Ubiquitous Recommendation Systems

David W. McDonald, University of Washington

they present, these systems refine the wealth of information, video feeds, online documents, and applications available through the pervasive infrastructure.

USER MODELS

Nevertheless, such systems face many difficult challenges before they can fulfill their supporting role in popular visions of ubiquitous computing. Constructing accurate user models and putting those models to proper use are among the most important challenges to developing next-generation recommendation systems.

In many popular visions of ubiquitous computing, the environment proactively responds to individuals who inhabit the space. For example, a display magically presents a personalized advertisement, the most relevant video feed, or the desired page from a secret government document. Such capability requires more than an abundance of networked displays, devices, and sensors; it relies implicitly on recommendation systems that either directly serve the end user or provide critical services to some other application.

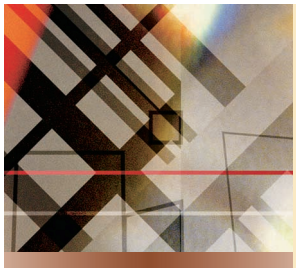

Recommendation systems can help people cope with information overload in ubiquitous computing environments.

NARROWING CHOICES

In general, recommendation systems manage information overload by helping a user choose from among an overwhelming number of possibilities. These systems broadly fall into three classes based on the techniques they use to narrow the range of likely choices:

- *Content-based filtering* systems utilize machine-learning techniques such as naïve Bayes to analyze Web pages, Usenet news, e-mail, and other types of electronic content amenable to automatic textual analysis. For example, such a system might compare words in an online film review with terms that characterize your movie-watching preferences to determine whether you are likely to enjoy that film.
- *Collaborative filtering* systems ignore descriptions and instead focus on ratings of items by multiple users. For example, if a movie you have not seen is highly rated by others who share your taste in films, such a system might recommend that movie to you.
- *Link-based* systems discover relations among items and then use graph-theoretic algorithms to find items in a set that are either the most exemplary or most referred to by other set members. Such approaches work well on hypertext and have been used to discover social networks. A link-based technique facilitates Google's good search results.

Recommendation systems mediate the user experience in the digital world, and they will be increasingly helpful in performing the same role in the physical world, thereby filling an important gap in ubiquitous computing. By tailoring the environment or information

Filtering systems rely heavily on user models. Collaborative filtering systems cluster individual profiles to find highly rated items to recommend, which requires models to share the same structure or otherwise be compatible with one another. Most current systems rely on an identical model structure for each user—for example, ratings of known items on a specific scale. However, user models should represent different aspects of people's day-to-day experience, not just their likes or dislikes; this may include implicit as well as explicit ratings.

Content filtering systems follow a predictive model that makes recommendations based on a user's past preferences. This type of model builds personal profiles and uses them to rate items. For example, if you often rent films starring Clint Eastwood, such a system might recommend *High Plains Drifter*; likewise, an e-commerce pharmacy might recommend throat lozenges because you just added a cold medication to the shopping cart.

Content filtering is also useful for managing unwanted e-mail and may help curtail the Short Message Service spam that is starting to appear on cell phones. One of the future horrors depicted in the sci-fi thriller *Minority Report* is the ubiquitous presence of pop-up ads, seemingly with no filters to control content. In instances where explicit computer input is generally not part of user activity, researchers are seeking answers to how content filtering systems might learn preferences and develop a personal profile.

Collaborative filtering systems make "birds-of-a-feather" assumptions that people who liked certain things in the past will like similar things in the future. This model incorporates people's tastes without relying on content, which in some cases, such as with music or movies, might be hard to compare. For example, when a person views a DVD on Amazon.com, the Web site lists other film titles purchased by people who bought the DVD without analyzing the movie itself.

Collaborative filtering has the potential to learn from many people and arrive at appropriate recommendations without requiring the system to construct a complete profile for each user. The challenge that researchers face is how to maintain privacy of all users as such systems become ubiquitous.

THE FUTURE

These and other possible models for recommendation systems may be adequate for current uses, but perhaps not for future ubiquitous computing situations. Researchers are exploring a wide range of applications such as intelligent tourist and restaurant guides, navigation aids, and bricks-and-mortar shopping systems that make recommendations based upon user activities and behavior patterns. For example, a car navigation system might recommend routes based on the driver's travel patterns, such as a tendency to avoid high-speed roadways.

In addition, it would be useful if Web sites could compute access patterns and provide users with location-dependent content recommendations. Ubiquitous computing systems with knowledge of more than locations—say, the tools a person is using—could greatly benefit that person by recommending others who have expertise with those tools.

Farther afield, imagine if a device such as an LCD projector could learn "best known methods" and supply this information to users—for example, "When other people hooked their IBM ThinkPad to this projector, they selected low resolution." Do the models developed for recommendation systems on the desktop computer also apply to new devices that are not typically regarded as having computing or communication capabilities?

Unlike current models, future models must also account for how people associate or interact with one another. These models should be based on social science rather than popular ideals—the world is not like *Cheers*, a place where everybody knows your name and says hello when you walk in the door. Serendipitous, meaningful interactions in public social spaces are quite rare in some cultures. How many times have you had a substantive conversation with a stranger in line at a grocery store? Good models of social spaces would make it possible to predict, for example, whether two people who frequent the same café are compatible.

As recommendation systems evolve to exploit new advances in ubiquitous computing technology, researchers and practitioners from technical and social science disciplines must collaborate to address the challenges to their effective implementation. Although it may be impossible to perfectly anticipate each individual's needs at any place or time, ubiquitous computing will enable such systems to help people cope with an expanding array of choices. ■

David W. McDonald is an assistant professor in the University of Washington's Information School. Contact him at dwmc@u.washington.edu.

Editor: Bill N. Schilit, Intel Research Seattle; bill.schilit@intel.com