Intelligent Systems for Tourism

"Is E-Commerce Dead, Past Its Prime, or Just Resting?"

This is the title of a recent call for papers for the *Journal of IT Theory and Application*. The special issue aims to discuss the problems of unfulfilled business forecasts and related stock market plunges. However, despite these problems, online transactions are increasing—at least in some sectors, such as the travel and tourism industry. In fact, tourism is the leading application in the B2C (business-to-consumer) arena. Although the slow economy and current political developments have negatively influenced e-commerce, it is still flourishing in the tourism sector.

In the first quarter of 2002, travel and tourism accounted for a total turnover of approximately US$7 billion—an increase of 87 percent from the same quarter in 2001 (see www.comscore.com/news/online_travel_q1_041602.htm). Consumers made 46.7 million hotel reservations worldwide in 2001, netting US$12.9 billion in revenue, according to the Hotel Electronic Distribution Network Association (see www.hedna.org). In addition, 32 percent of US travelers this year have used the Internet to book travel arrangements (see www.nua.com/surveys/index.cgi?f=VS&art_id=905357908&rel=true), and further growth is expected.


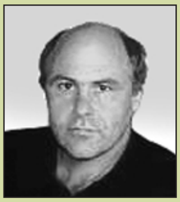

Steffen Staab, University of Karlsruhe, sst@aifb.uni-karlsruhe.de
Hannes Werthner, ITC-irst Research Center, werthner@itc.it.

Tourism industry features

The tourism industry has specific features that explain its importance for economic (regional) development and its inclination toward IT systems.[1]

Tourism is a leading industry worldwide, representing approximately 11 percent of the worldwide GDP (according to the World Travel & Tourism Council's tourism satellite account method[1]). There will be approximately one billion international arrivals in the year 2010 (according to the World Tourism Organization, www.world-tourism.org). Furthermore, tourism represents a cross-sector (umbrella) industry, including many related economic sectors such as culture, sports, and agriculture, where over 30 different industrial components have been identified that serve travelers. In addition, tourism greatly influences regional development, owing to its SME (small- and medium-sized enterprises) structure and relatively small entrance barriers. For example, in the European Union, the hotel and restaurant sector accounts for more than 1.3 million enterprises. This is approximately 8.5 percent of the total number of enterprises, and 95.5 percent of these enterprises are small (with one to nine employees).[2]

Also, because tourism is based on mobility, the supply and demand side forms a worldwide network, where production and distribution are based on cooperation. In addition, it is an information-based industry, so the tourism product is a confidence good, where at the moment of decision-making, only information about the product—not the product itself—is available.

The problem with these statistics is that they refer to different meanings and varying definitions of e-business and e-commerce. Some definitions distinguish between the two, while others view them as the same, and all have their own variables and measurement methods.[3] Even more problematic is that the definitions are all transaction- and business-oriented. They ignore that the Web is also a medium for creating communities, learning new things, and having fun—things that don't always result in business. The Web also encourages user interaction; users can build their own sites to share their travel experiences. Thus, another traveler—rather than hotel management or a travel agency—might provide the most valuable information about a vacation resort.

So, where does AI come into play?

Whereas other industries have a stronger hold on doing things traditionally, the travel and tourism industry has always been open to new technologies. For example, back in the 1960s, airline centralized reservation systems were among the first worldwide computer networks.

In addition, companies traditionally outside the tourism field are entering the sector, mainly from IT and media sectors. Industry features (mainly that IT and media are information-based businesses and are umbrella industries) might explain this trend, or even the change in consumer behavior. For example, consumers use IT not only for information gathering but also to order services over the Internet. A new type of user is emerging who doesn't just try one or two services but all kinds of travel and leisure services. Such users don't mind becoming their own travel agents, but given the extensive use of distributed systems on the Internet, there comes the urgent need to find, combine, and sift through the right pieces of information intelligently.

Today, AI-based developments in the field are at the forefront, such as individualized pricing (priceline.com), reversed multiattribute auctioning (mytraveldream.com), recommendations in bundling products (as described later), Semantic Web applications (Harmonise.org), and mobile applications (described later). In fact, IT developments and research have induced much change in this industry. We can expect this innovation to continue—at both the business level (such as dynamic market structures and prices) and the technology level. In addition, the IT and tourism field represents the nucleus of a new industry that will produce new products, skills, and jobs.

Challenges for intelligent systems in tourism

The industry's dominant features are its heterogeneous and worldwide distributed nature and its strong SME base (especially SMEs in tourist destinations). Another inherent characteristic is mobility, where the entire tourist life cycle is integrated with the respective supplier processes (see Figure A).

Obviously, suppliers' processes cross company borders, leading to enhanced B2BC (business-to-business and business-to-consumer) applications, enforcing cooperation between companies, and supporting mobile communication with the consumer. Given such a framework, future systems should

- Be heterogeneous, distributed, and cooperative
- Enable full autonomy of the respective participants
- Support the entire consumer life cycle and all business phases
- Allow dynamic network configurations
- Provide intelligence for customers and suppliers (interfaces and tools) as well as in the network (which would lead to a set of different services)
- Be scalable and open (with respect to geographical and functional extension)
- Focus on mobile communication (and the notion of *ambient intelligence*), enabling multichannel distribution

Ambient intelligence, in which the surroundings become the interface, is at the focus of European research (see www.cordis.lu). We can define it as the convergence of ubiquitous computing and communication and intelligent user-friendly interfaces. Such systems should be embedded, personalized, adaptive, and anticipatory, and they should provide access for everybody, anywhere, at any time. Whereas today the dominant mode of interaction is lean-forward (that is, tense and concentrated interaction with the computer), it will become laid-back (relaxed and enjoyable). People should enjoy computer interaction for travel planning, and technology should move to the background.

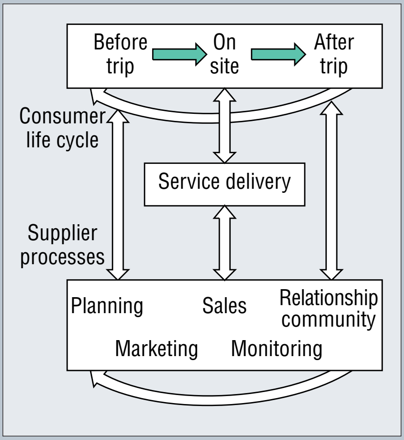

Figure A. The tourist life cycle and suppliers' processes.

Opportunities for research

The e-tourism market's dynamics and the requirements of future systems emphasize e-tourism's importance and raise several technical research issues:

- (Semantic) interoperability and mediated architectures (where we can distinguish between system integration and the semantics issue, related to the problem of too much information)
- E-business frameworks supporting processes across organizations—virtual organizations
- Mobility and embedded intelligence
- Natural multilingual interfaces (also novel interface technologies)
- Personalization and context-based services
- Information-to-knowledge transformations—data mining and knowledge management

The tourism domain is also an excellent example of the trend toward personalized services and a complex market mechanism. It reflects users becoming a part of product creation. So, researchers must also study nontechnical issues related to markets and users, such as

- Dynamic market and network structures
- Pricing and market design
- Design and experimenting business models
- User decision modeling and usage analysis

These research issues underline the importance of an interdisciplinary approach. Many different disciplines should contribute, including computer science, management science, economics, law, statistics, sociology, and psychology.

In this issue

The following essays tackle some of these thorny issues and stress the need to provide a holistic approach. For instance, recommender systems that simply propose itineraries without considering the trust a user must invest will fail. Francesco Ricci elaborates, arguing that recommender systems should not only filter information but also offer completely new suggestions. Analogously, Alexander Zipf discusses how location-based services for tourists must go beyond adding a geographical parameter to some database query. They must consider the personal context to provide an adequate location-based adaptation to maps, itineraries, or products. Furthermore, supporting the user requires providing a globally coherent view of tourism services (Ulrike Gretzel and Daniel R. Fesenmaier provide a general analysis of this topic, and Cécile Paris offers a concrete proposal). Finally, Craig Knoblock points out that users want not only advanced planning but also a system that will take care of their itineraries before, during, and after their travels.

—*Steffen Staab and Hannes Werthner*

References

1. H. Werthner and S. Klein, *Information Technology and Tourism—A Challenging Relationship*, Springer-Verlag, New York, 1999, p. 323.

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]

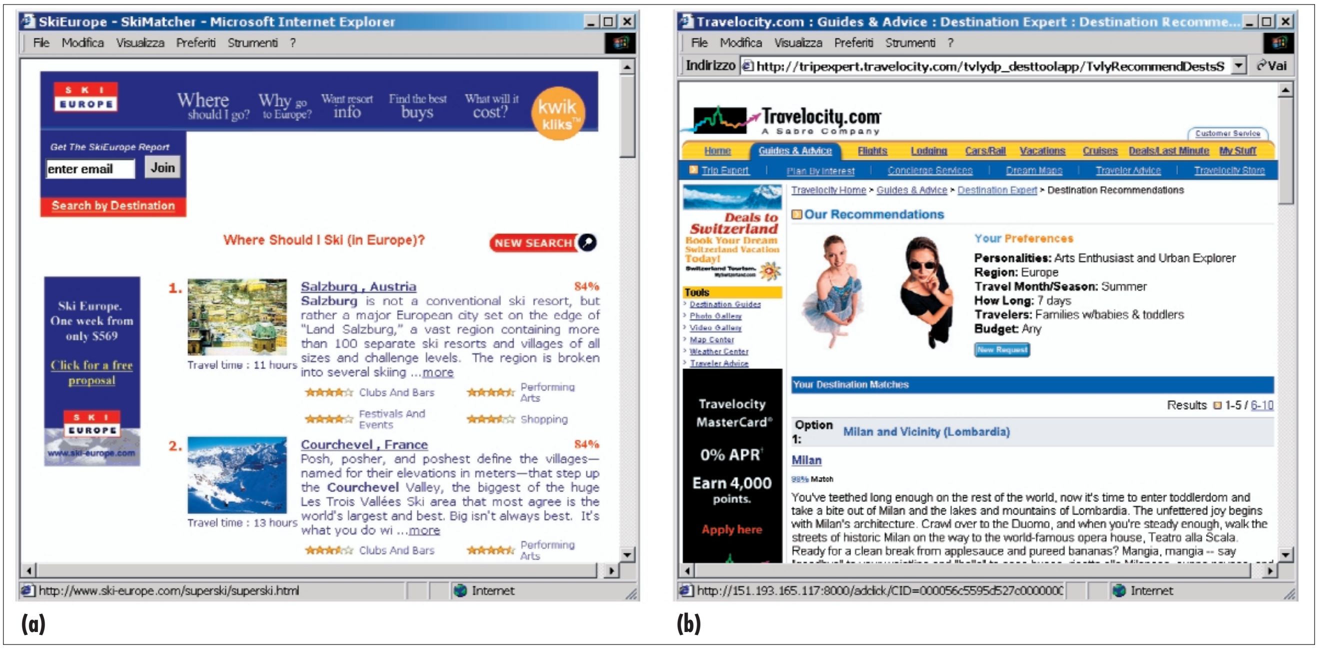

Figure 1. (a) Ski-Europe and (b) Travelocity destination recommendation tools.

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]

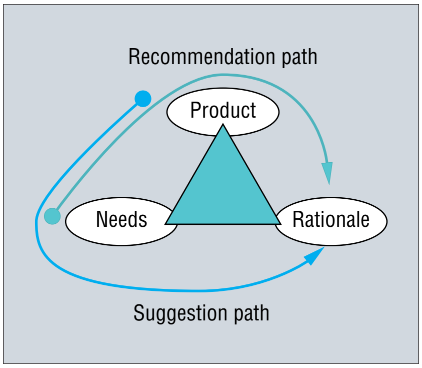

Figure 2. Recommendation and suggestion paths.

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 words does not assure that the lessons about personalization and context-awareness have been thoroughly learned. Personalization and localization are important prerequisites for successful tourist applications, but we must further combine them to better integrate contextual information.

A common assumption is that a system offering only the most relevant information in a given situation (determined by place, time, task, interests, and so forth) will be more successful than a system offering only default information. I argue not only that context-aware information includes standard tourist information (for example, on sights, hotels, and restaurants) but also that supporting services such as tour planning and dynamic maps must profit from personalization and context-awareness. So, we must develop location-aware proactive tips or aggregated services, such as personalized tour proposals, that geographic information systems (GIS) offer. This, in turn, will lead to new (user- and context-) adaptive mobile GIS services for tourists.

Mobile systems for tourists can strongly benefit from the power of GIS. Information interesting to tourists is location-dependent by nature, and GIS can offer this data in a location-aware way. The tourist's position then acts as a filter and parameter for system queries.

In addition, location awareness is a key factor for mobile commerce's success, because it can contribute to a system's ease of use in many ways. GIS can handle spatial and topological queries, allowing navigation and route finding. More advanced GIS data models also let us store and retrieve historical information, which gives much more power to possible queries regarding a region or site's development and history—knowledge that an electronic tour guide should be able to provide. Also, in many queries, the system must handle fuzzy and user-specific measures such as "interesting," "ugly," or "within reach."

Standard GIS functionality can implement personalization and context awareness through *adaptive map generation*, *personalized tour proposals*, or *context-aware proactive tips*, as has been realized in several research projects (such as Crumpet, Deep Map, and others).

Adaptive map generation

When orienting yourself in a foreign area or searching for some kind of business or sight, maps are of great value. They can represent large amounts of information about the area of interest in a single picture—in a (potentially) easily comprehensible form. However, to facilitate the correct reading and understanding of a (sometimes confusing) map, we must design it properly. This is still a challenge for AI and smart systems, because map design is a complex task involving cognitive and psychological aspects.

We must properly design a map to many factors, from technical conditions (screen size, network bandwidth, and so forth), interests, socioeconomic parameters, and the recipient's cognitive abilities, to task and use purpose (for example, if the user wants to receive information on topographic, historical, or other thematic aspects or if the map is solely used for navigation purposes). Obvious examples include map style, color, and use of symbols. Furthermore, we should be able to adapt the map's scale, alignment, and size according to the user's current location, travel method and velocity, or interests. First attempts to model the requirements and develop a framework for adaptive map production are under way, but we need improvements.[1] We have yet to formalize all parameters or discover all psychological or cognitive relationships, and known rules often contradict each other. This includes contradictions between maps needing to show the correct geometry of objects and needing to generalize objects in specific scales and work with limited displays. Oftentimes maps must change the object's position, change the boundary's appearance, and omit details.

> Personalization and localization are important prerequisites for successful tourist applications, but we must further combine these to better integrate contextual information.

Personalized tour proposals

A conventional tour planner computes a path through a street network, given two or more locations on the network. A tour proposal component's task is more complex. It tries to suggest individual sight-seeing tours according to the tourist. Recent developments find user-optimal tours that must respect hard time constraints via different heuristic approaches related to the prize-collecting traveling salesman problem.[2]

Several possibilities exist for including user interests in a tour-proposing algorithm. First, we must identify the set of possible attributes that could influence the choice for a particular section of a route. Such attributes include hard restrictions (or physically given attributes) and dynamic and soft user-specific parameters. Their weighting can vary strongly from one person to another or through time. Such parameters could include aesthetic aspects, the area's social milieu, dislike of traffic, or a preference for nice viewpoints. Such attributes are rarely available in street databases or GIS, nor are these parameters' dependencies modeled formally.

Context-aware proactive tips

A proactive spatial-context module gives tips based on the user's location and interests regarding nearby objects of interest. Marketing quickly adopted this idea of pushing location-dependent information (advertisements) to users, who were not always happy to receive this information. So, it is crucial to integrate fine-grained user and context models into such a service to raise acceptance and keep tourists from turning off such a feature. The system must use more than the user's position and the location of objects to deliver suggestions. Even resolving what "nearby" means to the user in the current situation involves a wide range of personal parameters and contextual information. Parameters that might influence the definition of "nearby" include

- The user's physical condition
- Available transportation
- The weather
- The task
- Travel speed
- Familiarity with the region
- The region's structure
- Terrain

In addition, we must draw the user's attention to these nearby objects of interest nonintrusively to help ensure acceptance. This is a usability challenge, and we must test and evaluate different possibilities.

Geodata handling on mobile devices

How do we get these services to run on mobile devices? To move the described GIS components from heavy servers to thin clients, we must develop lightweight versions of these components. The European Media Laboratory has implemented a first prototype providing geodata-handling-and-processing capabilities such as topological queries (for example, "Is it within …?" or "Does it intersect …?") on a PDA using a spatial access method using an R-Tree.[3] Interoperability is also important, because clients might need to work together with different back-end systems. So, the data model is based on an open standard—the OpenGIS Consortium's Geography Markup Language. Through such developments, we can realize future applications that might use client-side resources to perform GIS services. This would also reduce dependency on network availability. Next we must consider intelligent prefetching and caching strategies in combination with location awareness and resource (network, CPU) adaptivity. Further research on update and synchronization strategies or optimized protocols for adaptive geodata transmission over wireless networks would also be helpful.

To choose relevant information according to interests and context, we need fine-grained user and context models and interaction histories or the ability to analyze a user's movements. Only a few systems currently try to offer all this. We still need to prove that the effort to collect and exploit this information will pay off economically and in terms of improved usability and acceptance.

This leads to a remaining issue for adaptive mobile services: the availability (or lack thereof) of appropriate (timely, correct) data (also geodata) with appropriate semantic metadata. Usable content is generally important—the best adaptation techniques won't help much if the content is insufficient. We can view the lack of semantic understanding as a possible threat to the success of adaptive location-aware systems. Although the Semantic Web initiative is developing the necessary technologies,[4] it might still be a while before semantically enriched data is widespread. Incorporating context is important, but so is further researching sensors and context data.

Finally, we can't neglect security and privacy. Spatial privacy is a major issue for spatially enabled tourism services, because users usually only provide their position in exchange for some benefit. However, localization techniques now let system providers obtain information without much user intervention, and some party in the value chain will always be able to track the user's location. This has raised the awareness of possible fraud and created the need for legal and technical solutions.

Acknowledgments

The Klaus-Tschira Foundation and the Crumpet project (EU-IST-1999-20147) supported this work. I thank all my colleagues at the European Media Laboratory and partner institutions for their valuable input.

References

1. A. Zipf, "User-Adaptive Maps for Location-Based Services (LBS) for Tourism," *Proc. 9th Int'l Conf. Information and Comm. Technologies in Tourism* (ENTER 2002), Springer-Verlag, Heidelberg, Germany, 2002.

2. M. Jöst and W. Stille, "A User-Aware Tour Proposal Framework Using a Hybrid Optimization Approach," to be published in *Proc. 10th ACM Int'l Symp. Advances in Geographic Information Systems*, ACM Press, New York, 2002, pp. 329–338.

3. S. Schmitz, A. Zipf, and H. Aras, "Open GML-Based Mobile Geodata-Handling for PDAs," to be published in *Proc. Annual Conf. Int'l Assoc. Mathematical Geology* (IAMG 2002), Terra Nostra—Schriften der Alfred Wegner Stiftung, Berlin, 2002.

4. S. Staab, "Ontologies' KISSES in Standardization," *IEEE Intelligent Systems*, vol. 17, no. 2, Mar./Apr. 2002, pp. 70–71.

Building Narrative Logic into Tourism Information Systems

Ulrike Gretzel and Daniel R. Fesenmaier, *University of Illinois at Urbana-Champaign*

Tourism is one of the top e-commerce categories and one of the most experiential and complex products sold online. Neither holistic sensory experience nor complexity lends itself to prevailing Web site design and its underlying computing structures. Consequently, tourism experiences are almost exclusively captured as pieces of information that can be described in functional terms and thus easily translated into database structures.

You would assume that the interface's role would be to reintegrate these information fragments into consistent wholes; however, online encounters of tourism information are currently restricted to interactions with interfaces that more or less directly mirror the ontology of the database systems to which they are connected. The database principle's dominance in tourism information system design becomes apparent when you look at the search options and result displays that these systems offer. Users typically must express their information needs about travel destinations (accommodation, transportation, attractions, activities, and so forth) as highly structured queries or choices among search options that more or less reflect the rows and columns in which the data is stored. Even when systems support natural language query, database logic still largely drives the output's structure. The resulting display of bits and pieces of data in the form of item lists or collections of hyperlinks can only meet specific, functional information needs. It thus fails to reflect the complexity of tourism information's role.

We intend to challenge the database approach's dominance in tourism interface design by reflecting on its limitations in terms of effectively conveying relevant information about holistic vacation experiences. We suggest exploring narratives as a means to organize and display tourist information in a way that can communicate the different aspects of tourism experiences, including sensory and emotional components. Integrating narrative principles should lead to tourism information systems that can better meet a multitude of informational needs, provide information that more closely matches human memory, and support tourism decision-making throughout its various stages.

Searching for tourism experiences

It has long been recognized that experiences form the tourism industry's foundation. However, experience's role in consumption (including before, during, and after a purchase) is only now being considered for building effective marketing strategies. Travel activities such as skiing, hiking, shopping, and so forth provide the foundation for travel experiences, while the tourism industry acts as an experience facilitator. In addition, the setting in which activities occur

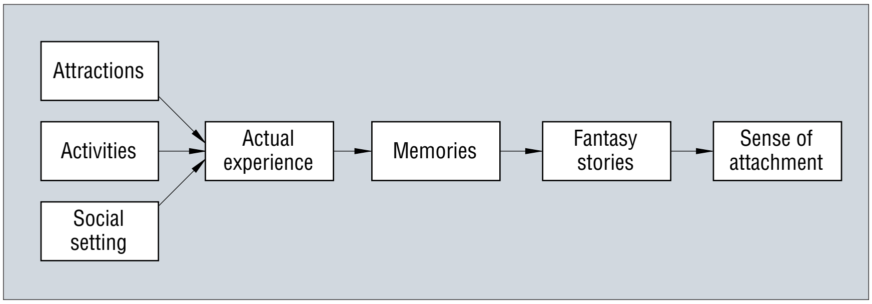

Figure 3. The sequence of travel experience.

contributes substantially to the nature of the experience (see Figure 3). Some have suggested that the memories stored as a result of travel experiences are key to attracting and retaining visitors. Furthermore, researchers have argued that stories—the prominent mechanism for communicating experiences—provide the path through which the tourism industry can build and extend markets.

Marketing and tourism scholars have argued that consumers often evaluate products more on experiential aspects than on objective features such as price and availability.[1] So, experiential information is not only entertaining and stimulating but also essential to the travel decision-making process, because it lets consumers understand and evaluate aspects of the travel product that cannot be easily described in functional terms or expressed as monetary values. Consequently, experiential information responds to the need for a holistic understanding of the specific travel experience to be evaluated.

Despite the travel decision-making process's sequential nature, whereby travelers move step by step through a series of hierarchically organized decision components, the information assimilated to serve as the basis for the various subdecisions must eventually make sense as a whole.[2] Information presented as unrelated items in a list or under separate categories makes it difficult for consumers to construct this cohesive picture of a travel experience. Furthermore, narrative situations—such as querying family and friends or consulting a travel agent—dominate traditional travel-information search. These human travel information providers typically supply contextual information and emphasize particular experiential aspects in a way that lets the information seeker establish mental connections among the various trip elements. Whereas existing tourism information systems appear to successfully provide functional information for specific components of travel decisions, they fail to address the need for holistic, experiential, and conversational ways of communicating travel information.

Database thinking in tourism information systems

The database logic's appeal lies in its clarity and suitability for computational purposes. Lev Manovich defines the database as a conceptual way to represent the world as a list of items.[3] Interacting with a database is a linear experience that differs considerably from viewing films or playing computer games. Database records are often displayed in arbitrary order or according to their relevance with regard to a certain search topic. Furthermore, interfaces following database logics essentially communicate information in fragmentary pieces. Although users can make mental connections between items that are displayed in a list, it requires additional cognitive effort. And if the user fails to establish connections, the number of items that he or she can successfully remember is rather limited. Whether these connections are made and how they are interpreted remain outside the tourism information system's control.

A lack of interpretation of the relationship between items is less problematic for unidimensional search concepts—for example, a search for room rates in a specific hotel. However, tourism experiences are typically multifaceted, and cognitively separating vacation aspects without losing coherence and meaning is often impossible. Searches for information on entire vacation trips are problematic from a database perspective because they are open-ended, vague, ill-defined, multidimensional, and unconstrained. Interfaces that simply provide access to databases, and that feature queries and information displays modeled after database structures, fail to acknowledge the complexity of tourism experiences.

The issue is to create an interface that can add relevance to the information it displays by supporting users in their efforts to imagine coherent experiences. A growing stream of research in psychology and AI suggests that narratives might connect seemingly unrelated items.[4] Narrative interfaces can translate the underlying database into a different kind of user experience that is not only more entertaining but also more informative, because it helps the user derive contextual information necessary to interpret coherent experiences.

Narrative concepts as organizing principles

We can describe narratives as event sequences that create a cause-and-effect trajectory of seemingly unordered items.[3,4] They not only allow for meaningful connections between pieces of information but also simultaneously afford the addition of emotional content and sensory details. Thus, they can convey great quantities of information, especially of experiential nature, in a format that users can quickly and easily assimilate.[5]

Recognizing the importance of narratives for communicating tourism information is not an issue of believing that narrative is the only organizing principle in human memory. Rather, it acknowledges the experiential nature and complexity of vacations and the importance of narratives as a means of communicating tourism experiences. Human beings can organize experience into narratives that help them make sense of the world.[6] The focus of narratives is on meaning and relevance, not on precision. Furthermore, narratives support mental imagery more successfully than other text genres.[7] Thus, narratives provide guidance in terms of interpreting search results but leave room for imagination.

Given the narrative framework's inclusive nature, the significance of restructuring Web site visits into narrative experiences becomes apparent from both a human–computer interaction and a marketing perspective. If narratives are a closer match to human knowledge and communication structures in the travel domain, then they should more effectively educate people about fuzzy or complex travel-related situations. Also, because narratives can link items into logical and consistent wholes, they can better represent bundles of information in contrast to single-item concepts. Furthermore, the inherent entertainment value of narratives promises to engage Web site visitors at a much higher level than interfaces that are direct representations of database structures. Finally, by providing a sequential path, plot, or storyline, narrative approaches can potentially facilitate navigation through unknown knowledge territory.

We should view the narrative concept as an integral part of the tourism experience, which begins with the information search process. Thus it should be understood as an underlying process for travel information search rather than an imposed design principle. Its integration into Web sites in general, and tourism information systems in particular, goes beyond adding yet another story. The importance of narratives lies in applying the narrative concept to both information displays and navigational space to provide users with destination information that affords immediate comprehension, sense-making, and ultimately high personal relevance.

Researchers have proposed alternative strategies for integrating narrative design concepts into tourism information systems through the means of story matching or narrative cues provided through an interactive display.[8] Other domains have already used narrative design; online gaming and social agent development, for instance, have successfully integrated narrative principles into their designs. Tourism information systems could greatly benefit from adopting the narrative approaches developed in these areas. The ultimate goal of tourism information system design is to provide an electronic environment that can meet all travel-related information needs, including searches for the purpose of satisfying hedonic needs and the need for coherence.

References

1. C.A. Vogt and D.R. Fesenmaier, "Expanding the Functional Information Search Model," *Annals of Tourism Research*, vol. 25, no. 3, July 1998, pp. 551–578.
2. D.R. Fesenmaier and J. Jeng, "Assessing Structure in the Pleasure Trip Planning Process," *Tourism Analysis*, vol. 5, no. 1, 2000, pp. 13–27.
3. L. Manovich, *The Language of New Media*, MIT Press, Cambridge, Mass., 2001.
4. R. Schank and R. Abelson, "Knowledge and Memory: The Real Story," *Knowledge and Memory: The Real Story*, R.S. Wyer, ed., Lawrence Erlbaum Associates, Hillsdale, N.J., 1995, pp. 1–86.
5. N. Gershon and W. Page, "What Storytelling Can Do for Information Visualization," *Comm. ACM*, vol. 44, no. 3, Mar. 2001, pp. 31–37.
6. M. Mateas and P. Sengers, "Narrative Intelligence," *Proc. AAAI Fall Symp. Narrative Intelligence*, AAAI Press, Menlo Park, Calif., 1999, pp. 1–10.
7. W.F. Brewer, "Imagery and Text Genre," *Text*, vol. 8, no. 4, 1988, pp. 431–438.
8. U. Gretzel and D.R. Fesenmaier, "Stor-*e*-telling in Destination Recommendation Systems: Concepts and Implications of Narrative Design," *User Modeling and Decision Making in Travel and Tourism Emergent Systems*, Electronic Commerce and Tourism Research Laboratory, Trento, Italy, 2002, pp. 5–7.

Information Delivery for Tourism

Cécile Paris, *Commonweath Scientific and Industrial Research Organization, Mathematical and Information Sciences, Australia*

Tourism has not escaped the information-overload trend, and we see an increasing number of tourism portals providing information filtered by users' specific requests and preferences. The information provided is diverse—including travel planning, route descriptions, and advice on sites to visit—and the filtering mechanisms often require sophisticated strategies.

Unfortunately, this is only half of the problem. Once collected, to facilitate understanding and use, the information must be presented in a manner that is appropriate and natural, tailored to the unique context. This context might include the fact that users are restricted either in their real estate (for example, the size of their handheld devices) or medium (for example, a phone only allows speech). In addition, to gather information on a topic, users typically must issue numerous queries (about hotels, the weather, possible activities, history, and so forth), with the best results achieved when they know about relevant Web sites and their structure.

Here, I argue that filtering is not enough and that tailored delivery includes delivering appropriate and relevant information in a coherent and meaningful way. Delivering a coherent document is an effective way to give users information they can reason about and act on, thus potentially turning information retrieval systems into more powerful communication tools.

I begin by examining two approaches to information delivery: *information filtering* and *natural language generation*. Most systems in the tourism domain use IF, but I argue that coupling it with NLG will produce more coherent and relevant output.

Filtering vs. language generation

The most common way to deliver information responsive to users' needs is based on information retrieval (IR) techniques that have an added filter. On the Web, we typically obtain a user's information by asking him or her to complete a form. A system then adapts the content of the pages presented to the user. (Similarly, when delivering information to a user on the go, the user profile is stored, and filtering happens given the additional location information that a GPS provides.) This type of system employs information-filtering techniques and appends the retrieved answers into a fixed document structure. This approach is fairly rigid, often does not produce a document that is coherent as a whole, and requires the user to issue numerous queries to satisfy his or her information needs.

The other approach uses NLG and user-modeling techniques to form a text by following typical rules of discourse answering a general-information need. These rules ensure that the system includes and coherently organizes all relevant information. This approach is flexible and ensures that the resulting document is coherent and akin to a naturally occurring document. Furthermore, this approach avoids having the user issue numerous specific individual queries. However, the data sources are often restricted to those constructed for the system.

In the tourism domain, where large databases and documents already exist, the second approach is unrealistic, if it requires manually reengineering the data sources. Thus, the IR approach is typically the one adopted. However, delivery would be more useful if the resulting output were organized coherently. In the tourism domain, such a document could be a travel guide. Indeed, simply providing a list of items (for example, as the output of a retrieval engine), in no natural order—or at least, not in an order that makes sense to a user—and requiring the user to issue a number of queries, is not enough. This is true even if items have been filtered to consider a user's profile.

Achieving coherence

In our work at CSIRO, we have followed this approach, coupling NGL techniques with IR. Given a discourse goal that represents a desired effect on the user's mental state and knowledge (such as to familiarize the user with his or her travel location), our system—the Virtual Document Planner (VDP)[1] and Tiddler, its application to the tourism domain[2]—explicitly plans a text. Tiddler obtains relevant information from heterogeneous data sources, on the basis of its discourse rules and constrained by the user model. Tiddler also organizes the retrieved information as a coherent whole, potentially synthesizing it when required.

The system achieves coherence by exploiting coherence relations representing how sections of text relate to each other. We use Rhetorical Structure Theory, which lends itself to computational planning and has been used in many generation and multimedia systems.[3–7] Our text planner uses a library of discourse plans, indicating how to achieve a discourse (or communicative) goal.[3] Such a planner selects, synthesizes, and assembles only relevant content for the user and coherently presents the information. Typical of this approach,[3, 5–6] the discourse plans are designed on the basis of a corpus analysis and, in this domain, represent a travel guide's prototypical structure and contain the information typically included in a travel guide. Documents produced by following these plans thus resemble a travel guide, with tailored information based on the user model.

Complementary to discourse planning is the problem of pulling out bits of information from various sources and combining them to form a valid document. Instead of manually building the knowledge base, as generation-based approaches often do,[3,5] or automatically building the knowledge bases from other sources,[8] we use IR techniques to retrieve information from heterogeneous databases and Web pages. In particular, we use Norfolk, a scripting language developed for synthesizing virtual documents from databases and existing Web pages.[9,10] Norfolk constitutes the interface between the discourse planner and the data sources, retrieving the information the discourse component requests and pulling it together, forming XML fragments.

Coupling an NLG approach to an IR system has several advantages. First, the output's overall organization and content is coherent, reflecting naturally occurring documents and discourse. This output is thus more natural to the reader. Indeed, Tiddler produces a (tailored) travel guide that looks like traditional guides and that users can take along or use as a starting point to get additional information. This provides more effective output than merely filtering information.

Second, our architecture is such that the tailored guide's overall organization and content is constant regardless of the delivery medium. We achieve this by decoupling the content and organization planning from the presentation planning, the stage at which the medium is considered. It ensures that, should users choose a different medium as they move from one setting to another (such as from a desktop to a handheld device), the information is still accessible and easy to browse.[11]

Coupling an NLG approach to an IR system has several advantages. First, the output's overall organization and content is coherent, reflecting naturally occurring documents and discourse.

Third, the system can participate in a meaningful and natural interaction (or dialog) with the user, because it understands what it produces and can reason about the user's further requests, in the context of the current discourse. It achieves this by keeping a *discourse tree* at the end of the planning process. This intermediate structure represents the document content and contains the intermediate discourse goals and the coherence relations that hold among them.[3]

Fourth, the system pulls together information addressing the user's information need, as opposed to answering specific queries. This alleviates the need for the user to issue numerous queries.

Although the results of our initial evaluations were not statistically significant owing to a lack of participants, they indicated that, as a whole, users prefer tailored documents.[2] We have had difficulty, however, proving that tailoring is a more effective communication tool using quantitative measures reflecting actions (for example, did the users use the information?) as opposed to qualitative measures reflecting user judgment (did the users find the output relevant?). Effectiveness is much harder to test than, for example, precision, recall, or readability, because it requires a task's context. In turn, the tourism domain (and many others) requires that we track numerous participants over an extended time period. These conditions are difficult (and costly) to satisfy.

We also must acknowledge concerns about systems that tailor (or filter) information. First, can someone hijack the technology and turn it into a marketing trick? There is a delicate balance between providing relevant information and pushing a company's product. Similarly, the goals of providing the user with useful information and appropriately highlighting a product or service provide conflicting constraints for system implementation. Having the latter override the former will likely annoy users and push them to stop using the system.

A similar concern can be raised regarding coherence and the discourse strategies I discussed: a discourse goal might be to convince the user of something, thus influencing the user's opinion. Furthermore, tailoring information to a user model might be too prescriptive and restrictive, preventing serendipity. It is still possible, however, to develop mechanisms that allow for an element of surprise or that let users bypass the tailored information. Finally, user privacy is an important issue, with systems storing user information.

Acknowledgments

These colleagues participated in the design and development of the Virtual Document Planner and Tiddler: Nathalie Colineau, Francois

Paradis, Anne-Marie Vercoustre, Stephen Wan, Ross Wilkinson, and Ming Fang Wu. I thank Ross Wilkinson, Ming Fang Wu, and Keith Vander Linden for their feedback on this essay.

References

1. N. Colineau and S. Wan, "Mobile Delivery of Customized Information Using Natural Language Generation," *Monitor*, vol. 26, no. 3, Sept.–Nov. 2001, pp. 27–31.
2. C. Paris et al., "Generating Personal Travel Guides—And Who Wants Them?" *Proc. Int'l Conf. User Modeling* (UM 2001), Lecture Notes in Computer Science 2109, Springer-Verlag, New York, 2001, pp. 251–253.
3. J.D. Moore and C.L. Paris, "Planning Text for Advisory Dialogues: Capturing Intentional and Rhetorical Information," *Computational Linguistics*, vol. 19, no. 4, 1993, pp. 651–694.
4. W.C. Mann and S.A Thompson, "Rhetorical Structure Theory: Towards a Functional Theory of Text Organization," *Text*, vol. 8, no. 3, 1988, pp. 243–281.
5. E. André and T. Rist, "Designing Coherent Multimedia Presentations," *Proc. Human-Computer Interaction: Software and Hardware Interfaces* (HCI 93), Elsevier, Amsterdam, 1993, pp. 434–439.
6. B. De Carolis et al., "The Dynamic Generation of Hypertext Presentations of Medical Guidelines," *The New Rev. Hypermedia and Multimedia*, vol. 4, 1998, pp. 67–88.
7. N. Green, G. Carenini, and J. Moore, "A Principled Representation of Attributive Descriptions for Generating Integrated Text and Information Graphics Presentations," *Proc. 9th Int'l Workshop Natural Language Generation*, Assoc. for Computational Linguistics, Ontario, Canada, 1998, pp. 18–27.
8. R. Dale et al., "Dynamic Document Delivery: Generating Natural Language Texts on Demand," *Proc. 9th Int'l Conf. and Workshop Database and Expert Systems Applications* (DEXA 98), IEEE Press, Piscataway, N.J., 1998, pp. 131–136.
9. F. Paradis, "Information Extraction and Gathering for Search Engines: The Taylor Approach," *RIAO* (Recherche d'Informations Assistée par Ordinateur), 2000.
10. F. Paradis, A.M. Vercoustre, and B. Hills, "A Virtual Document Interpreter for Reuse of Information," *Proc. Electronic Publishing*, Lecture Notes in Computer Science 1375, Springer-Verlag, New York, 1998, pp. 487–498.
11. D. Chincholle, "Designing Effective Mobile Services on Small Communication Devices," *Proc. OzCHI 2000: Interfacing Reality in the New Millennium* (OzCHI 2000), CHISIG, Sydney, 2000.

Agents for Gathering, Integrating, and Monitoring Information for Travel Planning

Craig Knoblock, *University of Southern California*

The standard approach to planning business trips is to select the flights, reserve a hotel, and possibly reserve a car. Choosing the airports, deciding whether to park at the airport or take a taxi, and deciding whether to rent a car are often ad hoc decisions based on past experience. The time and effort required to make more informed decisions usually outweighs the cost. Similarly, once we've planned a trip, many of us forget about it until a few hours before the flight. We might check the status of our flights or use a service that automatically sends updated flight information, but other than that, most of us cope with problems as they arise. Beyond flight delays and cancellations, there are a variety of possible events that travelers would like to anticipate, but again, the cost and effort required to monitor these events are not usually deemed worth the trouble. Schedules can change, prices can decrease after purchasing a ticket (if a ticket price goes down, many tickets can be returned and repurchased for a small fee), flight delays can result in missed connections, and hotel rooms and rental cars can be given away because a traveler arrives late.

To address these issues, at USC we developed the Travel Assistant,[1] an integrated travel planning and monitoring system. This system provides an interactive approach to making travel plans where all the information required to make informed choices is available to the user. For example, if the user is deciding whether to park at the airport or take a taxi, the system compares the cost of parking and the cost of a taxi given the choice of airport, the selected parking lot, and the traveler's starting location. Likewise, when the user is deciding which airport to fly into, the system not only provides the cost of the flights but also determines the cost of ground transportation at the destination. Once a traveler has planned a trip, the system monitors various aspects of the trip using a set of information agents that can attend to details for which it would be impractical for a human assistant to monitor. For example, beyond simply notifying a traveler of flight delays, an agent also sends faxes to the hotel and car rental agencies to notify them of a delay and ensure that the room and car will be available. Likewise, when a traveler arrives in a city for a connecting flight, an agent notifies the traveler of any earlier connecting flights and provides the flights' arrival and departure gates.

These innovations in travel planning and monitoring are made possible by two underlying AI technologies. The first is the Heracles interactive constraint-based planner,[2] which captures the interrelationships of the data and user choices using a set of constraint rules. Using Heracles, we can easily define a system for interactively planning a trip. The second is the Theseus information-agent execution system,[3,4] which facilitates the rapid creation of efficient information gathering and monitoring agents. These agents provide data to Heracles and keep track of information changes relevant to the travel plans. On the basis of these technologies, we have developed a complete end-to-end travel planning and monitoring system that is in use today.

Should I drive or take a taxi?

Heracles integrates a wide variety of travel-related data from Web sources to provide the data that users need to plan their trips. It uses information agents (described later) to provide real-time access to any of the online sources related to travel. The system starts by checking a traveler's calendar to identify upcoming trips. The user then selects a trip, and the system interactively helps the user plan the trip, providing the choices of flights, hotels, ground transportation and so forth. For each decision, the system makes a recommendation that the user then selects or overrides. Relationships about the different choices are made explicit using constraints such that choices in one part of a plan are immediately reflected in other parts. For example, if the traveler decides to depart from Long Beach Airport instead of Los Angeles International Airport, then the system would immediately retrieve new directions to the airport and update when the traveler should leave his or her house on the basis of both the distance and updated flight time. Figure 4 shows the initial information about driving to the airport and the result of changing to a different airport. The propagation of updates occurs automatically as part of the constraint network.

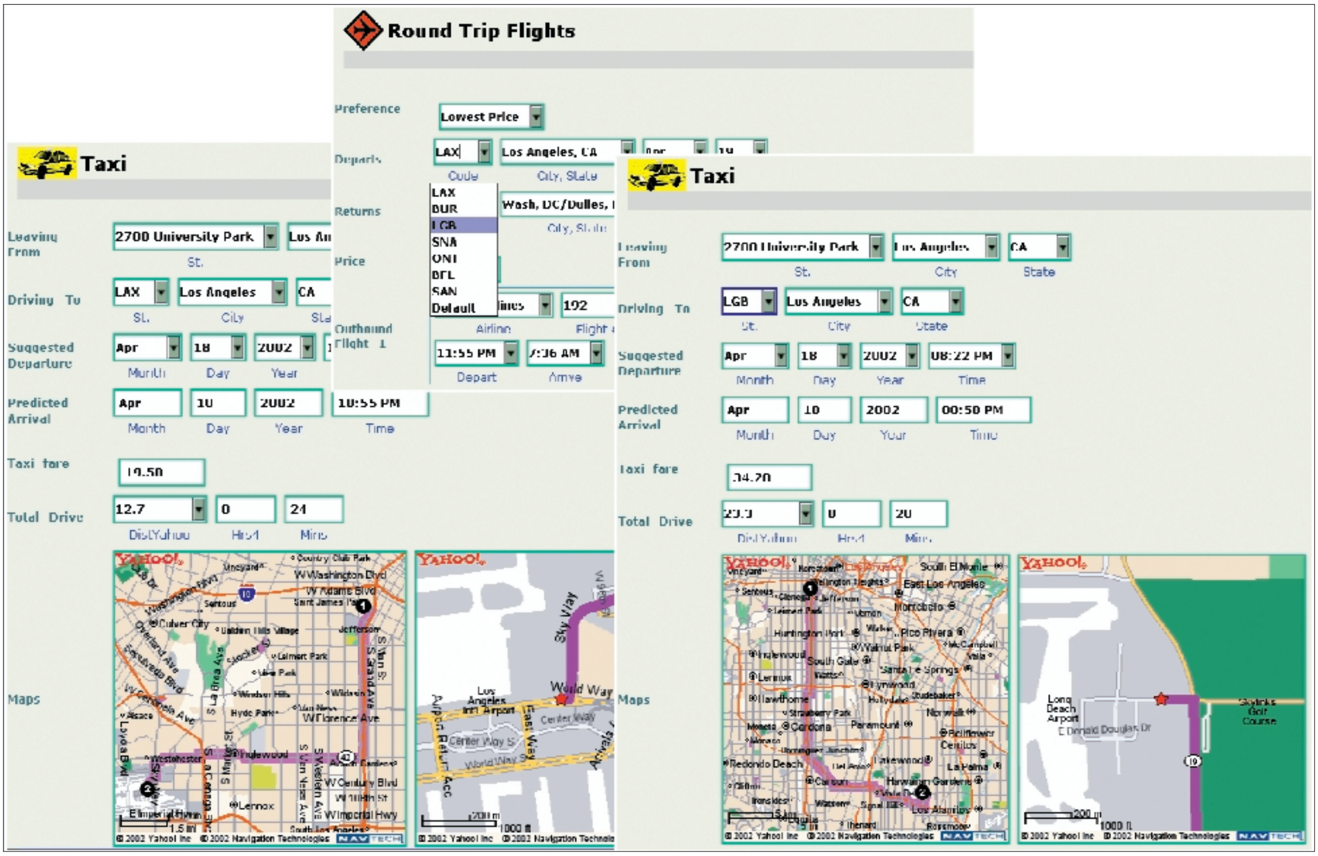

Figure 4. Choices are propagated during the planning process.

Consider the choice of parking your car at the airport or taking a taxi. The appropriate choice depends on a variety of factors, and most people make suboptimal decisions based on simple heuristics. Figure 5 shows an example constraint network for making these types of decisions, where the system calculates the total cost of parking your car, based on the selected parking lot and the number of days you will be gone, and compares this to the cost of taking a taxi, which depends on the distance and taxi rate. Instead of simply guessing about these types of choices, the system can carefully evaluate them and make a recommendation, which the user can still override.

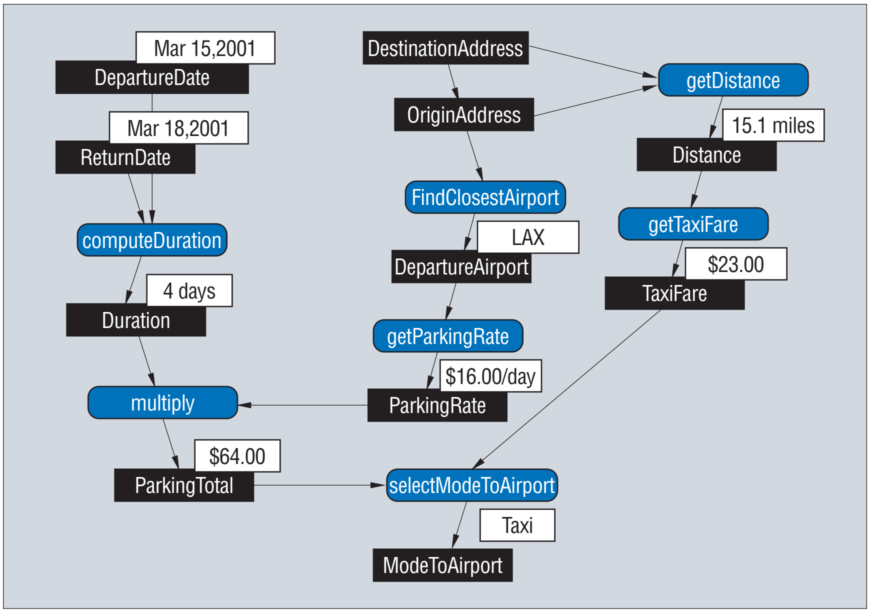

Figure 5. The constraint network for comparing taking a taxi versus driving.

Besides integrating the data to help users evaluate tradeoffs, the system also makes more information available directly to the user. For example, instead of simply choosing arbitrarily between two flights, the system can tell the user the types of plane, how the seats are configured, and the flight's on-time performance. If the system has access to the user's frequent flier accounts, it can even tell the user which airline to fly to accumulate enough miles to qualify for a free ticket. All of this information is organized into the constraint network to give users the information they need to make informed decisions to plan their trips.

You gave away my room but charged my card?

The Travel Assistant uses information

Table 1. Actual messages monitoring agents send.

Agent	Message type	Message text
Flight status	Flight delayed	Your United Airlines flight 190 has been delayed. It was originally scheduled to depart at 11:45 AM and is now scheduled to depart at 12:30 PM. The new arrival time is 7:59 PM.
Flight status	Flight cancelled	Your Delta Air Lines flight 200 has been cancelled.
Flight status	Fax to a hotel	Attention: Registration Desk. I am sending this message on behalf of David Pynadath, who has a reservation at your hotel. David Pynadath is on United Airlines 190, which is now scheduled to arrive at IAD at 7:59 PM. Since the flight will be arriving late, I would like to request that you indicate this in the reservation so that the room is not given away.
Airfare	Airfare dropped	The airfare for your American Airlines itinerary (IAD - LAX) dropped to $281.
Earlier flight	Earlier flights	The status of your currently scheduled flight is: # 190 LAX (11:45 AM)–IAD (7:29 PM) 45 minutes late. The following United Airlines flight arrives earlier than your flight: # 946 LAX (8:31 AM)–IAD (3:35 PM) 11 minutes late.

agents to support both planning and monitoring. Although information agents are similar to other types of software agents, their plans are distinguished by a focus on gathering, integrating, and monitoring data from distributed and remote sources. To efficiently perform these tasks, we use Theseus.

Theseus provides a framework for the efficiently executing these agents. The system is based on a streaming dataflow model, which means it executes the agent's operations as information streams into the system. This is critical in a networked environment, where the main delay in executing an agent is waiting for information from remote sources. A dataflow system maximizes the operations' parallelism by executing as many of the actions as possible in parallel. A streaming dataflow system sends results from one action (for the next action to process) even before the first action has completed. By exploiting these two types of parallelism, the system issues information requests as early as possible to minimize the execution time.

There are two types of agents written in Theseus: *information agents* and *monitoring agents*. Information agents provide all the data needed in Heracles. Each information agent extracts information from a specific Web site. These agents take a particular information request and navigate to the appropriate page on a Web site, locate the required information, and return it as an XML document for processing by another agent or application. We build agents for extracting data from Web sites using machine-learning techniques to train an agent by example on which information to extract from a given Web site.[5,6]

The monitoring agents are built on top of the information agents and perform their tasks at regular intervals. Table 1 shows a set of example messages from the monitoring agents in the Travel Assistant. These agents often must maintain state to keep track of previously returned results. For example, to track prices or schedule changes, the agents must know about previous prices and schedules. The monitoring agents also need to communicate with people, so they support the ability to send email, text messages, or faxes. For example, as Table 1 shows, when a flight is delayed or will arrive after 5 pm, an agent sends a fax to a hotel to ensure it doesn't give away the hotel room.

Isn't this being done commercially?

Most commercial systems for travel planning take the traditional approach of providing tools for selecting flights, hotels, and car rentals in separate steps. There are two integrated approaches to this problem. The first is a system called MyTrip from XTRA On-line. On the basis of personal calendar information, the system automatically produces a complete plan that includes the flights, hotel, and car rental. Once it has produced a plan, the user can then edit the system's individual selections. Unlike with the Travel Assistant, the user cannot interactively modify the plan, such as constraining the airlines or departure airports. Also, MyTrip is limited to only the selection of flights, hotels, and car rentals.

The second approach, which i:FAO Switzerland is commercially developing, uses constraint satisfaction to find a complete itinerary.[7] However, this system assumes that all the relevant data has already been retrieved before the constraint satisfaction process and does not address how to interleave information gathering with constraint satisfaction to handle the enormous amount of potentially relevant information.

For monitoring a trip, some commercial systems (such as the one run by United Airlines) provide basic flight status and notification. However, these systems do not actually track changes in the flight status (they merely notify passengers a fixed number of hours before flights), and they do not notify hotels about flight delays or suggest earlier flights or better connections when unexpected events occur (such as bad weather).

While the Travel Assistant provides a useful set of functionalities, we could improve it in many ways. First, a natural extension would be to more tightly integrate the monitoring and planning capabilities, so that when a flight is cancelled or delayed the system would automatically provide an alternate travel plan. Another limitation is that the set of monitoring tasks are fixed by the system builders. We are working on the Agent Wizard, which is a user interface that would let an end user specify his or her own agents. For example, a user might want to specify an agent that monitors the FAA site and sends notification whenever there is more than a 30-minute flight delay in the user's connecting city so that he or she could consider alternate routes. Or another user might want to define an agent that records real-time flight data every five minutes for a user's flight to build a detailed map of the actual flight path. With the huge amount of information available on the Web, there is no end to the set of monitoring agents that users could specify. ■

Acknowledgments

Thanks to my colleagues who participated in the development of the Travel Assistant: Jose Luis Ambite, Greg Barish, Steve Minton, Maria Muslea, and Jean Oh. This material is based on work supported in part by DARPA and the Air Force Research Laboratory under agreement F30602-00-1-0504, and in part by the Air Force Office of Scientific Research under grant number F49620-01-1-0053.

Coming Next

"Web Services: Been There, Done That?"

References

1. J.L. Ambite et al., "Getting from Here to There: Interactive Planning and Agent Execution for Optimizing Travel," *Proc. 14th Conf. Innovative Applications of Artificial Intelligence* (IAAI 2002), AAAI Press, Menlo Park, Calif., 2002, pp. 862–869.
2. C.A. Knoblock et al., "Mixed-Initiative, Multi-Source Information Assistants," *Proc. World Wide Web Conf.*, ACM Press, New York, 2001, pp. 697–707.
3. G. Barish et al., "A Dataflow Approach to Agent-Based Information Management," *Proc. 2000 Int'l Conf. Artificial Intelligence* (IC-AI 2000), CSREA Press, Athens, Ga., 2000, pp. 657–664.
4. G. Barish and C.A. Knoblock, "Speculative Execution for Information Gathering Plans," *Proc. 6th Int'l Conf. Artificial Intelligence Planning and Scheduling* (AIPS 2002)**,** AAAI Press, Menlo Park, Calif., 2002, pp. 184–193.
5. N. Kushmerick, *Wrapper Induction for Information Extraction*, tech report UW-CSE-97-11-04, Dept. of Computer Science and Eng., Univ. of Washington, Seattle, 1997.
6. C.A. Knoblock et al., "Accurately and Reliably Extracting Data from the Web: A Machine Learning Approach," *IEEE Data Eng. Bulletin*, vol. 23, no. 4, Dec. 2000, pp. 33–41.
7. M. Torrens, B. Faltings, and P. Pu, "SmartClients: Constraint Satisfaction as a Paradigm for Scaleable Intelligent Information Systems," *Constraints*, vol. 7, no. 1, Jan. 2002, pp. 49–69.

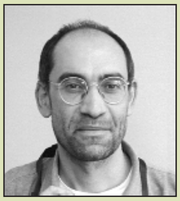

Francesco Ricci is a senior researcher and the technical director of the eCommerce and Tourism Research Lab at ITC-irst. His research interests include recommender systems, constraint satisfaction problems, machine learning, case-based reasoning, and software architectures. He received his PhD in mathematics from the University of Padova. Contact him at ricci@itc.it.

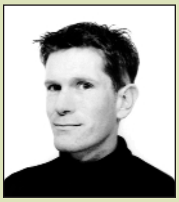

Alexander Zipf is a research associate at the European Media Laboratory in Heidelberg, Germany. Recent projects he's worked on include SmartKom, on multimodal and mobile human-computer interaction, and the IST EU project Crumpet (*cr*eation of *u*ser-friendly *m*obile services *pe*rsonalized for *t*ourism), which is on user- and context-adaptive agent-based location-based services. He studied mathematics and geography at the Universities of Heidelberg and of Manchester, and he received his PhD from the University of Heidelberg. He is member of the International Federation for IT and Travel & Tourism, and of several working groups on mobile and geographical information systems in Germany. Contact him at the European Media Laboratory–EML, Schloss-Wolfsbrunnenweg 33a, 69118 Heidelberg, Germany; zipf@eml.org; www.eml.org/english/homes/zipf/zipf.html.

Daniel R. Fesenmaier is a professor and the director of the National Laboratory for Tourism and eCommerce, University of Illinois at Urbana-Champaign. In addition, he directs research in the areas of internet marketing and the development of knowledge-based decision support systems for tourism organizations. His primary research interests are in tourism marketing, information search, and consumer decision making. He received his PhD in geography from the University of Western Ontario. Contact him at the Nat'l Laboratory for Tourism and eCommerce, Univ. of Illinois at Urbana-Champaign, 104 Huff Hall, 1206 South Fourth St., Champaign, IL 61820; drfez@uiuc.edu.

Ulrike Gretzel is a PhD student at the Institute of Communications Research and a research assistant at the National Laboratory for Tourism and eCommerce, University of Illinois at Urbana-Champaign. Her research efforts focus on the integration of information technologies in tourism organizations and the communication of tourism experiences through online tourism advertising, with a focus on embodied cognition and narrative design. She received her master's with emphases in tourism marketing and international business from the Vienna University of Economics and Business Administration. Contact her at the Nat'l Laboratory for Tourism and eCommerce, Univ. of Illinois at Urbana-Champaign, 104 Huff Hall, 1206 South Fourth St., Champaign, IL 61820; gretzel@uiuc.edu.

Cécile Paris is a principal senior research scientist and R&D group leader at CSIRO/Mathematical and Information Sciences. Her research interests include language technology and multimedia presentation, user modeling, human–computer interaction, authoring tools, and intelligent tutoring systems. She received her BA in computer science from the University of California, Berkeley, and an MS and a PhD in computer science from Columbia University. She is the chair of CHISIG, the Australian Computer Human Interaction Special Interest Group of the Ergonomics Society of Australia. Contact her at cecile.paris@csiro.au; www.cmis.csiro.au/Cecile.Paris.

Craig Knoblock is a research associate professor at the University of Southern California and a senior project leader at the Information Sciences Institute. He is also the chief scientist for Fetch Technologies. His research interests include information agents, information integration, automated planning, machine learning, and constraint reasoning. He received his BS in computer science from Syracuse University and his MS and PhD in computer science from Carnegie Mellon. He leads the Information Agents Research Group. Contact him at Univ. of Southern California, Information Sciences Inst., 4676 Admiralty Way, Marina del Rey, CA 90292; knoblock@isi.usc.edu; www.isi.edu/~knoblock.