2. H. Werthner et al., “Information Society Technologies for Tourism,” *Report of the Strategic Advisory Group on the 5th Framework Program on Information Society Applications for Transport and Associated Services*, 1997.

3. H. Werthner, “Just Business—Shouldn’t We Have Some Fun?” *Proc. 3rd Int’l Conf. Electronic Commerce & Web Technologies* (Proc. ECWEB-DEXA), Springer-Verlag, New York, 2001.

Hannes Werthner is head of the eCommerce and Tourism Research Lab at the IRST Research Center, a professor at the University of Trento, and founder of the eCommerce Competence Center in Vienna. His research activities cover decision support systems, simulation, artificial intelligence, and Internet-based information systems, especially in the field of tourism. He earned an MS and PhD in computer science from the Technical University Vienna. He is a member of the strategic advisory board for the European research program IST, acts as the editor in chief of the journal *Information Technology and Tourism*, and is Honorary President of the International Federation for IT and Travel/Tourism (IFITT). Contact him at werthner@itc.it.

Travel Recommender Systems

Francesco Ricci, *eCommerce and Tourism Research Laboratory*

Recommender systems are commonly defined as applications that e-commerce sites exploit to suggest products and provide consumers with information to facilitate their decision-making processes.[1] They implicitly assume that we can map user needs and constraints, through appropriate recommendation algorithms, and convert them into product selections using knowledge compiled into the intelligent recommender. Knowledge is extracted from either domain experts (content- or knowledge-based approaches) or extensive logs of previous purchases (collaborative-based approaches). Furthermore, the interaction process, which turns needs into products, is presented to the user with a rationale that depends on the underlying recommendation technology and algorithms. For example, if the system funnels the behavior of other users in the recommendation, it explicitly shows reviews of the selected products or quotes from a similar user.

Recommender systems are now a popular research area[2] and are increasingly used by e-commerce sites.[1] For travel and tourism,[3] the two most successful recommender system technologies (see Figure 1) are Triplehop’s TripMatcher (used by www.ski-europe.com, among others) and VacationCoach’s expert advice platform, Me-Print (used by travelocity.com).

Both of these recommender systems try to mimic the interactivity observed in traditional counselling sessions with travel agents when users search for advice on a possible holiday destination. From a technical viewpoint, they primarily use a content-based approach, in which the user expresses needs, benefits, and constraints using the offered language (attributes). The system then matches the user preferences with items in a catalog of destinations (described with the same language). VacationCoach exploits user profiling by explicitly asking the user to classify himself or herself in one profile (for example, as a “culture creature,” “beach bum,” or “trail trekker”), which induces implicit needs that the user doesn’t provide. The user can even input precise profile information by completing the appropriate form.

TripleHop’s matching engine uses a more sophisticated approach to reduce user input. It guesses importance of attributes that the user does not explicitly mention. It then combines statistics on past user queries with a prediction computed as a weighted average of importance assigned by similar users.[4]

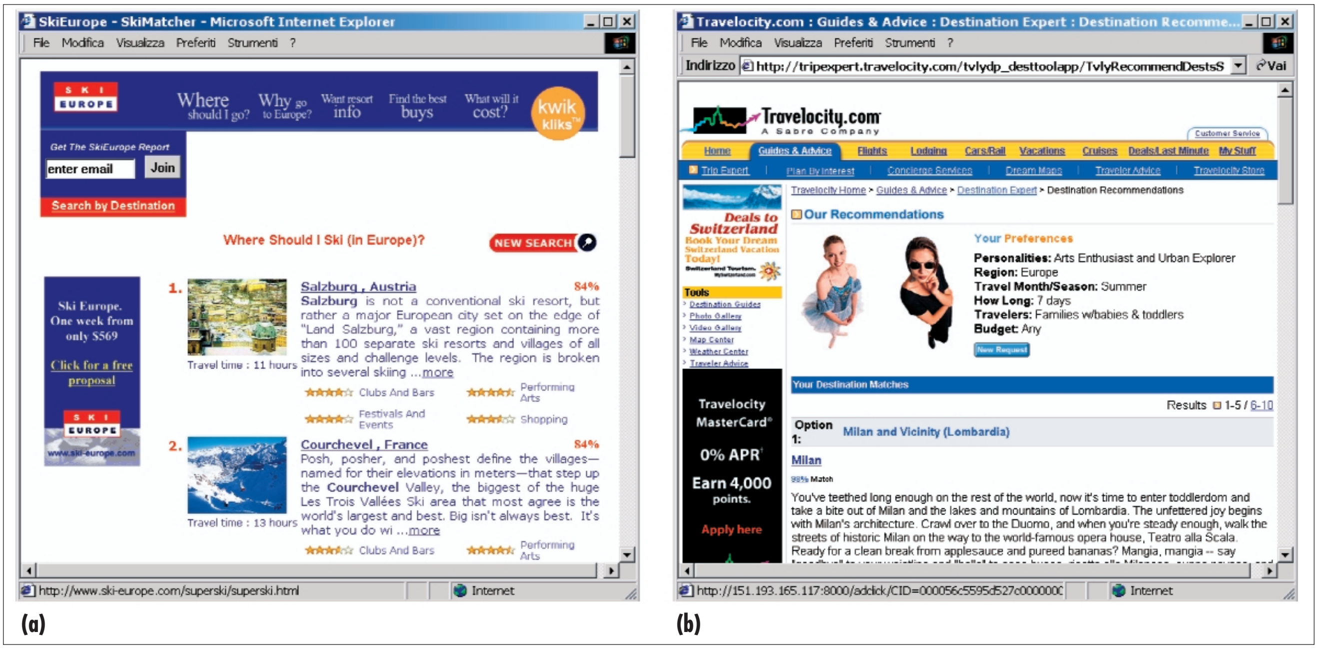

Figure 1. (a) Ski-Europe and (b) Travelocity destination recommendation tools.

Caveats and limitations

Neither system supports the user in building a "user defined" trip, consisting of one or more locations to visit, accommodations, and plans to visit additional attractions (a museum, the theater, and so forth). Although travel planning is a complex decision process, these systems support only the first stage—deciding the destination.

Researchers have proposed several *choice models*,[5] which identify two groups of factors that influence destination choice: personal features and travel features. The first group contains both socioeconomic factors (such as age, education, and income) and psychological and cognitive ones (experience, personality, involvement, and so forth). The second group might list travel purpose, travel-party size, length of travel, distance, and transportation mode. These various factors affect all stages of the traveller's decision-making process, which is a complex constructive activity.

Another reason why these systems focus on destination selection relates to the filtering (content-based) approach. Even if we could apply the same filtering technology to other tourism objects, such as cruises, the system would have to describe a catalog of cruises—that is, build a catalog using a selected set of features (decision variables). The approach does not scale unless we pursue a costly knowledge-engineering activity for each product type. So, these systems must have a particular catalog—in this case, a catalog of destinations—which requires extensive domain knowledge and must be built for the particular application. Currently, the focus is on destinations because they are rather stable, reusable concepts (many recommender systems can exploit the same destinations knowledge base).

Pure *collaborative filtering* approaches do not suffer from this problem, but, unfortunately, we cannot readily implement them in the travel domain. The major issue is the complexity of travel objects; we can't simplify a trip to the point where two travellers' trips are the same. Surely two people have bought the same book, but it is less likely that two people have experienced the same trip. This points to a basic requirement of CF approaches: one user's purchase history must be comparable to that of another. Thus, one user's travel list must somehow overlap that of another user. One approach could be to simplify the travel description to a certain point—for instance, representing just the destination—but then we will discover that the already visited destinations are insufficient to predict the next one. Additional context information must be included, so we must query the user about the content of his or her trip. Hybrid approaches that combine content- and collaborative-based approaches will more likely succeed.[6]

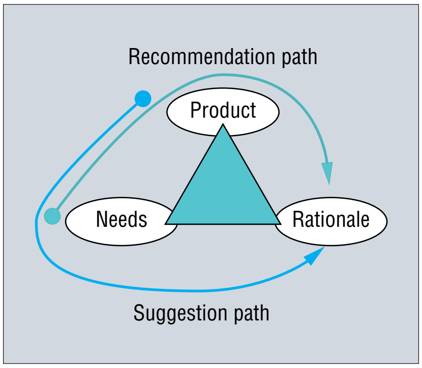

Figure 2. Recommendation and suggestion paths.

Broadening the scope

Going back to the basic recommendation process (moving from needs to products with explanations), this apparently linear process is far from being straightforward in the real world.

Catching user needs and decision styles

Recommender systems struggle to catch user needs, and companies have implemented different approaches to tackle this issue. Amazon.com, for instance, immediately recognizes the user's identity and recommends a book, without asking for any user input. In contrast (similar to the two travel recommender systems mentioned earlier), www.activebuyersguide.com involves a user searching for a vacation in a multistage interaction. First, the site asks about the vacation's general characteristics (type of vacation, activities, accommodation, and so forth). Second, it asks for details related to these characteristics, then for tradeoffs between characteristics. Finally, it recommends destinations. Both approaches have drawbacks, but an adaptive approach, where questions are fine-tuned as the human–machine interaction unfolds, has more potential.

Researchers have recently argued that recommender systems should support multiple decision styles.[5] The DieToRecs recommender (a case-based travel planning system) supports these decision styles by letting the user enter the system through three main doors: *iterative single-item selection*, *complete travel selection*, and *inspiration-driven selection*.

Iterative single-item selection lets the most experienced user efficiently navigate in the potentially overwhelming information space. The user can select whatever products he or she likes and in the preferred order, using the selections done up to a certain point (and in the past) to personalize the next stage. For example, if the user selects a particular destination, that destination is used to recommend a particular accommodation.

Complete travel selection lets the user select a personalized travel plan that bundles items available in the catalog. The personalized plan is constructed "reusing" the structure of travels built by other users in similar sessions.

Inspiration-driven selection lets the user choose a complete trip by means of a simpler user interface (icon based) and an interaction that is as short as possible. The technology behind this approach is provided by integrating case-based reasoning with interactive query refinement. Interactive query refinement allows a more flexible dialogue management—the system tackles failures due to over- or underspecified user needs, suggesting precise repair actions (constraint relaxation or tightening, respectively). Case-based reasoning provides the framework to cast a recommendation session into a case- and similarity-based ordering of both complete trips and single products.[6,7]

Generating recommendations

The mechanistic idea that from needs (problems), the recommender's intelligent algorithm can deduce the right products (solution) is far too simple. Marketers state that needs can be created such that products can be sold. This motivates the suggestion path in Figure 2. Products shown on a Web site can help create needs by offering examples to users who might not have enough experience to formulate the query as the recommender system might require (see, for example, www.activebuyersguide.com). In other words, an effective travel recommender system should not only notice the user's main needs or constraints in a top-down way but also allow the exploration of the option space and support the active construction of user preferences (in a bottom-up way).

Recent research has emphasized this change of perspective, defining it as *navigation by proposing*.[8] In this approach, the system shows the user examples of products, selected from those that the initial query retrieved. The user can choose a product as the current best choice, which updates the initial query and lets the recommender identify a new set of suggestions. The relevance feedback technique used in information retrieval (for example, Rocchio's method) has influenced this approach, which basically injects new constraints or terms—extracted from the selected item or a corresponding cluster—into the original query. In addition, the approach is conversational in that it supports either a multistage interaction or a dialog that interleaves needs elicitation with products.[9] In multistage interaction, example recommendations elicit user needs by exploiting a dialog control component, which poses only focused questions, determined by the previous interaction steps.

Speaking the right language

As I mentioned earlier, recommender systems must carefully manage the human–machine dialogue such that even a naive user can effectively use the system. Rephrasing a user-centered design slogan: "Recommender systems are about people, not machines." Thus, usability issues, such as choosing the product description language, come to the fore. For instance, asking if the user needs a "hot shoe" or a "manual white balance" in a digital camera could be a "hard to say" question for a naive photographer.

A recommender system's ultimate effectiveness relies on its algorithms and their ability to extract useful and novel products from the catalog.[10] However, even if the recommendations are useful, users will struggle if the help system is poor, the item descriptions are too terse, or the site navigation support is confusing. System usability is such an important issue that even a recommendation that is not useful but correct (for example, a place already visited) can increase a user's trust in the system—a necessary condition for recommendation acceptance.

Recommender systems could become learning environments or simpler information presentation tools, but we must design them to support surplus learning and user behavioral changes; again, usability comes first. Furthermore, the interaction and interface design can deeply affect the user's decision-making process. Different design choices can induce distinct decision strategies and influence the user's affective state (emotions, level of involvement, quality of the flow experience) in peculiar ways.

Recommender systems emerged initially as filtering tools, where the primary concern was to discard, in a large database of products, items inappropriate to user needs. Now, experiences with real recommender systems and research prototypes show that the user tasks and functions supported by such systems are much more varied. We thus should focus on new support functions for expanding the user's horizon.

Acknowledgments

I thank Fabio Del Misser, Elena Not, and Hannes Werthner for comments on an earlier version of this essay. Thanks to Daniel R. Fesenmaier and Josef Mazanec for their inspiring critiques and encouragement. This work is partially funded by the CARITRO foundation (under Contract "eCommerce and Tourism") and the European Union's Fifth Research and Technology Development Framework Programme (under Contract DIETORECS IST-2000-29474).

References

1. J.B. Schafer, J.A. Konstan, and J. Riedl, "E-Commerce Recommendation Applications," *Data Mining and Knowledge Discovery*, vol. 5, nos. 1–2, Jan.–Apr. 2001, pp. 115–153.
2. F. Ricci and B. Smyth, "Recommendation and Personalization in eCommerce," *Proc. Adaptive Hypermedia 2002 Workshop* (AH 2002), Univ. of Malaga, Malaga, 2002.
3. H. Werthner and S. Klein, *Information Technology and Tourism—A Challenging Relationship*, Springer-Verlag, New York, 1999.
4. J. Delgado and R. Davidson, "Knowledge Bases and User Profiling in Travel and Hospitality Recommender Systems," *Proc. 9th Int'l Conf. Information and Comm. Technologies in Tourism* (ENTER 2002), K. Woeber, A. Frew, and M. Hitz, eds., Springer-Verlag, Heidelberg, Germany, 2002, pp. 1–16.
5. D. Fesenmaier et al., *Tourist Decision Model*, tech. report D2.2 DieToRecs IST-2000-29474, EU IST project, 2002; http://dietorecs.itc.it/PubDeliverables/D2.2-V1.0.pdf.
6. F. Ricci and H. Werthner, "Case-Based Querying for Travel Planning Recommendation," *Information Technology and Tourism*, vol. 4, nos. 3–4, 2002, pp. 215–226.
7. D.R. Fesenmaier et al., "DieToRecs: Travel Advisory for Multiple Decision Styles," to be published in *Information and Communication Technologies in Tourism 2003*, A. Frew, ed., Springer-Verlag, New York, 2003.
8. H. Shimazu, "ExpertClerk: Navigating Shoppers Buying Process with the Combination of Asking and Proposing," *Proc. 17th Int'l Joint Conf. Artificial Intelligence* (IJCAI 2001), Morgan Kaufmann, San Francisco, 2001, pp. 1443–1448.
9. M. H. Göker and C. A. Thomson, "Personalized Conversational Case-Based Recommendation," *Proc. Advances in Case-Based Reasoning: 5th European Workshop* (EWCBR 2000), Springer-Verlag, New York, 2000, pp. 99–111.
10. K. Swearingen and R. Sinha, "Beyond Algorithms: An HCI Perspective on Recommender Systems," *Proc. Recommender Systems: Papers from the 2001 ACM SIGIR Workshop*, 2001, http://cs.oregonstate.edu/~herlock/rsw2001.

Adaptive context-aware mobility support for tourists

Alexander Zipf, *European Media Laboratory*

As mobile devices decrease in size, weight, and price and increase in power, storage, connectivity, and positioning capabilities, tourists will increasingly use them as electronic personal tour guides. However, to make such mobile tourist services a success, a range of factors must work together, from technical issues (such as bandwidth, positioning availability, and supported interaction paradigms) to user interface and security issues. We must also consider issues such as the availability of accurate, timely, and localized data, end-user costs (business models), and trust.

Location awareness for mobile users

Resolving these issues becomes more urgent as time-to-market gains importance. However, the danger exists of investing a lot of money into solutions that tourists will not accept. For example, many companies have already started developing mobile city information and navigation systems targeted at tourists (in particular, during the Universal Mobile Telecommunications System (UMTS) hype in Europe). These companies often claim to provide personalized, location-aware solutions, but using buzz