodie *Research Division, IBM Thomas J. Watson Research Center, 19 Skyline Drive, Hawthorne, NY 10532 (brodiec@us.ibm.com).* Dr. Brodie is a research staff member at the Watson Research Center. She received her Ph.D. in computer science in 1999 from the University of Illinois at Urbana-Champaign. For her thesis, she developed a methodology for designing planning tools and she used this methodology to develop an interactive battle planning system for the U.S. Army. Since joining the Watson Research Center, her research has focused on methods for automating privacy policies in organizations, on personalization of Web sites and the ways it can benefit both the e-commerce companies and their customers, and on the use of collaboration tools to enhance information flow in large, distributed organizations.

John Karat *IBM Research Division, IBM Thomas J. Watson Research Center, 19 Skyline Drive, Hawthorne, NY 10532 (jkarat@us.ibm.com).* Dr. Karat is a research staff member at the Watson Research Center. He received a B.S. degree in mechanical engineering from Lehigh University, Bethlehem, PA, and a Ph.D. in cognitive psychology from the University of Colorado at Boulder. Dr. Karat conducts human-computer interaction (HCI) research in a variety of areas including privacy, personalization, and information management. He is chairman of the International Federation for Information Processing Technical Committee on HCI (IFIP TC13), North American editor for *Behaviour and Information Technology*, and Editor-in-Chief of the Kluwer Academic Publishers series on HCI. He is a member of the ACM SIGCHI Executive Committee and has been actively involved in ACM CHI and DIS, and the IFIP INTERACT conferences in the HCI field.

John Vergo *Research Division, IBM Thomas J. Watson Research Center, PO Box 218, Yorktown Heights, NY 10598 USA (jvergo@us.ibm.com).* Mr. Vergo is a member of the Technical Strategy team at the Watson Research Center in Yorktown Heights, New York. His research interests include human-computer interaction, User-Centered Design methods, multimodal user interfaces, e-commerce user experiences, speech recognition, natural language understanding, scientific visualization, 3D graphics, and software development methods. He has a B.S. degree in mathematics and psychology from the University at Albany and an M.S. degree in computer science from Polytechnic University, Brooklyn, NY.

Sherman R. Alpert *Research Division, IBM Thomas J. Watson Research Center, PO Box 704, Yorktown Heights, NY 10598 (salpert@us.ibm.com).* Mr. Alpert has been with the Watson Research Center since 1987. He is currently a Senior Software Engineer involved in user experience research. He received a B.S. degree in computer science from the State University of New York at Stony Brook, and an M.A. degree in computing in education from Columbia University Teachers College, where he has pursued additional studies toward a doctorate. He has been involved in research and software development in a variety of domains including educational technology, human-computer interaction, multimedia, and object-oriented programming and design. He serves on the editorial and review boards of several journals and on a number of conference and program committees.