Personalization Overview

Objectives

This document covers the following topics:

- Personalization as a concept
- Uses of personalization
- Different types of personalization
- LikeMinds personalization and how it differs from other approaches
- LikeMinds' general capabilities

Key Terms

Personalization

Personalization is the process of deciding - given a large set of possible choices - what has the highest value to an individual. These choices can range from a customized home page look and feel to product recommendations, from banner advertisements to news content. As we will see, personalization is a very broad term and must be further qualified.

User-profile based personalization

This is the process of making decisions based upon predefined group membership or stored user profile information. For example, on a portal site, users may choose to see particular news feeds, select certain categories of sports scores, and be shown ads that match their user profile.

Rules-based personalization

This approach makes decisions based on pre-defined business rules as they apply to a segmentation of users. For example, a computer hardware maker may have different levels of technical support membership. Their website will then expose different support content to the user based on which level of support they have purchased. So if user ‘X’ is in group ‘Y’, we will show them content from content group ‘P’.

Collaborative filtering

This is a personalization approach that is based upon mathematical analysis of collected user preference or activity data. This analysis leads to predictions of what a given user will like or will be likely to do next. For example, in the case of a DVD direct sales site, recommendations might be made based on how you and users similar to you rated movies.

Mentors

A collaborative filtering (CF) term: These are other users in a CF database who have expressed preferences. Those users with compatible tastes or "like-mindedness" are used to generate recommendations. Mentor recommendations are drawn from a mentor "pool" based on most-qualified users. The qualification to become a mentor depends on the quantity of information (in CF terms, user transactions) which is recorded about the user.

Ratings

These are explicit user opinions expressed via a web application interface. A user is shown a list of products and asked to assign each product a value, such as "favorite," "liked," "did not like" and so on. LikeMinds uses these to generate recommendations for other users via collaborative filtering.

Clickstream

A data gathering method that tracks content-related user behavior -- detailed product views, shopping cart inserts, product purchases and the like -- and interprets this behavior as a preference by assigning it a rating-like weight.

Event weight

A clickstream input term, event weight refers to the relative importance assigned to customer behaviors. Behaviors could be product searches, detailed product views, or purchases. For example, in a direct-sales situation, a product search might count be weighted as a “1”, while a detailed product view might be weighted as a “3”, hence the activity search is weighted more heavily.

Market Basket Analysis (MBA)

A recommendation method that uses historical purchase data to identify purchasing patterns and trends.

Product matching

Also a recommendation method, product matching relies on fixed rules--for example, "granola goes with yogurt,” or “all movies directed by Martin Scorsese go together” to offer recommendations. Also a rarely used LikeMinds engine of that name.

LikeMinds [Anonymous] User Profile

LikeMinds stores a user profile in it’s database for each persistent profile made. This profile is anonymous and does not neccessarily contain any demographic or other such direct marketing data. A LikeMinds profile consists of a user ID plus some additional activity or preference data. This data is largely only useful for the LikeMinds software.

Introduction

This document provides a general description of personalization, its implementation, and the types of personalization offered by LikeMinds.

Personalization has become a key concern when attracting customers and working to maintaining their loyalty. As web content matures and multiplies, sites need more and more sophisticated methods of delivering relevant content to their users. Personalization plays a crucial role in delivering user-centric content and experiences.

Through the use of personalization technology, both e-businesses and content-based sites can offer customers a user-specific experience mimicking and at times exceeding such potential in the offline world.

What is personalization?

Personalization is the process of deciding - given a large set of possible choices - what has the highest value to an individual. This adds both utility and warmth to a web application, as users find what they seek faster and feel “recognized” by a site.

On a more practical level Personalization is an overloaded term: There are many mechanisms and approaches (both automated and marketing rules controlled) whereby content can be focused to an audience in a one to one manner. This section delineates between the various approaches providing the student with a terminology to describe each approach in isolation. Furthermore it describes how the approaches can be combined such as using LikeMinds “engines” to prioritize results from a rules-based recommendation or filtering LikeMinds recommendations using business rules.

Note: as indicated above, this collected information need not be tied to the user's actual identity, only to the user's online presence represented by a unique user ID.

What can be personalized?

Personalized content may be advertising, recommended items, screen layout, menus, news articles, or anything else accessed via a web page or software application.

Business benefits

Personalization contributes to a variety of e-business goals: increasing site usability, replicating offline experience, converting browsers to buyers, retaining current customers, re-engaging customers, and penetrating new markets.

Increase site usability

By limiting navigation options, and providing direct links to desired content, personalization automatically makes a site more navigable, allowing users to find desired information, products, and services more quickly.

Replicate offline experience

Replicating familiar offline experiences is a key goal/ benefit of personalization. Ideally, personalization acts as a stand-in for the Friendly Store Clerk, the person behind the counter at the corner hardware store who remembers you, suggests purchases, and helps you solve your particular problems.

Conversion (increased sales)

Research shows that converting browsers to buyers has a significant impact on site revenues. Toward this end, personalization brings targeted, high-value purchase opportunities directly to the user. By positioning desired content in front of a user, personalization increases the odds that a browser will become a buyer.

Retention

As the Internet matures, and success becomes measured in more than strict traffic numbers, retaining customers is crucial for any site's success. Personalization enhances site "stickiness," that is, an increased likelihood that customers will bookmark and return to your site. Customers return more frequently to sites where they receive specific benefits, and personalization provides these benefits.

Re-engagement

Often times, a customer will shop or consume information from a number of sites on the Internet. Re-engagement is the process of reaching back out to a customer via email or other means to let them know you have something that they may be interested in. If such notifications are personalized, the customer will learn to trust re-engagement attempts and they will more likely be successful.

Penetration

Usable sites inevitably attract more users, as word-of-mouth spreads about the site's utility.

Measuring success

Although determining the success of any particular web application feature can be quite challenging, personalization should lead to measurable increases in sales, items per sale, site traffic, and conversions. Ideally, site tracking tools, such as IBM's Site Analyzer, can be used to develop this data.

It is a good idea to set measurable goals for what changes personalization will effect on your site, e.g., “convert x% of browsers to buyers,” then measure for these goals before and after implementation of personalization.

Privacy issues

In any personalization scenario, privacy must be balanced against increased user benefit. Tracking users’ buying habits, of course, requires collecting some fairly sensitive information. As a counter to these privacy issues, though, personalization provides benefits, in terms of increased usability and less unwanted content.

LikeMinds provides personalization without imposing on user privacy. It requires no data other than a user ID and password in order to work. Users can effectively be anonymous and still benefit from personalization.

Types of Personalization

We've established the benefits of personalization–it helps your site attract, keep, and re-engage customers and increase the quality and quantity of customer sales. In fact, personalization is one of the primary advantages provided by Internet-enabled businesses.

Let's now examine the different ways in which personalization is implemented. We start with the most basic and proceed to the more complex.

Types of personalization

Web-based personalization can be implemented in a variety of ways, such as:

- User-profile based
- Rules-based
- Market-basket analysis
- Collaborative filtering

User-profile based

These allow users to determine the content on a personalized page, e.g. "sports," "news," "entertainment," and so on. "My.___.com" sites generally employ this method. Personalized content of this sort provides a great user benefit; users can customize their browser home page with their most frequently-accessed content, for example. However, since these features require users to actively configure options, users tend to take advantage of them in only limited ways.

User-profile based personalization is generally implemented with an application server such as IBM WebSphere.

Figure 2-1 Example of user profile personalized page.

Advantages

Users can select exactly what they want to see. they control their own profile and the information in that profile.

Disadvantages

This process is extremely manual. users may often miss new content if profile gets stale etc.

Rule based

Another form of personalization implements rules based on a customer's demographics, past purchases, or product attributes.

These use an if-then process to decide what content to present. Rules can be based on promotional logic, item dependencies, demographic analysis, or on simple item-to-item pairings (computers and printers, for example).

In the IBM WebSphere environment, the Personalization Rules Engine provides rule implementation. See product documentation for more details.

For example, a drugstore site may know, through research, that males aged 24-50 (note not a particular male, myself for example but a class of males) tend to purchase a certain type of razor. The site can then recommend the razor based on information about the user's age.

Advantages

- This method leaves marketers in control. With rule-based recommendations, marketers determine content recommendations by implementing site-wide rules. This allows particular content to be segmented by audience, useful when your users fall into broad categories, such as support level or geographic region.
- This method can be implemented consistently with other marketing programs. In other words, if a marketing campaign targets a particular customer demographic, the site can recommend campaign products directly to users in that demographic.
- Since it requires little user information, rule-based recommendations can be used in conjunction with other personalization solutions, as a "fall back" method.
- Rules-based personalization works effectively in situations where particular marketing knowledge must be leveraged, where a "blind," algorithm-based method will not work (for example enforcing political sanctions)
- This method is great when you know your prospective customer base really well and it is relatively stable..

Disadvantages

- Rules-based personalization involves an extensive, research-intensive setup, which can be difficult to manage, particularly with a large inventory. Essentially, new rules must be added for each new product.
- This involves complex rule maintenance. Effectively, you must change rules to account for every user situation, or risk failing to meet user expectations and decreasing your site's usefulness.

Market-basket Analysis

Also mathematically generated, market-basket analysis provides "item to item" recommendations, based on items purchased together by past users. Since item-to-item associations involve less knowledge of the user than item-to-user affinities this personalization mechanism can have an almost instantaneous response.

This has 'historically' identified some interesting combinations, such as the "diapers and beer" example, where analysis of traditional shopping carts unearthed the surprising truth that diapers and beer are often purchased together.

Advantages

- Market basket analysis improves on rule-based recommendations, in that identifies some less obvious combinations of items.
- Since it can be used with less data, but doesn't require rules configuration, MBA acts as an effective fall back method for other recommendation methods. Because personalization can be provided "first visit" this method of personalization is very popular in sites with less loyal repeat visitors and anonymous browsers.

Disadvantages

- Market-basket analysis assumes that people tend to think alike–that is, that since user x bought two items in combination, user y will also desire these products in combination.
- Requires a base of transactional data from which to work. Since most sites track what users place in their “shopping basket,” this generally does not pose a problem.
- In some situations, the concept of a user “basket” is not immediately apparent, and must be developed programatically. (Help desk situation, for example.)

Collaborative filtering (CF)

This method involves gathering data on user preferences and behavior, and then using that data to algorithmically produce recommendations for new users. These similar users are often known as “mentors” or “neighbors.”

This method, known as collaborative filtering, solves the problem of personalizing content for an essentially unknown user, avoids the problem of grouping users in categories to which they really don’t belong, and allows for unique combinations of tastes in one individual (for example, sailing and yoga).

This is often referred to as “community-based” recommendations. It mimics the word-of-mouth product and service recommendations we receive daily from friends and associates: people like us, but not exactly like us.

Collaborative filtering allows the creation of “markets of one,” that is, individuals to whom items are custom-targeted. In fact, this process occurs naturally with this method: customers are automatically recommended products that suit them and them alone, because collaborative draws recommendations from a variety of other users with similar, but not exactly corresponding tastes.

Furthermore, collaborative filtering offers levels of recommendations; aside from simply telling users that they'd like products, it indicates how much they might like them.

Inevitably, this provides the user with new ideas that don't radically differ from the old ideas, increasing the likelihood that the user will find the new ideas good ones.

Advantages

- Since collaborative filtering does not attempt to make broad judgments about groups of people, it recommends products more likely to specifically suit a user's wants and needs.
- Collaborative filtering can tell your users not simply that they'll like an item, but how much they'll like it, e.g. "we predict that you'll give this CD a four out of five."
- Once a CF system is put in place, the system itself selects items to "push." This removes mystery from the process of deciding which products to market.
- Once the system is in place, it automatically generates new recommendations, without any additional input besides a continual stream of new users.

Disadvantages

- Because CF requires a body of data before it can make recommendations, it has difficulty making recommendations for new items. This is one half of the problem with "cold-starting" collaborative filtering on a particular site.
- Similarly, CF poorly suits an environment with rapid inventory churn; it depends on items being currently stocked in order to rate them.

- New users represent the other half of the "cold-starting" problem; since with new users, LikeMinds doesn't know anything about their preferences, it cannot make new recommendations for these users.

Using methods in combination

Given the complexity of the business models and user bases of most e-business sites, generally sites will want to employ a combination of personalization methods.

For example, the section below details how rules-based and collaborative filtering personalization work together.

Using rules-based personalization with collaborative filtering

Since these represent two of the major methods of personalization available, let's compare the two:

When rules work

- With a limited number of items.
- When price points are high or purchasing frequency is low.
- When a dependency exists between items, such as in the case of homeowner's insurance and disability insurance.
- A hierarchical promotional structure is at work, as in the case of an affinity credit card or customer loyalty program. If a customer flies Bob's Airline, we always want to recommend a credit card associated with Bob's Airline.

Why

- Algorithmically-generated recommendations work poorly in finite, limited arenas.
- Promotional or dependency logic is difficult to integrate into these generated recommendations. For example, recommending disability insurance without homeowner's insurance is wrong, even if "other users like you" have chosen disability insurance.

When collaborative filtering supplements rules

- When the number of items offered and users purchasing are high; this requires many, many rules to be created.
- When price points are low, items are all quite dissimilar, or the products offered have a wide range of user appeal.
- New products or new users are commonplace.

Why

- Cannot write rules covering all items
- A wide variance within the purchasing space enhances the similarity approach, as well as having lower risk for making "bad" recommendations.
- Rules require "specific history" to function well. LikeMinds requires only that any history be available and can use its various similarity metrics to find affinities between either users or items.

Summary

The personalization method you choose depends on the demands and constraints of your application, but also on the particular character of your inventory.

The next section describes LikeMinds' specific approach to personalization.

LikeMinds Personalization

Overview

This section describes how LikeMinds implements personalization, beginning with collaborative filtering and proceeding through market basket analysis.

Collaborative filtering

LikeMinds provides collaborative filtering through its Preference and Clickstream engines.

Method

The way collaborative filtering works is relatively simple:

Table 2-1 How collaborative filtering works

Step	Description
1	A Web site or application maintains and populates a database of implicit and explicit user preferences or their “significant” navigational and/or purchase behaviors.
2	Using this data, LikeMinds then creates a “user profile” based on that user’s choices and behaviors, then dynamically groups each individual user with a group of other users, called mentors. Mentors form a specific community created uniquely for every user in the database.

Step	Description
3	Analysis algorithms are then used to weigh the relationships between, say, User X and every member of User X's community. Mentors weighted more heavily have more in common with User X, and thus their recommendations count more heavily.
4	Drawing from this mentor pool, LikeMinds identifies items not yet rated or seen by User X, but which seem desirable to the user based on mentor preferences/behaviors.
5	In other words, the mentors in User X's specific community will recommend items for that user that they have interacted with but the targeted user has not. Thus, we have the expression "community based recommendations."
6	As User X expresses more preferences, LikeMinds can produce more exact recommendations by more effectively matching the user with mentors.

The assumption with collaborative filtering is that User X will like or be interested in interacting with the products or items that her mentor community has liked or has positively interacted with. This assumption has largely been proven to be valid.

Mentors

Mentors are other users who have expressed preferences similar to the targeted user. To select these users, LikeMinds first identifies best-qualified mentors and groups them in a **mentor pool**. Then, when recommendations are needed for a particular user, LikeMinds selects mentors based on their similarity to the targeted user. Mentors figure in both explicit and implicit recommendations.

Mentors provide the raw material for recommendations, through their own preferences and behaviors. For example, if User X rated three movies highly, and another user rated the same three movies highly and also rated a fourth movie, unseen by User X, highly, it's a good bet that User X will like the fourth movie as well.

This is, of course, a simplified way of describing the LikeMinds process; in reality, preferences from multiple users are taken into account.

The diagram below visually represents mentors at work. The closer a mentor is to the target individual in tastes, the more heavily that mentor's recommendations are weighted.

Figure 2-2 Mentor Pool Selection

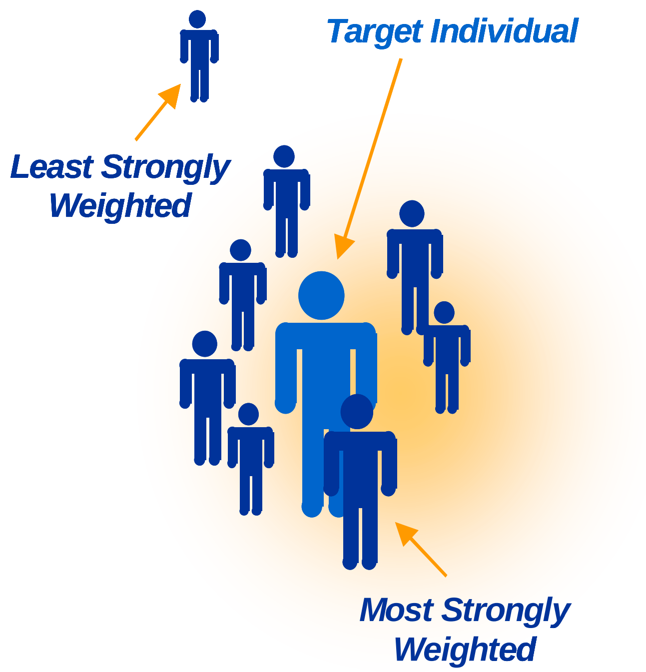

How do they generate recommendations?

Again, LikeMinds draws recommendations from mentor-expressed preferences.

Brief selection overview

Mentors are selected from a "mentor pool" of most suitable users. From this mentor pool, mentors are selected based on their similarity with the targeted user.

Examples of use

Collaborative filtering can be used to produce meaningful product recommendations in a variety of e-business situations, such as: a customized home page, an explicit recommendation page, an “additional suggestions” section added to a shopping basket, or content recommendation page.

For example, say a clothing retailer wants to implement a site section that recommends clothing. It could ask users to rate different clothing styles the first time they log onto the site. Then, the LikeMinds engine aggregates these ratings and uses them to create user profiles describing each user’s tastes.

Based on these user profiles, LikeMinds creates mentor relationships for each user.

Differentiation from other personalization solutions

Let’s explore the difference between collaborative filtering-based recommendations and other forms in more detail.

User based model

Users make explicit or implicit choices, which are collected into a database. This data then become the basis of

- The user’s profile, which allows recommendations to be targeted to the user.
- Recommendations for other users, as the user enters the “mentor pool.”

Dynamic

With LikeMinds in operation, new user choices, and new users, automatically change the recommendations made by the system as a whole.

Sensitive

Based on a user's profile, even small system changes can produce new recommendations.

User-based marketing

With this model, taste is not dictated from above. Recommendations are user-based and dynamically generated: no one needs to make rules for every user-product interaction. This marketing also occurs instantaneously, without extensive planning.

Types of collaborative filtering

Ratings

Ratings are an explicit means for users to express their preferences; the site must actively ask users how they feel about particular items.

In this model, "you are who you say you are," that is, the recommendations returned will reflect your explicit choices.

Clickstream

With this method, a site collects data based on user behaviors such as searching for a product, viewing a product in detail, or purchasing a product.

In this model, "you are what you do,' in other words, your profile is based on your actual behavior in relation to an item.

Ratings

Definition

Ratings are expressions of product value. An individual rates products through a web application, and the results are used to create recommendations for the user and others.

Implementation

With explicit ratings, users are shown items and asked to assign them a value-based rating, e.g. “excellent,” “very good,” etc. This method often works well with sites that have a large inventory of one product type, such as movies, music or books. With the right user interface, sites can constantly request ratings from users, building a set of preferences that allow for more and more sophisticated recommendations.

Users must rate a variety of items in order for LikeMinds to make good recommendations--this is known as “mentor coverage.”

Advantages

Ratings provide an exact measure of user preference, and are easy to measure against other users’ ratings.

Also, this allows users to participate in a community, in letting others know what they think.

Disadvantages

Requiring users to rate items before they can receive recommendations. This requires explicit collaboration from the user.

Clickstream

Definition

This refers to the process of tracking user preferences via common e-business events weighted so that they act as ratings. These events, common to most e-business sites, represent those that best evidence user choice, such as detailed product view and product purchase.

Implementation

In order to collect data via clickstream, you must integrate LikeMinds with your application server, which provides the business logic for such events as detailed product views and shopping cart ads.

Advantages

Clickstream collects preference data without interrupting a user's experience.

Since clickstream tracks actual user events, these may better indicate what future purchases a user will make.

Disadvantages

Since Clickstream relies on what users "do," not what they say they do, it is possible that recommendations may not match users perceived desires.

Clickstream requires more planning. Tuning the site is critical, to be sure that events are assigned and weighted to produce optimal recommendations.

Event weight

Each predefined event is given a weight, in order to develop it into a ratings system. For example, a product view may carry weight “1,” a detailed product view weight “3,” and a product purchase weight “5.”

Market Basket Analysis

Definition

This method uses analysis of user “shopping baskets” that makes recommendations based on past purchase combinations. This analysis yields common combinations, such as olive oil and vinegar, but also non-obvious ones, such as the “diapers and beer” example.

How it works

To work, market basket analysis tracks user purchases across time, noting what items are bought in combination. Item to item relationship strength is calculated by counting the number of times every item occurs with every other item (in pairs, triplets, quadruplets, etc.) in the same shopping cart or user session.

Market Basket Analysis can use existing transactional data--often stored by sites for marketing purposes.

Since these recommendations depend entirely on past transactions to work, they do not require any user data. MBA is thus very effective at recommending items to new users, since it need not know anything about these users except the product(s) they’ve already chosen.

Table 2-2 How market basket analysis works

Step	Description
1	Site gathers data on user transactions. (Often, this has been previously stored for marketing purposes.)
2	Engine sifts through this data, identifying items purchased or viewed in groups.
3	When user purchases or views an item, engine recommends item grouped with chosen item in the past.

LikeMinds implementation

LikeMinds' specific method of MBA allows for a more flexible definition of "shopping cart," including, for example, detailed product views. LikeMinds also allows for tuning the strength of a recommendation.

How it works with collaborative filtering

Since, as noted above, MBA makes rather general assumptions about customers' buying habits, it is generally preferable to use collaborative filtering once you've built a mature database with lots of data from which to draw collaborative-filtering-based recommendations.

In this scenario, market basket analysis provides an effective "backup mechanism" for cases where collaborative filtering won't work, e.g. as with a new user or new product. Ideally, these methods can be combined to create a complete personalization solution.

With LikeMinds, you can tune your application to use market basket analysis when recommendation data falls below a certain level, or when few mentors are available.

Summary

This session has covered the following:

- Personalization is the process of delivering content targeted to a particular individual.
- A variety of means exist to implement personalization, including user profiles, rules, collaborative filtering and market basket analysis.
- Each of these methods offers particular benefits, and their use should be determined by the specifics of a personalization situation.
- The LikeMinds system works in conjunction with application servers, such as IBM WebSphere, to provide collaborative filtering, market basket analysis, and product matching.
- LikeMinds' collaborative filtering allows you to collect data via either explicit user choices (ratings) or implicit user choices (Clickstream). It then generates recommendations by algorithmically calculating best-fit recommendations from other user choices.
- Market basket analysis gathers data on user "baskets" and then uses basket groupings to produce recommendations.
- Product matching produces recommendations based on rules relating products by specific attribute