

Chapter XII

Creating Adaptive Web Sites Using Personalization Techniques: A Unified, Integrated Approach and the Role of Evaluation

Charalampos Vassiliou, Dimitrios Stamoulis, Anastasios
Spiliotopoulos and Drakoulis Martakos
National and Kapodistrian University of Athens, Greece

INTRODUCTION

“Personalization uses information from tracking, mining and data analysis to customize a person’s interaction with a company’s products, services, web site and employees. Consumers and companies can benefit from the unique treatment resulting from personalization. Providing content of special interest to your visitor can help establish a relationship that you can build upon each time that person returns to your site” (Deitel, 2001). Not only time constraints, but also the ‘lost into the cyberspace’ phenomenon creates pressure over the user to browse as fast as possible the web content to arrive at what s/he is really looking for. Therefore, personalization aims at satisfying the user by presenting those items, which are indeed valuable for him/her, and optimise the interaction for maximum efficiency and effectiveness. Subsequently, the creation of adaptive web sites emerges as a ‘sine qua non’ criterion for contemporary high quality e-services.

Personalization can be applied either at the form or at the content. Consequently, it is usually implemented at two levels: at the level of the interface or context through which the user interacts with the content, and at the level of the content itself. Most of the contemporary web sites offer the user the capability of creating their own web site, e.g. ‘My Banking Page’, not only to endow it with the feeling of ownership—obviously what you own, you also care for—but also to allow the user to interact in the way s/he defines as more appropriate and comfortable. “Excite is a search engine that offers ‘My Excite Start Page’”. This allows you to select the content and style that appears on your Excite home page” (Deitel, 2001). Such a user-initiated process is sometimes referred to in the literature as ‘personalization’, whereas the customisation, which is used in a more generic meaning, that results “from the site’s ability to tailor itself to each use” and is “designed to be altered by the organization” is called ‘tailoring’ (Rayport, 2001). This chapter prefers to use the term personalization to deal only with the organization-initiated customisation of the web content and services and does not at all discuss how to configure a web site according to an individual’s preferences that are explicitly input.

Personalization or “tailoring by site enables the site to reconfigure itself based on past behaviour by the user or by other users with similar profiles. These sites can make recommendations based on past purchases, can filter marketing messages based on user interests and adjust prices and products based on user profiles. Amazon makes recommendations across product categories. For example, based on a user’s history of book purchases, the site recommends CDs or DVDs that others with similar book interests have bought.” (Rayport, 2001). In essence, personalization is about correct guessing about what the user perceives as having added value for him/her. Since either a customer appreciates an offering or the customer is getting annoyed, no recommendations at all are probably better than

wrong and disturbing recommendations to the user. Hence, a need arises to evaluate the personalized offerings in terms of customer/user satisfaction. However, evaluation is an organizational process that starts long before a set of personalized e-services is being launched. This organizational activity of evaluation tends to be neglected, despite that the contemporary logging capabilities of web servers make the task of recording the traces of digital interactions an almost trivial task. However, “the fact that activity on the web is “measurable” is alone not enough: measurements need to be *meaningful*. Much work is needed to develop analytical approaches to the data that yield meaningful statistics. As impressive as these tools are, the analysis of web activity has only just begun to scratch the surface of what may be possible in the future. Using server logs as a foundation, additional data can be gathered via online user surveys (in conjunction with user registration) or third party data collection agencies, such as those which runs multi-site banner ad campaigns for clients” (Digital Enterprise). More than so, evaluating the success of a personalized e-service is truly a virgin ground.

How can success be defined in the context of the provision of personalized e-services? Theoretically, successful personalization means that the automatic propositions by a web site fully address the core of the visitor’s / customer’s desires, expectations and needs, and really reflects the user’s prioritisation of the information provided. One would wonder if any kind of tool could always ‘poll the visitor’s mind’ to find out whether he is satisfied with the personalized propositions or not. Therefore, implicit measures of the user’s satisfaction, and, equivalently, the degree of the personalization success, have to be identified. Moreover, these implicit measurement results have to be recorded and, even worst, to be fed back to the personalization techniques so that the personalization engine learns while interacting with the user, and, thus, avoids making the same wrong suggestions for the user’s likes and needs. Possible sets of metrics criteria are discussed in this chapter, as well as their relevance to the problem of evaluating personalized e-services; some mechanisms that seem appropriate for that task are also presented.

Concluding, this chapter tutorially presents how systems adaptivity is best exemplified in the domain of web sites by the premise of the personalization techniques. The main contribution of this chapter is a unified approach that integrates the main categories of the personalization techniques as well as the identification of the need to evaluate the automated propositions made by adaptive web sites. Adaptivity relates to the site’s capability of creating tailor-made form and content. Moreover, the importance of defining success factors for the assessment of the personalization techniques and the role of evaluation is stressed. The chapter proposes the modification of existing personalization algorithms/ techniques so that they can accept as another input the post-personalization evaluation results. To insert feedback into existing personalization algorithms/ techniques is our ongoing

further research effort. Summarizing all of the above, the main outcomes presented in this chapter include a state of the art critical review in personalization techniques, a framework for integrating different personalization techniques into a single, unified approach along the phases of the customer decision process, the enhancement of the traditional personalization chain with the evaluation phase which results are feedback into the personalization techniques. To indicatively show how to extend personalization architectures with the proposed evaluation layer, an existing architecture has been adapted accordingly. The vision of this research is to embed a learning capability (through a fuzzy logic system, a neural network etc.) into the personalization techniques so that they avoid making the same unsuccessful suggestions in terms of links and features not being valued by the user. By achieving this goal, personalization-enabled web sites will behave as adaptive and evolutionary information systems.

BACKGROUND

From the human perspective, representation ultimately affects presentation and interaction with a service or product. True personalization implies not only adapting content to the individual, but also how that content is communicated for maximum effect. According to Pednault (Pednault, 2000), representation in personalization is divided in the technology and the human aspect. On the technology side, representation involves data structures used to implement a personalized service or product. Appropriate data structures are required not only to support the various media employed, but also to encode user-specific data needed to determine the current user context and goals, and to define what information to present to the user when, and in what media or interactive form. On the human side, just as people adapt to each other, similarly a system should begin to adapt to the way the interacting user wishes to communicate, receive and organize information. The issues arising outline what needs to be represented in order to achieve the desired level of adaptation, and how to represent it in the form of data structures that yield efficient algorithms for carrying out the adaptation. The blending of these two representation aspects produces the final personalized end-user experience.

A personalized interactive service or product, by its very nature, should respond in real time to user inputs. To do so, a system must have quick access to the right information at the right time to decide what to do next. Satisfying this requirement might well entail the use of specific data structures that conflict with the needs of other components of the system, such as learning algorithms. For the purpose of learning predictive models of user behaviour, large volumes of data must

be recorded for later analysis. Compactness of the representation and ease of storing the data are the main design considerations for this task. In all likelihood, the best data structures for storing data would be quite different from the best data structures for quickly accessing the relevant subsets of data needed in each interactive context. Likewise, the best data structures for applying the learning results might be quite different from those used during the learning process itself.

The key to personalization is the understanding of the customer's desires and needs and in data-processing terms that means organizing and building customer data stores and attaching those stores close to the customer interaction to effect the interaction and personalize, really personalize the experience. These multiple logical data stores necessary to fulfil the total requirements of the customer experience are defined in an information model (Wells and Wolfers, 2000). The model encompasses traditional customer details (name, address, place of work, income), family situation (children), financial history, transaction history, and behaviour. Customer behaviour is the piece that is not reliable and consistently captured today, but it is a critical element. It must be captured across every channel touched by a customer, organized and stored. Information can be gathered in many ways. One of the best ways is simply to ask the customer for it. Further, the model also captures information describing the style of customer behaviour. The information model will also have several traditional data warehouse-style processes applied to it. These processes will, for example, cluster customers in behaviour domains.

Computer systems that perform personalization on the user needs require recognizing patterns in these users behaviour. Understanding user behaviour and learning to personalize require the building of a model for the user. Assuming no additional information exists, by simply assuming that user actions in the past are repeated in the future, we can predict a user's future interactions with a computer system, finding in the past actions similar to those taken in the present. These predictions construct a user profile containing information on what the user has previously seen or done, based on the user's history of past actions. Nevertheless predictions are limited by the fact that this captures one dimension of user's behaviour. Hirsh argues (Hirsh et al, 2000) that we are compelled to search for other sources of data that can complement a model based on the shallow history of user's past actions. Each user action takes place in the context of a specific task, and the main question is what information is available so that the process of predicting the future can be maximally informed by the past. Our ultimate goal is to predict a user's actions while exploiting as much information as can be obtained. Every method for predicting a user's future actions is based on some form of user profile or model that links information about the user or the task to expectations about the user's behaviour.

A typical user interaction exhibits many patterns. Machine-learning algorithms are being used to recognize such regularities and integrate them into the system, to personalize the system's propositions to its user. Systems that achieve such automatic personalization have been called "self-customizing software," in that the system's responses to a user are automatically customized to the personal characteristics of the user (Schlimmer and Hermens, 1993).

A number of researchers have developed methods that predict user's actions based solely on the user's history of past interactions with a computer system. Interestingly, methods that look solely at the user's single immediately preceding action, comparing it to past situations where the user previously took that action, often give rise to surprisingly good predictions (Davidson & Hirsh, 1998) (Padmanabhan & Mogul, 1996).

According to (Mulvenna et al, 2000), personalization is the provision to the individual of tailored products, service, information or information relating to products or services. This broad area also covers recommender systems, customisation, and adaptive Web sites.

Three aspects of a Web site affect its utility in providing the intended service to its users. These are the content provided on the Web site, the layout of the individual pages, and the structure of the entire Web site itself. Personalization technology involves software that learns patterns, habits and preferences. Initial attempts at achieving personalization on the Internet have been limited to check-box personalization, where portals allow the user to select the links they would like on their "personal" page. However, this has limited use since it depends on the users knowing beforehand the content of interest to them. Furthermore this "suggestive" form of personalization cannot be considered true personalization since it is not based on user profiles indicating behaviour and preferences.

On the other hand according to (Wells & Wolfers, 2000) currently an emergence of two types of personalization can be seen on the Internet: One offers users the ability to become GUI editors by allowing them to construct personalized pages; the other targets marketing of products and services based on information held about an individual. Wells argues that neither really meets the needs of the customer. In the first case, the Internet site will often present a rather long and confusing list of series of boxes to check, resulting in an uneasy feeling, or worse, an unpleasant feeling on the part of customers. Equally awkward, the resulting sites often separate their personalized sections 'my.com' from the rest of the site giving the feeling the personalized piece has simply been tacked onto the side of the site. In the second more advanced case, sites will often collect data about the customer's behaviours, apply a series of rules to the site, and as a result radically change the content of the site when the customer returns. This results in a jarring experience for the customer and the very thing the site was trying to achieve was not achieved

because of a clumsy implementation. However, Wells and Wolfers (Wells & Wolfers, 2000) argue that these two approaches, however, hold kernels of goodness and when refined and combined will produce the type of experience that encourages repeat visits, and more importantly, encourages visitors to turn into actual revenue producing customers.

In contrast, content-based filtering (also content-based prediction) was the first attempt at using AI for achieving personalization in a more intelligent manner. Consider a person reading online sport news. What we would like is a system that observed what sport news the user has read - and, more importantly, has not read - and learns to present the user with new articles the user will want to read. Although there is clearly systematization in the user's actions, here the patterns require deeper analysis. Rather than mimicking user actions taken in the past, a system that effectively personalizes itself to a user's sport news interests must look inside the news to understand how to distinguish those news that interest the user from those that do not. Systems that personalize in such a fashion are often said to be "content-based," in that they base their predictions on the contents of the artefacts about which they are concerned.

Content-based filtering originates from information retrieval and case-based reasoning research (Hammond et al, 1996). The success of the content-based method relies on an ability to accurately represent recommendable items in terms of a suitable set of content features, and to represent user profile information in terms of a similar feature set. The relevance of a given content item to a specific target user is proportional to the similarity of this item to the user's profile; content-based filtering methods select content items that have a high degree of similarity to the user's profile. A major drawback of the above-mentioned technique is that the content description requirement can be problematic and time consuming. Additionally, content-based techniques also suffer from a number of shortcomings in the way they select items for recommendation. Limited set of items representing user profiles, especially immature new user profiles, will result into future recommendations to display limited diversity.

In summary, content-based approaches typically offer us a means for describing items of user interest and a means for comparing item descriptions to locate close matches. However, when using these approaches, we also usually find ourselves considering the preferences of a single user. This is in contrast to our expectations, that we should be able to exploit such additional information in learning to predict a user's interests (Hirsh et al, 2000).

A recent alternative to content-based strategies is collaborative filtering techniques (Balabanovic & Shoham, 1997) (Coldberg et al, 1992) (Maltz & Ehrlich, 1995). This allows users to take advantage of other users' behavioural activities based on a measure of similarity between them. These techniques require

users to divulge some personal information on their interests, likes and dislikes, information that many Web users would not necessarily wish to divulge (Mulvenna et al, 2000). Consider how you decide whether to read a particular book. Sometimes our decisions are based on publishers reviews, but often we are triggered by “word-of-mouth”, based our decisions on feedback from others whose opinions we value and share. This is the basic idea underlying the collaborative filtering method. According to (Smyth & Cotter, 2000), the basic idea is to move beyond the experience of an individual user profile; instead, to draw on the experiences of a population or community of users. Typically, each target user is associated with a set of nearest-neighbour users by comparing the profile information provided by the target user to the profiles of other users. Collaborative filtering techniques look for correlations between users in terms of their ratings assigned to items in a user profile. The nearest-neighbour users are those that display the strongest correlation to the target user. These users then act as “recommendation partners” for the target user, and items that occur in their profiles (but not in the target user profile) can be recommended to the target user. In this way, items are recommended on the basis of user similarity rather than item similarity. “In summary, collaborative filtering compares ratings of a present user’s interests and decisions with those of past users to offer content relative to the present user’s interests. Music and book sites often use collaborative filtering to make recommendations to their customers” (Deitel, 2001).

Finally, another alternative is observational personalization, which attempts to circumvent the need for users to divulge any personal information. The underlying assumption in this approach is that within records of a user’s previous navigation behaviour there are hidden clues to how services, products, and information need to be personalized for enhanced Web interaction. According to (Mulvenna et al, 2000), there are three principal components to observational personalization: analytics, representation, and deployment. Web mining provides the tools to analyse Web log data in a user-centric manner such as segmentation, profiling, and clickstream discovery. The knowledge mined by using these tools is increasingly being represented using W3C standards such as XML and the deployment of the knowledge on Web servers may be carried out through personalization or recommender systems.

In general, personalization techniques can be analysed in Offline and Online. Offline personalization is based on simple user profiling and manual decision rule systems. Manual decision rule systems, allow Web site administrators/marketers to specify business rules based on user demographics or static profiles, collected through a registration process or session history. “Rules-based personalization is the delivery of personalized content based on the subjection of a user’s profile to set rules or assumptions” (Deitel, 2001). The rules are used to affect the content

served to a particular user, based on relationship analysis. Online personalization demands advanced real-time adaptive user profiles in order to identify and observe the customer, define the objectives, identify the value and provide the personalized content. A slightly different description, hierarchy and categorization of personalization techniques are presented by a report from Gartner Group (Gartner Group, 2000). More specifically the personalization techniques according to Gartner Group are:

- *Relationship Analysis*: Analyses of previous interactions with the customer based on online and offline purchases.
- *Contextual Inference*: Analysis of the content being viewed and then displaying related content.
- *Clickstream Analysis*: Collects data about what the visitor is viewing and then displays related content.
- *Profiling (Content-based filtering)*: Collects explicit preference data from visitor and then matches the resulting profile to predefined content.
- *Preference Matching (Collaborative filtering)*: Explicit collection of preferences that are then matched to other people's preferences.

MAIN THRUST OF THE CHAPTER

Looking critically at the advantages and disadvantages of the personalization techniques mentioned above, the following can be identified.

Collaborative filtering has a number of advantages over content-based methods.

- Firstly, since explicit content representations are not needed, the knowledge engineering problem associated with content-based methods is mitigated.
- Secondly and more importantly, perhaps, the quality of collaborative filtering typically increases with the size of the user population, and collaborative recommendations benefit from improved diversity when compared to content-based recommendations.

However collaborative filtering does suffer from a number of significant downsides.

- Firstly, it is not suitable for recommending new items or one-off content items because these techniques can only recommend items already rated by other users. If a new or one-off item is added to the content database, there can be a significant delay before this item will be considered for recommendation. It is only until many users have seen and rated an item, that this item will find its way into enough user profiles to become available for recommendation. This so-called "latency problem" is a serious limitation that often renders a pure

collaborative recommendation strategy inappropriate for a given application domain.

- Secondly, collaborative recommendation can also prove unsatisfactory in dealing with what might be termed an “unusual user.” In short, there is no guarantee a set of recommendation partners will be available for a given target user, especially if there is insufficient overlap between the target profile and other profiles. If a target profile contains only a small number of ratings or contains ratings for a set of items that nobody else has reviewed, it may not be possible to make a reliable recommendation using the collaborative technique.

So if content-based methods exploit one kind of information (about the contents of each item the user accesses) and collaborative methods exploit a second kind of information (about what others thought of each item), then combining both sources of information should do even better? Hirsh argues positively (Hirsh et al, 2000), but how to do so is not immediately clear. Content-based methods don't provide obvious ways to exploit information about other users, and collaborative methods don't provide obvious ways to exploit information about the contents of the items under consideration. Both content-based and collaborative personalization methods suffer from a number of significant disadvantages. However, taken together, both techniques complement each other perfectly. For example, content based filtering can solve the latency problems associated with collaborative filtering. Furthermore, introducing a collaborative component solves the diversity problem associated with content-based methods.

Content-based methods build models that link information about the contents of items a user manipulates to the user's preferences concerning those items. Collaborative filtering methods build models that link information about other users' preferences to those of a given user. By integrating both content-based and collaborative filtering strategies, a personalization engine could provide a unique and powerful personalization solution.

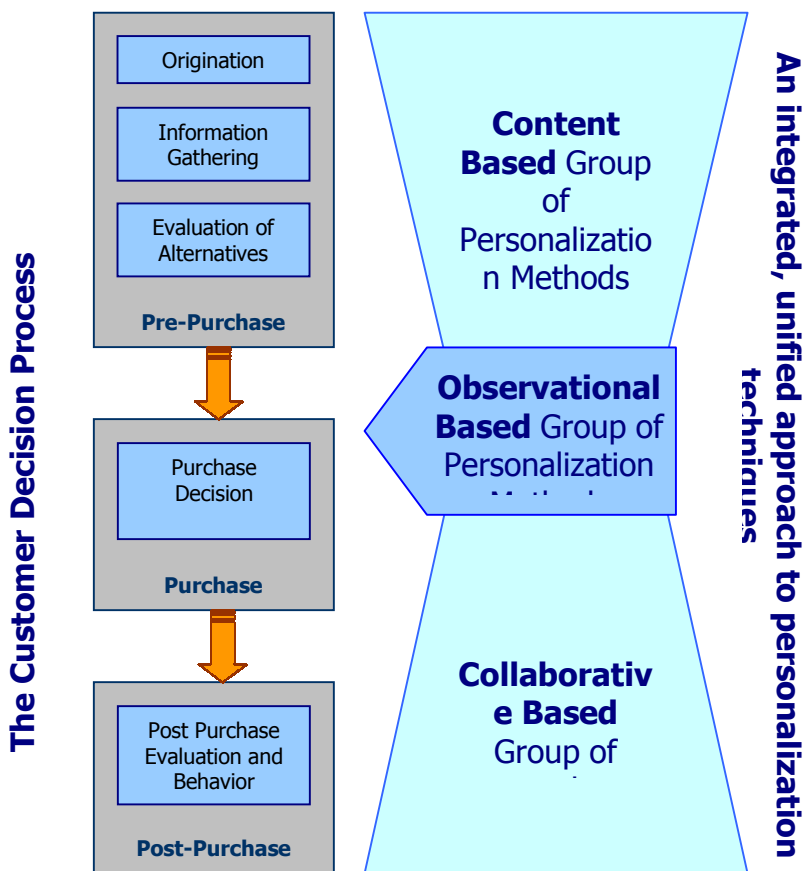
Critically contributing towards the positive argument of Hirsh, we are suggesting a framework of combining various personalization techniques to the customer decision process, as described by Rayport (Rayport 2001). “The customer decision process maps the activities and the choices customers make in accessing a specific experience within a value system, then lays out the series of steps from awareness of the experience to the purchase experience and through the use experience”. Successful selling over the Web will be measured not only on the capability to sell to a consumer once, but also by the capability to play a significant role in a long-term customer relationship. Similar to other customer interfaces, such as retail brick-and-mortar and call centres, success on the Web site must involve adding value throughout marketing, sales and customer service interactions,

ultimately using various personalization techniques. The three distinctive steps for building online consumer relationships are:

- Get Consumers to Come (Pre-Purchase)
- Once There, Add Value (Purchase), and
- Get Them to Come Back (Post-Purchase)

In any of these three phases, consumers have different objectives, scopes, intentions and shopping habits. In order to achieve and build a more efficient online consumer relationship, a different personalization technique is suggested for each of the three phases described above. The figure below presents a framework of understanding for personalization methods complementarities through the various phases of building online consumer relationships as well as the optimum use of personalization techniques per phase.

Figure 1: Building online consumer relationships using personalization techniques



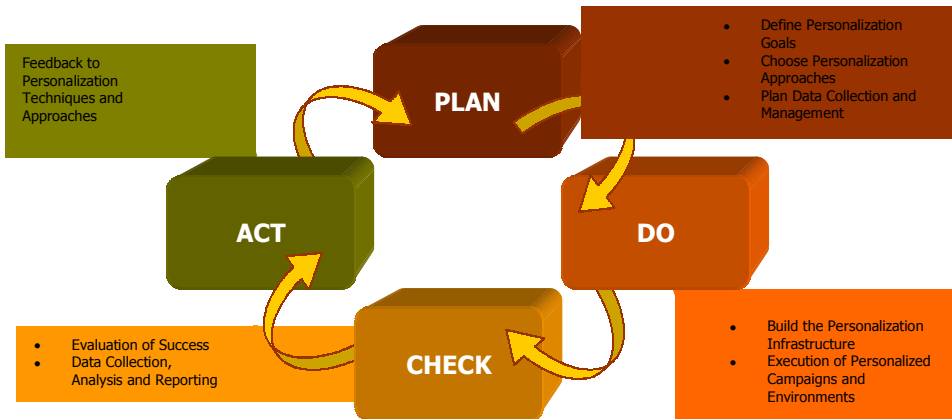
There are seven stages to implementing an effective personalization effort. The section that follows covers three of the initial steps: Defining online personalization goals, evaluating personalization approaches, and planning for data collection and management. Later, we will explore building the technical infrastructure, delivering personalized campaigns and environments, measuring success, and conducting ongoing data collection, analysis and reporting.

The core value of online personalization lies at delivering businesses the capability of establishing customer relationship and customer value management lifecycles. It may seem obvious; however, no personalization effort is complete without a mechanism for assessing each initiative's effectiveness and a process by which subsequent efforts can be optimised to achieve higher level of success. Companies must first assess the impact of each initiative - whether that is a site-based promotion or an e-Mail campaign - and second, they must be diligent in modifying and improving initiatives in an iterative fashion. To accomplish this in a more general approach, organizations should follow the steps outlined below - the first few of which are the basis of the initial data collection and management effort:

- Identify business objectives for which personalization should have a leveraging effect
- Define personalization goals
- Determine the metrics the organization is looking to apply (e.g., customer profit per ad campaign)
- Identify the data the organization needs to determine the data required for the evaluation process (e.g., number of customers, sales volume, gross margin, ad campaign cost, operating overhead, click through rates, conversion rates)
- Develop a solution which is appropriate for the specific personalization goals to analyse the data (e.g., NetGenesis, OLAP tools)
- Identify the location of that data (e.g., Web logs, application logs, event logs)
- Deploy the solution to target the right customer segment
- Produce metrics reports, and
- Automatically feed the metric results back to the personalization goals and techniques used.

Illustrated below is a view of the staged approach for online personalization initiatives that this chapter proposes. Since the personalization initiative should be a continuous improvement effort towards achieving the ultimate fit with the customer's explicit and latent needs, desires and preferences, the full personalization initiative should follow the PDCA cycle. "Plan-Do-Check-Act refers to Deming/ Shewart cycle of continuous improvement, which is fundamental total quality management concept" (Cohen, 1995).

Figure 2: The full iterative cycle for personalization success



In more detail the steps of building a personalization initiative include:

- **PLAN:**
 - o **Defining Personalization Goals:** Before an organization begins to plan their approach to personalization, they must first define they we are setting out to accomplish, because the organization cannot measure what they cannot define. Personalization initiatives must begin with an understanding of what personalization means to the organization. The type of business they are in, the customers they sell to and even the products they bring to market will all impact the ways in which they will use personalization and the benefits they seek to gain from it. Personalization initiatives should be tied to discrete business goals.
 - o **Choosing Personalization Approaches:** After determining the scope and magnitude of the effort, organizations must then match the ends with the means. Having set the benchmarks related to specific desired outcomes, companies must then assess which approach or combination of approaches to personalization will best suit their needs. Earlier in this section we have presented a critical extensive description of various distinct personalization approaches as well as a framework for building online consumer relationships using personalization techniques.
 - o **Planning for Data Collection and Management:** It is a way to deliver to the organizations customers' information, incentives and sales opportunities that are timely and pertinent. In order to do this well, it is necessary to understand who the organizations customers are, what they like and dislike and how they interact with the organization through each

one of their channels.

- **DO:**
 - o **Data Sources for Online Personalization Efforts:** The primary means for data collection to support the online effort are research, site behaviour and usage, marketing campaigns and enterprise data. Through research, both primary and secondary, companies can gain information that will assist them in building the foundation for initial customer management and messaging strategies. Secondary research is often considered a good starting point.
 - o **Customer Profiles:** The product of these combined efforts will be information about the organizations' customers that can be used to establish customer profiles - centralized sources of information about each customer. Profiles are the collection of attributes that characterize the explicit, implicit, demographic and psychographic elements of each customer's interaction with the organization. These profiles are the product of the entire data collection and management effort and will be the foundation of their personalization strategies.
 - o **Data Management:** Significant energy will need to be dedicated to determining the best strategies for storing the information that organizations collect. The information that will be relied upon most for generating reports and driving personalization should be stored in such a way that it remains easily accessible. Less frequently used data can be stored elsewhere.
- **CHECK:** Through this evaluation phase, organizations will be able to compare each initiative's results with the business objectives initially sought and, hopefully, collect the metrics' results to feed the personalization techniques for optimisation.
- **ACT:** As a result, companies must continue to optimise campaigns, sites and business initiatives over time, always making improvements based on proven success records and the information provided by ongoing data collection and analysis.

Depending on the decisions the organization makes upon each of the steps of building its personalization initiative, the desired result is obviously the creation of a stable and scalable personalization infrastructure.

As an organization builds the infrastructure to support online personalization, it is necessary to plan for the extensibility of the solution. As they would with any new system, be sure to lay the groundwork for integrating additional infrastructure components and leveraging new data sources. When they choose to extend the initiative into the rest of the enterprise, they will need such capabilities.

While current personalization systems may use sophisticated algorithms and techniques, they also hardwire the interaction sequences in their interfaces. For example, a personalization facility at an online bookseller may have some users who think of books primarily by title, others who look for a particular author, and still others who would like to personalize with respect to a combination of features. To cover all potential scenarios, the system designer must anticipate every type of situation beforehand and implement customisation interfaces (algorithms) for all of them. Ramakrishnan (Ramakrishnan, 2000) argues that the absence of an adequate programming model means that designers must make many assumptions and simplifications in the interface design. Some of these result from necessity (“This site is organized in this manner and I can’t help it”); others may reflect a lack of understanding or appreciation of user needs (“I think this is the best interface to my customers”). In either case, the user may experience serious cognitive and representational frustrations, because the modes of interaction are hardwired (for example, “This interface works only if you specify both the ISBN number and the title”).

Additionally, designing a complex Web site so that it readily yields its information is a difficult task. The designer must anticipate or even predict the users’ needs and structure the site accordingly. However, users may have vastly differing views of the site’s information, their needs may change over time, and their usage patterns may violate the designer’s initial expectations. As a result, Web sites are all too often fossils cast in HTML, while user navigation is idiosyncratic and evolving. Understanding user needs requires understanding how users view the data available and how they actually use the site. For a complex site, this can be difficult since user tests are expensive and time-consuming, and the site’s server logs contain massive amounts of data. As an indicative example, imagine a site devoted to information about automobiles. The Web master initially decides to organize the site by manufacturer; each auto company will be represented by a dedicated page, containing links to each of their models. However, many visitors to this site intend to comparison shop; a visitor wanting to compare minivans, for example, would have to go to each manufacturer and look up their minivan offerings.

While adaptive Web sites are potentially valuable, Perkowitz & Etzioni (Perkowitz & Etzioni, 2000) state that their feasibility is still unproven: can nontrivial adaptations be automated? Will adaptive Web sites run amok, yielding chaos rather than improvement? What is an appropriate division of labour between the automated system and the human Web master? To investigate these issues empirically, Perkowitz & Etzioni considered a case study where they analysed the problem of automatic index page synthesis (an index page is a page consisting of links to a set of pages that cover a particular topic) based on visitor access patterns and suggest an approach for a Web management assistant: a system that can

process massive amounts of data about site usage and suggest useful adaptations to the Web master. The use of such assistants is one way to develop adaptive Web sites—sites that semi-automatically improve their organization and presentation by learning from visitor access patterns.

Looking at the evaluation of personalised web services through a more e-business perspective the criteria for measuring success and feedback differ. How can we measure success in the design and evolution of personalized interactive services for e-business? The ability to design, implement, and maintain user interfaces and user navigation in personalized interactive services requires defining meaningful metrics and feedback techniques.

In order to enhance the evaluation methodology of e-business personalised web sites we need to utilise e-business intelligence. E-business intelligence is the analysis and use of information collected about visitors to an e-business Web site. According to Schonberg (Schonberg et al., 2000) good business practice dictates the use of effectiveness measurements to guide the design of all Web site features. For Web sites with personalized interactive content, the process must take the highly dynamic nature of the content into account and the outline of a complete process for a *design-measure-analyse feedback cycle*.

In order to measure and evaluate the successful provision of personalised e-services to the end customer we need to understand what success means. Success of an e-business site usually resides to the answers on a set of questions such as: What types of visitors does an e-business want to attract, what messages need to be conveyed, what should the visitor be able to accomplish, and what does the e-business want the visitor to do? The metrics required to evaluate success follow directly from the goals.

Notably, measuring the success of personalization initiatives can now expand the simple customer acquisition metrics that dominated the 1990s. Focus is now on correlating campaign/promotion metrics, such as acquisition and conversion rates, to the primary goal of each initiative (e.g., actual or ongoing sales, registration, data collection, in-store traffic, etc.). In the same manner that the strength of online personalization efforts can be bolstered through the use of enterprise data, their impact can now be better understood, determined and justified by evaluating them, in part, with traditional business metrics (e.g., sales volume, gross profit, ROI, etc.).

By basing evaluation on the same metrics, this approach to measurement enables organizations to align their online and enterprise initiatives and to also learn the most successful tactics for managing profitable customer relationships. At the end, it is a company's ability to effectively fine-tune its personalized approach to customer management with its most profitable segments that will more definitively result in the benefits described earlier. By closely monitoring the effectiveness of certain marketing campaigns or discrete initiatives as well as the behavioural and

transaction history of customers, companies will also have the added benefit of being able to track ROI at a much more granular level than in the past.

An interesting question is what metrics are best for evaluating the effectiveness of Web site design features? An interesting and worthy approach for evaluating the effectiveness of Web site design and personalization features can be based on *click-through* and *look-to-buy* metrics. Using an example from the online ad-banner industry, click-through data measures the ratio of clicks to impressions, where an impression is simply the display of an ad banner on a Web page. A high click-through rate means visitors who see the ad click on it frequently, therefore, the ad is bringing many visitors to the site. Look-to-buy data compares ad banner impressions with sales transactions and revenue directly attributable to the ad banner. It is a better measure of ad banner effectiveness, since the quality of visitors coming from the ad banner is captured and return on investment more accurately measured.

Look-to-buy metrics work well for dynamic, personalized content. In fact, ad banners fall into this category - ads typically are dynamically rotated and may also be personalized. With look-to-buy metrics, each personalized component on a page can be counted and its effectiveness evaluated. Generally, however, if the goal is something other than maximizing sales, the appropriate metric would be look-to-X, where X is the goal. In addition to the metrics mentioned above, additional supportive metrics should be defined in order to provide a more structured and concrete evaluation feedback. Such metrics include: Repeat business, Click through ratio, Time spent, Order Size, Buying frequency, Satisfaction/return rate, Web-influenced purchases.

Since success can only be implicitly inferred from the user actions, evaluating the success of personalised propositions is unavoidably based on assumptions. For example a newspaper filtering and personalization system that “Re-orders each major index session (e.g. international news), including the front page, according to the user preferences”, assumes that “since the user followed a link to the article’s body they must have found the lead relevant, even if the actual body proved not to be interesting upon further reading” (Kolcz, 1999).

According to Schonberg (Schonberg et al., 2000), the ability to collect and combine customer data from multiple sources enables richer analysis. Click - stream data, which captures the sequence of Web pages seen by each visitor to a Web site, is the standard data source for tracking visitors browsing behaviour. However, voluminous as this data is, it is low level and contains limited information. Many useful metrics cannot be calculated with click - stream data alone. Integrating click - stream data with other sources considerably expands the quality of information. Furthermore, newer technologies and services make large-scale collection and sharing of data possible. Once goals, metrics, and data sources are

identified, the Web site must be designed to collect and correlate data, extract information, and calculate metrics. When considering metrics and building user profiles from the visitors' perspective, it is imperative to consider the entire user experience at the Web site. In addition to personalization features, the user experience includes the tasks, services provided, navigation, design, and the overall value the visitor gains by visiting the site. To the extent that metrics gathered can be interpreted to enhance the user experience in these areas the more satisfied the visitor would be with the site, which will encourage future returns to the site.


The solution proposed by this section, takes into account the state-of-the-art critical review in personalization techniques conducted in this chapter as well as the evaluation methodologies e-metrics as they have been described above. Already existing personalization architectures should be extended to include an evaluation layer. For example, an extension of an information architecture-based framework for a personalization system by Instone (Instone, 2000) is provided in figure 3, including an extra layer for measuring a personalization initiative's success. In a similar manner, all personalization frameworks should be extended to accommodate the optimisation/self-adaptivity stage. Otherwise, the personalization initiatives' success cannot be improved. To achieve optimised results, the personalization techniques are fed with the results of the evaluation metrics, and, thus, any personalization architecture and/or framework has to reflect the plan-do-check-act cycle explained previously. The feedback input to the personalization techniques is an open research issue that we are currently working on.

Within the context of personalization, attributes and attribute values provide the "glue" which links together the users and the content and forms the personalized user interface. Attributes of the content are matched up with attributes of users. Specific attribute values about a user are paired with content meta-information to determine which content to display and how to present it at any given time. In this framework, we have users and the content meeting at the user interface through the process of personalization.

In more detail:

- Users: Users have profiles that represent their interests and behaviours. Specific values for a profile are determined by the set of defined attributes and the possible values for each attribute.
- Content: Likewise, content is profiled, based on a set of attributes and assigned specific values.
- Underneath the user interface is the profile layer, where specific values for the attributes are used to determine what content to present to which user under what conditions. A user's profile exists here and can be changed explicitly by user actions (such as filling out a form that requests particular profile information), or implicitly by certain actions (such as buying certain products).

Figure 3: An Information architecture – based framework for evaluative personalization systems (Adapted from Instone)



Likewise, a profile of the content exists and is matched with user profiles through a set of rules.

- Beneath the profile layer are the vocabularies which regulate the assignment of attribute values. At the vocabulary layer, the attributes themselves are defined and the set of acceptable values (preferred terms) are specified. The relationships between attributes are defined, such as child and parent attributes.
- The personalization rules are what leverage the profiles, attributes and values in order to make the personalized user experience. The most powerful rules operate on the set of attributes as a whole, at the “vocabulary” layer. When user and content profiles share the same attributes, then we can make rules that work for all values of those attributes.
- The evaluation layer filters the interactivity of the end user through its personalised interface and collects results based on the predefined e-metrics defined at design level and the feedback provided by the end user/customer. The data collected in conjunction with Web server log files are analysed using

Web Usage Mining techniques and the results obtained are filtered back into the personalization rules that were described above.

In general, this architecture defines a personalization system as any piece of software that applies business rules to profiles of users and content to provide a variable set of user interfaces. Nevertheless, one should always bear in mind that certain issues might affect the evaluation procedure such as the sequence of sessions used by a user to contact a site, the fact that eCRM should observe the individual sessions of a user as well as the whole lifecycle of the user and the web-structure usually affects the behaviour and navigation of a user.

FUTURE TRENDS

As the Internet matures and user interfaces change and become more sophisticated and complex, the role of personalization, the goals and the means for achieving these goals have to change, as well.

In order to enhance the personalization offered to the users, it is important to develop a way of measuring the success of the personalization approaches and to be able to use the results for improving its quality and the experience of the users. The evaluation of a personalization technique is carried out with the use of certain sets of metrics criteria. Then, in order to offer really adaptive and evolutionary web user experiences through dynamic personalization, the results of such an evaluation are fed into the personalization process. Technically, this means that a modification of the existing personalization algorithms and techniques is required, so that they can accept as another input the post-personalization evaluation results. Ideally, a learning capability will be embedded in the personalization techniques of the future, providing the users with more effective personalization.

A refinement and evolution of existing technologies for enabling and enhancing personalization, as well as introduction of new ones is certain to happen in the near future. The changes in the content, the user interface and the use of the Internet, as it matures, will require new technologies and techniques for personalizing the new content for the new uses and users and presenting it through the new interface. In addition, current technologies for supporting personalization (e.g. data mining, web-usage mining, machine learning, OLAP, data warehousing) will be refined and current personalization techniques are expected to evolve.

Another emerging trend for the future of personalization and certainly an important research opportunity is the application of personalization on alternative fields of IT, such as 3rd generation mobile telephony, call centres and PDAs. These are all rapidly evolving technologies, but more importantly they are being quickly adopted by the users. Furthermore, personalization techniques can be imported

from web sites for use in these technologies, but the reverse is also true; web site personalization and adaptability can benefit from research on personalization on these technologies. However, a major research challenge in this area will be to personalize combined user experiences from these technologies. For instance, when content is offered to a user that uses both a desktop computer in the office and a PDA when on the move, it should be examined how to personalize the content for maximum efficiency in this person's work, taking into account the capabilities, the limitations and the characteristics of each device and technology. To support this kind of device independent, cross channel personalise contact with their customers companies have to set up appropriate systems infrastructure.

The privacy of the personal data that are gathered through the personalization process is a very delicate research topic. A multidisciplinary effort will have to be carried out in order to ensure the confidentiality of the data. A characteristic example of the difficulty of preserving the privacy of the data is the ubiquitous span mail that most Internet users are familiar with. Law experts are already worried about the constant monitoring of what people are doing over the internet and have started investigating whether law should intervene (Volkh, 2000).

Making the application of personalization as seamless as possible is another emerging trend for personalization. The more experience users have with using the Internet, the easier it is for them to understand when they are being the subject of personalization. The web sites will tend to not let the users know that they are being offered personalized content by forcing them to answer questions regarding their preferences, but instead personalization will take place discreetly. The users should have the impression that a particular web site actually specializes in and offers the content that they are interested in.

As the Internet matures, the scope of personalization on web sites becomes broader. Currently, the main topics that personalization takes place successfully are specific "personalization-friendly" products (such as computer systems) and content (such as news services). However, there is a need for the provision of personalization in more complex products and services. There are times that the system has to "guess" the user's preferences and behaviour, when the user hasn't explicitly stated them or even doesn't know they exist. It is easy to offer a somewhat personalized experience to someone who has declared himself a sport fanatic, by offering sport news services and sport products. But what happens when we want to promote less straightforward products and services to someone, items that the users have a subjective way of picking, such as clothes and music? Or, even more, when we want to offer products that the customer hasn't used before? This can be very important from a marketing point of view because, if successful, it would mean the introduction of the user to a new market, with the particular web site and company as their first choice. That's why this can be considered a future

multidisciplinary research topic of paramount importance. Finding out the criteria that a user may employ for making subjective observations and decisions, such as what looks aesthetically correct in a web site, what kind of music is better or what kind of clothes are in fashion, can be an extremely difficult and complicated task. For achieving this, further and more sophisticated application of psychological/cognitive methodologies for discerning customer behaviour and clustering customers in behaviour domains is needed. In fact, it is expected that in the near future web sites will be initially designed by psychologists and cognitive scientists, instead of IT professionals. Then evolutionary and adaptive web sites will provide “End-users with the power to instruct their own computers and agents what to do, what they want and need and when they want and need” (Riecken, 2000).

CONCLUSIONS

The advent of Internet and especially of World Wide Web technologies has heralded the era of mass customisation. Mass customisation is the ability to offer tailor-made, individualized and personalized offerings at a massive scale, which is a key premise of adaptive Information Systems that need to evolve along with their users’ needs, preferences and requirements. Internet as a direct communication medium, coupled with World Wide Web’s capabilities of collecting detailed data at the granularity of individual mouse clicks, provide a tremendous opportunity for personalizing and enhancing the Web experience for users: presenting the information that is valuable to each individual user and optimising the interaction for maximum efficiency and effectiveness. Personalization enables and facilitates the creation of adaptive web sites, which is a prerequisite for the provision of high quality e-services.

Beyond e-commerce, the advances on Web personalization may provide useful insights to the problem of implementing adaptive systems. . Recently there has been an increasing amount of research activity on various aspects of the personalization problem. Most current approaches to personalization supported by various Web-based companies rely heavily on human participation to collect profile information about users. This suffers from the problems of the profile data being subjective, as well as getting out of date as user preferences change over time.

Personalizing computer systems to the needs of their individual users requires recognizing patterns in users’ behaviour. Central to this task of learning to personalize is the task of building a model of the user. Each user action takes place in the context of a specific task, and the main question is what information is available so that the process of predicting the future can be maximally informed by the past. The ultimate goal is to predict a user’s actions while exploiting as much information

as can be obtained with a view to proactively suggesting and/or facilitating those interactions that either the user intends to do, or values most if offered to him/her. Every method for predicting a user's future actions is based on some form of user profile or model that links information about the user or the task to expectations about the user's behaviour. The expectation is the more cues we have been linking our individual as well as collective experiences with the present conditions, the better chance we have of predicting the user's action. Developing techniques capable of using such a range of information sources is the next challenge to be met in building computer systems that effectively learn from and about their users.

In this chapter we have reviewed current and future technologies for personalization and identified key aspects of successfully delivering personalized e-services. To be able to run such an adaptive web site, a company should plan and execute personalization initiatives under the proposed unified, integrated approach. Our approach is a response to the open research issue of the personalization techniques' complementary; by integrating them across the customer decision process in a way that leverages the advantages of each group of techniques. The chapter also proposes the evaluation of a personalization initiative in order to assess its success. Thus, the personalization campaigns can be optimised; hence the evolutionary aspect of the adaptive web sites. The core issue of our ongoing research is the implementation of the unified, integrated approach together with the extension of personalization techniques to include a post-personalization evaluation stage.

Noteworthy future trends in personalization and, in general, web sites' adaptivity, include the application of psychological and cognitive methodologies for the enhancement of the personalization process, as well as personalization for different and new types of content and personalization in alternative fields of IT (e.g. mobile telephony). We are currently working towards a systematic definition of potentially useful metrics for evaluating the success of personalization efforts. Combining metrics results to the existing personalization techniques will form a new breed of evolutionary and adaptive information systems that can effectively support personalized web sites and personalized e-services.

REFERENCES

- Balabanovic, M. and Shoham, Y. (1997). FAB: Content-based collaborative recomender. *Communications of the ACM*, March Issue **40 (3)**, 66-72.
- Bamshad Mobasher, Robert Cooley and Jaideep Srivastava (2000), Automatic Personalization Based on Web Usage Mining, *Communications of the ACM*, March Issue **40 (3)**, 142-151.
- Choi, S.Y., Stahl, D.O. and Whinston, A. (1997). *The economics of e-commerce*. Macmillan Technical Publishing.

- Cohen, L. (1995). *Quality Function Deployment: how to make QFD work for you*. Addison-Wesley.
- Coldberg, D., Nichols, D., Oki, B. M., Terry, D. (1992). Using collaborative filtering to weave an information tapestry. *Communications of the ACM*, December Issue **35 (12)**, 61-70.
- Davidson, B.D. and Hirsh, H. (1998). Predicting sequences of user actions. *Predicting the future: AI approaches to Time-Series Problems. Tech. Report WS-98-07*, AAAI Press.
- Deitel, H. M., Deitel, P.J. and Steinbuhler, K. (2001). *e-Business and e-commerce for managers*. Prentice – Hall.
- Digital Enterprise: Available at: <http://digitalenterprise.org/metrics/metrics.html>
- Edwin, P.D. Pednault (2000). Representation is Everything. *Communications of the ACM*, August Issue **43 (8)**, 80-83.
- Gartner Group, (2000). *Web Site Personalization Strategies / Technologies*. Conference Presentation.
- Gartner Group, (2001). . Research Note, February 25, 2000
- Hammond, K.J., Burke, R. and Schmitt, K. (1996). *A case based approach to knowledge navigation. Case based reasoning experiences lessons and future directions*. D. B. Leake, ed. MIT Press, Cambridge, Mass., 125-136.
- Haym Hirsh, Chumki Basu and Brian D. Davison (2000), Learning to Personalize. *Communications of the ACM*, August Issue **43 (8)**, 102-106.
- Instone, K. (2000). *An Information Architecture – Based Framework for Personalization Systems*. Argus Associates
- Kalakota, R. and Robinson M. (1999). *e-Business: a roadmap for success*. Addison Wesley.
- Kolcz, A. and Alspector, J. (1999). An internet based newspaper filtering and personalization system. *Proceeding of the 22nd International ACM SIGIR Conference*. 332.
- Kotler, P. and Armstrong G. (1997). *Marketing - An introduction*. Prentice Hall.
- Maurice D. Mulvenna, Sarabjot S. Ananad and Alex G. Buchner (2000), Personalization on the Net using Web Mining. *Communications of the ACM*, August Issue **43 (8)**, 122-125.
- Padmanabhan, V.N. and Mogul, J.C. (1996). Using predictive perfecting to improve World Wide Web latency. *Comput. Commun. Rev.*, July Issues, **26(3)**, 22-36.
- Pine, B.J. (1993). *Mass Customisation: the new frontier in business competition*. Harvard Business School Press.
- Ramakrishnan, N. (2000). PIPE: Web Personalization by Partial Evaluation, *IEEE Internet Computing*, November – December Issue, 21-31.

- Rayport, J.F. and Jaworski, B.J. (2001). *e-Commerce*. McGraw-Hill Higher Education.
- Schlimmer, J.C. and Hermens, L.A. (1993). Software agents: Completing patterns and constructing user interfaces. *Artificial Intelligence Research* **1**, 61-89.
- Schonberg, E., Cofino, T., Hoch, R., Podlaseck, M. and Spraragen, S.L. (2000). Measuring Success, *Communications of the ACM*, March Issue **40 (3)**, 53-57.
- Smyth, B. and Cotter, P. (2000), A personalised television listing services. *Communications of the ACM*, August Issue **43 (8)**, 107-111.
- Spilopoulou, M. (2000), Web Usage Mining for Web Site Evaluation, *Communications of the ACM*, March Issue **40 (3)**, 127-134.
- Turban, E., McLean, E.R. and Wetherbe, J.C. (2001). *Information Technology for Management Making Connections for Strategic Advantage: Making Connections for Strategic Advantage*. J. Wiley.
- Vassos, J. (1996). *Strategic Internet Marketing*. Indianapolis, Que Publishing Co.
- Wigard, R. T. (1997). Electronic commerce: definition, theory and context. *The Information Society* **13(1)**, 1-16.

PART IV:

**EVOLUTION
AND
ADAPTATION
IN BUSINESS
AND
EDUCATION SYSTEMS**