Applying Trip@dvice Recommendation Technology to www.visiteurope.com

Adriano Venturini [1] and **Francesco Ricci** [2]

Abstract. Implementing decision support technologies in a real commercial tourism destination portal is challenging. First of all, the peculiar problems related to the tourism domain, which have been studied in the recent years in eCommerce and tourism research, must be considered. But, to provide an effective and useful tool, one must tackle additional requirements arising from the technical and operational environment, which influence not only the software development and architectural issues, but also methodological aspects. This paper describes the main choices we taken and the approach we followed to integrate the Trip@dvice recommendation technology in www.visiteurope.com, the major European tourism portal launched in March 2006, with the goal of promoting Europe as a tourist destination.

1 INTRODUCTION

There is a growing number of web sites that support a traveller in the selection of travel destinations or travel products and services (e.g., events, attractions, flight, hotel). Travellers enter tourism portals searching for information about destinations and related topics, such as, points of interests, historical data, weather conditions, travel packages, flights and hotels. This wide spectrum of information is currently provided by a number of web actors belonging to different categories: on-line travel agencies; tour operators; cruise operators; destination management organizations (multi-destination, regional, city); airlines; hotel chains; convention and visitors bureau [2, 15].

To help visitors to find the information they need and plan their vacation, decision aid technologies are being adopted by tourism portals. These tools must take into account the complexity of the tourist decision process, that has already been described in several studies (see [5] for a survey). A destination is a complex concept, it is an aggregation of more elementary components, such as attractions or cultural themes, more than a geographic area. In addition, a "travel plan", i.e., the output of the decision process, may vary greatly in the structure and content, it could be a flight and hotel, or a complete day by day schedule of the holiday. Because of this, the straightforward implementation of general decision aid and recommendation technologies, already proved to be successful in other sectors [12, 1], cannot be successfully applied to travel planning and destination choice [10].

To overcome these limitations, we have developed Trip@dvice, a travel recommendation methodology that incorporates a human choice model derived from specialized literature on traveler's behavior [4], and extends recommender systems designed for simpler products, and in particular those based on Case Base Reasoning [6]. Trip@dvice supports the selection of travel products (e.g., a hotel or a visit to a museum or a climbing school) and the building of a *travel plan*. In our approach the case base is composed of travel plans built by a community of users. A case is structured hierarchically including components that represent the search/decision problem definition, i.e. the travel's and travellers' characteristics, and the problem solution, that is, the set of products and services included in the plan [11].

Trip@dvice has been chosen by the European Union and by the European Travel Commission (ETC) to provide personalized travel recommendations in the European tourism destination portal, www.visiteurope.com. The main goal of visiteurope.com, which has been launched on March 2006, is to promote Europe as a tourist destination. The Travel Planner tool, based on Trip@dvice technology, is an advanced service allowing tourists to self-bundle their travel plan by receiving personalized recommendations about interesting places and activities that can be practiced in Europe.

To integrate a recommendation technology such as Trip@dvice in a real portal we faced several challenges, normally not considered when dealing with research tools. A tourism destination portal is a complex system composed of a wide set of components and services designed to help the user to find, examine and select the tourist products she is looking for. Supporting the portal visitors in their decisional process means to integrate the recommendation process in the portal concept itself, sharing content, functions and tools of the overall portal. Visitors should be able to bundle their travel plan at any time in their preferred way, either browsing the content through the classical navigation functions, or by selecting the preferred tourist products shown on the home page, or by accepting a suggestion which has been sent via email by a friend.

Therefore the recommender system should exploit the same content repository available in the portal. The European tourism destination portal, stores content related to the 34 European official national tourism organization belonging to ETC, and supports a weakly structured content model representing a wide sets of concepts. Recommendation functions must be adapted to deal with items represented with semi-structured models annotated with metadata information. Furthermore, the tourist target market of the portal is the whole world, thus the portal has to support a wide spectrum of visitors with different characteristics and decision styles [16]. Most of visiteurope.com visitors do not know much about Europe and they are supposed to increase their "destination" knowledge with the portal support.

This paper is organized as follows: Section 2 describes the functions implemented by Travel Planner service and the typical interaction supported by the tool. Section 3 summarizes the methodology implemented by Trip@dvice and how it has been adapted to visiteu-

[1] eCTRL Solutions, Italy, email: venturini@ectrlsolutions.com
[2] ITC-irst, Italy, email: ricci@itc.it

rope.com. Section 4 gives an overview of the architecture we adopted to integrate the recommendation technologies within the portal. Section 5 deals with some open questions and next research activities.

2 THE TRAVEL PLANNER SERVICE AND SUPPORTED SCENARIO

The Travel Planner service provides a set of functions that allow users to plan their vacation in Europe. It manages the current user travel plan, which is a set of tourist items (e.g. regions, cities, attractions, activities, events) selected by the user to self-bundle the travel. Tourist items can be selected from any section of the portal and added to the user's personal travel plan: each tourist item shown by any information service of the portal can be selected and added to the user's personal travel plan. All the items that the user has added to his personal travel plan contribute to create his user model (of the current travel) and will be exploited by the recommendation functions to provide relevant recommendations (see Section 3).

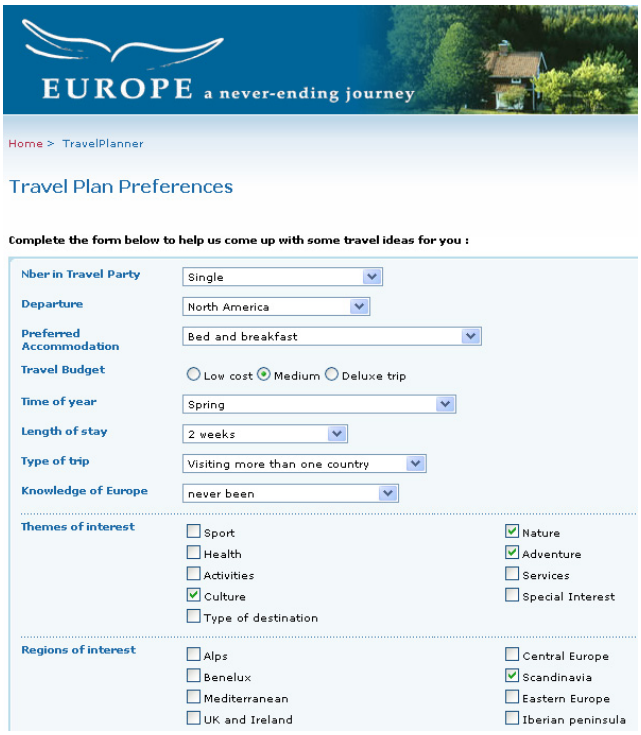

Figure 1. Specifying Travel Preferences

The recommendation functions we have implemented in visiteurope.com are the followings:

- **Search Activity**. Exploiting the information and preferences acquired during the user-system interaction and the experiences contained in past recommendation sessions, the system identifies and recommends ranked tourist items.
- **Seeking for Inspiration**. This function is intended for those users that rather than specifying preferences and constraints, prefer to browse the portal content, looking and rating proposals made by the system. Each system proposal is a coherent bundle of tourist items. If the user is interested in one proposal, she can either add all or some of its components to the travel plan, or get additional proposals similar to the one that she likes.

The initial step of the recommendation process consists in collecting the user's travel preferences, as shown in Figure 1. Here the user can enter general needs and preferences about the travel he is going to plan. Users are not forced to answer all these questions to get recommendations or to use the travel planner services, since this is not the unique source of information for recommendations. Moreover, users can at anytime refine their preferences to better specify their profile. It is important to highlight that the features that the user can here specify (e.g. themes and regions of interests) are not predefined by the recommender system. In fact, they belong to the classification trees used to classify items in the content management system of the portal. Features can be added and removed by the portal content manager, adapting the tool to the evolving needs of the tourism organization. Travel preferences become part of the current case and will be exploited by the Search Activity and by the Seeking for Inspiration functions to deliver product recommendations.

In the Search Activity function the user further specifies detailed preferences about the activities sought (Figure 2). Additional criteria can also be set, like the preferred countries and themes of interests. Note that the top level categories (e.g. Nature) already selected in the Travel Preference section are automatically expanded to allow the user to select children categories (e.g. Natural Park, Landscape) and refine her interests definition. By pushing the search button, activities matching the detailed preferences are ranked and shown to the user, who can add the preferred ones to his travel plan (Figure 3).

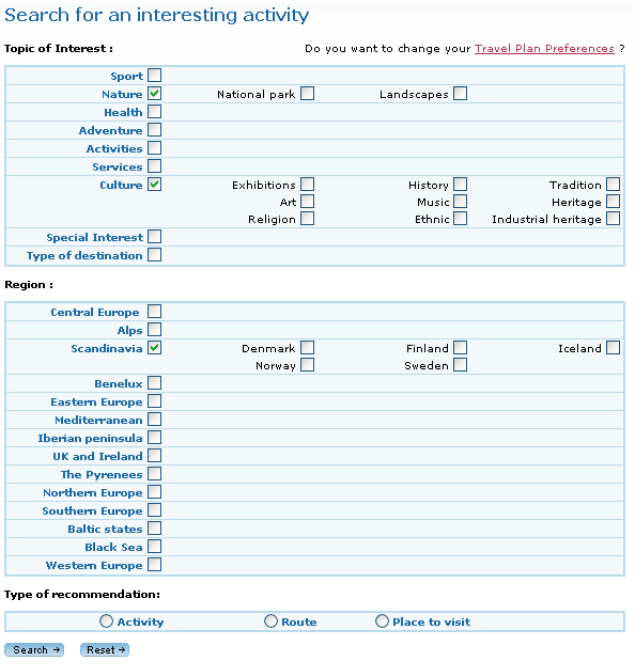

Figure 2. Specifying Detailed Preferences

The other main recommendation function which has been implemented is called "Seeking for Inspiration" (Figure 4). Here the system initially proposes some alternative bundles of tourist items. The user can examine each proposal in details and add all or part of it to her travel plan, or can get additional recommendations by selecting one bundle ("My favorite is this"). In this case, the system, exploiting the whole history of alternative bundles preferred by the user, proposes new bundles that are supposed to be closer to the user ideal travel.

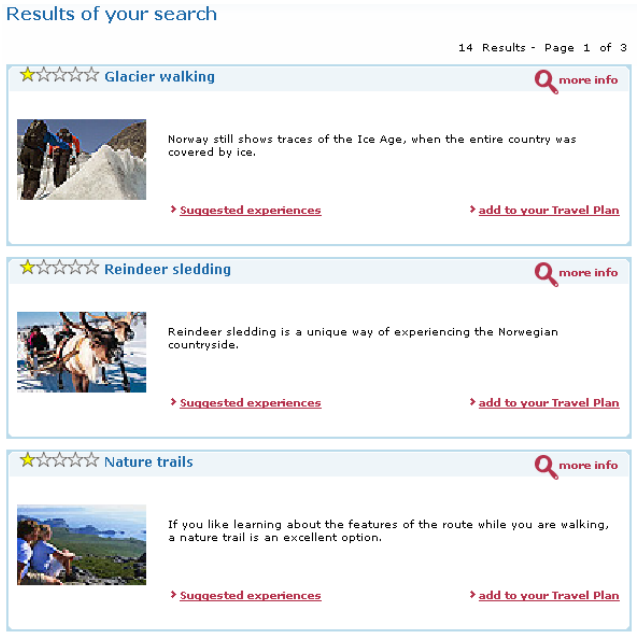

Figure 3. Search Results

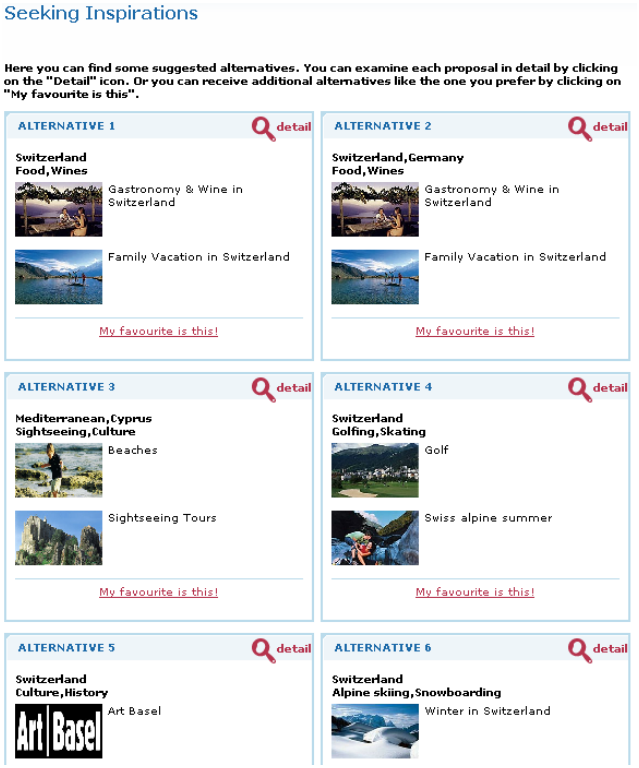

Figure 4. Seeking Inspiration

3 TRIP@DVICE RECOMMENDATION TECHNOLOGY

Trip@dvice is a recommendation tecnology which integrates case base reasoning and cooperative query answering to provide personalized suggestions about interesting tourist products. This section briefly describe the case structure and the methodology exploited in the two implemented recommendation functions, Single Item Recommendations and Seeking for Inspiration. Full details of Trip@dvice recommendation methodology can be found in [11].

3.1 Case Model

Trip@dvice bases its recommendations on a case model that captures a unique human-machine interaction session. A case collects: information provided by the user during the session, the products selected, and some long-term preferences and demographic data of the user if he is registered. Recommendation sessions are stored as cases in a repository (case base) [6]. In addition to the case base, product catalogues are stored in a relational database.

In Trip@dvice, a case represents a user interaction with the system, and it is built incrementally during the recommendation session. A case comprises the following three main components:

- **Collaborative Features (clf)** are features that describe general user's characteristics, wishes, constraints or goals (e.g. desire to relax or to practise sports). They capture general travel preferences and characteristics that are relevant to the decision-making process, but cannot be considered as product features (e.g. the traveller group composition). These features are used to measure case similarity.
- **Content Queries (cnq)** are queries posed over the catalogues of products. Content queries are built by constraining (content) features, i.e., descriptors of the products listed in the catalogues.
- **Cart** contains the set of products and services chosen by the user during the recommendation session. A cart represents a meaningful (from the user's point of view) bundling of different products. For instance, a travel cart may contain some destinations, some accommodations, and some additional attractions.

Figure 5 shows a case example. It represents a user, who is single, has a medium budget, interested in nature and culture and in visiting Scandinavia or Central Europe. These are the collaborative features. Then there is one query, where the user has further specified some countries he is interested in (Norway and Sweden). From the result set of that query, the user has selected the Raindeer sledding and the Trono old church site.

3.2 Single Item Iterative Selection

The overall process supported by Trip@dvice for the single item iterative selection is shown in Figure 6. The full description of this process can be found in [11], here we provide just a brief summary of the methodology. The user interacts with the recommender system by asking for recommendations about a product type (e.g., an activity)(1: AskRecommendation(q) in Figure 6). The system replies to this query q either recommending some products, or, in case the query fails, suggesting some query refinements. The RecEngine module manages the request. First, it invokes the EvaluateQuery function (2) of the Intelligent Query Manager module (IQM), by passing the query. This function searches the catalogue for products matching the query. If too many or no product matches the input

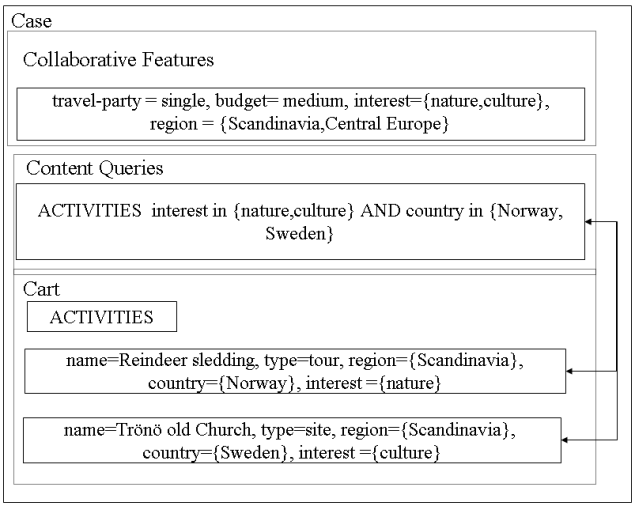

Figure 5. Example of a case

query q, then IQM analyzes q and determines a set of query refinements to suggest.

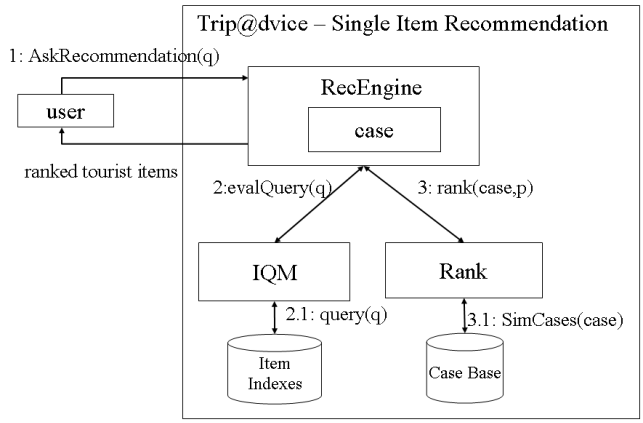

Figure 6. Single item recommendation with Trip@dvice.

The products selected by the user's query are ranked by invoking the Rank method (3). Rank receives as input the current case and the set of products p to be ranked. The current case is used to retrieve the set of K most similar cases from the case base, and the products contained in these retrieved cases are then used to rank the user selected products p. Finally, the ranked products are returned to the user. Ranking is computed exploiting two similarity metrics: first, the case base is accessed to retrieve the K most similar cases (reference cases) to the current one. Then, the products contained in the carts of the retrieved reference cases (reference products) are used to sort the products selected by the user's query. The basic idea is that among the products in the result set of the query one will get a better score if it is similar/equal to a product chosen by a user with similar needs and preferences. This function implements a hybrid "cascade" recommendation methodology [3], where first a conversational approach is used to select a set of options and then a collaborative via content approach is used for ranking [8].

3.3 Seeking for Inspiration

In this section we describe the seeking for inspiration function, designed for users that would prefer to browse travel options and get inspired by them before taking some decision. The recommendation proceeds according to the following loop (see Figure 7):

- **Retrieval.** The process is started with an initial case c that could be either the current case or a random case taken from the case base. It is the current user's case if the user has already created it in previous stages of the interaction, for instance by specifying his general travel preferences or adding some tourist products to the cart. Conversely, if the current case does not yet contain any useful information, then the recommender does not know anything about the user and could only make an initial random guess of what could be a good travel for the user. The retrieval module searches for the M most similar cases in the case base and passes this set to the next module, the Selection.
- **Case Selection.** In the second step the M cases retrieved from the case base are analyzed to select a very small subset of candidates to be presented to the user (six in the European tourism destination portal). To select this small set of candidates a greedy algorithm is used [14]. It iteratively selects a case starting from the initial case and puts the selected case in the result set. The case added at each iteration is the one that minimizes the sum of the similarities between itself and the cases already in the result set. In this way, it selects from the M cases those that are enough different among themselves to provide a reasonable variety of options to the user.

The selected cases are shown to the user, who can choose the preferred one. The retrieval step is executed again, with some notable changes. The seed case is now the case that received positive feedback from the user, and the number of cases retrieved from the case base, M in the first retrieval, is decreased by a factor $0 < \lambda < 1$. The rationale for decreasing the number of retrieved cases is to stronger focus the retrieval around the input case, since at this new stage we must count on the positive evaluation of the user on the selected case. This recommendation functionality follow the recommendation by proposing and comparison-based recommendation approaches introduced in [13, 7]

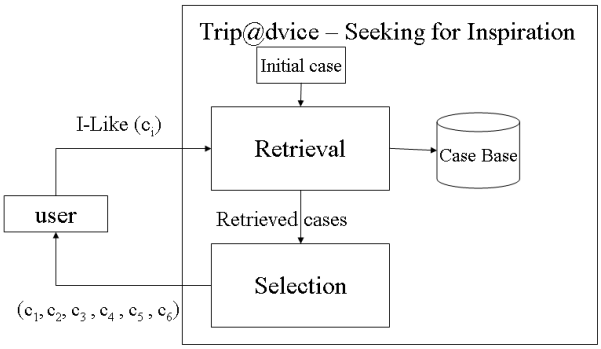

Figure 7. Seeking for inspiration recommendation cycle.

4 INTEGRATION ARCHITECTURE

As already said in the Introduction, integrating the recommendation technology in a full portal presents several issues that should be considered. [9] describes several options for deploying a recommender

system. One, is to build an independent web site, which provides recommendations about products described in a second portal. TripAdvisor is a popular example of this approach, it recommends hotels that are sold by a number of merchant web sites (e.g. expedia and travelocity). Another approach consists of deploying the recommendation technology inside the system itself, and tightly integrate the recommendation functions with the full system. In this way the designer can mix new recommendation functions with more traditional one (e.g. browsing and search), controlling the full human-computer interaction. In visiteurope.com we adopted the second option. Figure 8 shows a high level view of the modules relevant to our discussion. In particular it shows some components which are parts of a typical web portal architecture: the Navigation component, allowing the user to navigate the content of the portal following a navigation tree; the Content Management System (CMS), which is in charge of managing the content items of the portal; and Trip@dvice, providing the recommendation and travel plan management functions and the Travel Planner graphical user interface.

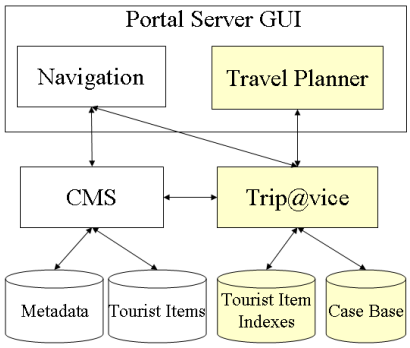

Figure 8. Single item recommendation with Trip@dvice.

These components should communicate among them. The Travel Planner GUI uses the Trip@dvice functions to display the current user travel plan and recommend items to the users. The Navigation service interacts with Trip@dvice to add to the user' travel plan the tourist items that the user chooses navigating in the portal or exploiting the recommendation functions. In this architecture, Trip@dvice must retrieve from the CMS the content items and their metadata to build the recommendations and must refer to them when storing cases.

5 CONCLUSION

In this paper we have illustrated the design and the implementation of a decision aid tool aimed at supporting the traveller's decisional process in the European tourism destination portal www.visiteurope.com. We have presented the main project goals, the interaction we have supported, its methodology and the integration architecture. The proposed recommendation technology, before being deployed in visiteurope.com, has been evaluated in a number of empirical studies involving hundreds of users [16]. The travel plans case base has been initialized with some cases for each country and is now growing, acquiring new travel cases as the system is used by the visitors. A case management tool, supporting the content manager in the validation and selective inclusion of new cases in the repository, has been developed and integrated in the portal infrastructure.

A number of open problems must be still faced in the near future. Techniques to analyze the case base and select a subset of cases among the potentially huge set of those created by the portal visitors must be developed. Moreover, tourism organizations should be supported in better understanding their users and their tourism offers by exploiting the knowledge stored in the case base.

REFERENCES

[1] Gediminas Adomavicius and Alexander Tuzhilin, 'Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions', *IEEE Transactions on Knowledge and Data Engineering*, **17**(6), 734–749, (2005).

[2] Dimitrios Buhalis, *eTourism: Information Technology for Strategic Tourism Management*, Prentice Hall, 2003.

[3] Robin Burke, 'Hybrid recommender systems: Survey and experiments', *User Modeling and User-Adapted Interaction*, **12**(4), 331–370, (2002).

[4] K. Grabler and A.H. Zins, 'Vacation trip decision styles as basis for an automated recommendation system: Lessons from observational studies', in *Proceedings of the ENTER 2002 Conference*, Innsbruck, Austria, (January 22-25 2002). Springer Verlag.

[5] Y.-H. Hwang, U. Gretzel, and D. R. Fesenmaier, 'Behavioral foundations for human-centric travel decision-aid systems', in *Proceedings of the ENTER 2002 Conference*, Innsbruck, Austria, (January 22-25 2002). Springer Verlag.

[6] Fabiana Lorenzi and Francesco Ricci, 'Case-based recommender systems: A unifying view', in *Intelligent Techniques for Web Personalization, IJCAI 2003 Workshop, ITWP 2003, Acapulco, Mexico, August 11, 2003, Revised Selected Papers*, pp. 89–113. Springer, (2005).

[7] Lorraine McGinty and Barry Smyth, 'Comparison-based recommendation', in *Advances in Case-Based Reasoning, Proceedings of the 6th European Conference on Case Based Reasoning, ECCBR 2002*, eds., S. Craw and A. Preece, pp. 575–589, Aberdeen, Scotland, (4 - 7 September 2002). Springer Verlag.

[8] Michael J. Pazzani, 'A framework for collaborative, content-based and demographic filtering', *Artificial Intelligence Review*, **13**, 393–408, (1999).

[9] Ulrich Rabanser and Francesco Ricci, 'Recommender systems: Do they have a viable business model in e-tourism?', in *Information and Communication Technologies in Tourism 2005*, Springer Verlag, Wien - New York, (2005).

[10] Francesco Ricci, 'Travel recommender systems', *IEEE Intelligent Systems*, **17**(6), 55–57, (2002).

[11] Francesco Ricci, Dario Cavada, Nader Mirzadeh, and Adriano Venturini, 'Case-based travel recommendations', in *Destination Recommendation Systems: Behavioural Foundations and Applications*, eds., D. R. Fesenmaier, H. Werthner, and K. W. Wober, CABI Publishig, (2006).

[12] J. Ben Schafer, Joseph A. Konstan, and John Riedl, 'E-commerce recommendation applications', *Data Mining and Knowledge Discovery*, **5**(1/2), 115–153, (2001).

[13] Hideo Shimazu, 'ExpertClerk: Navigating shoppers buying process with the combination of asking and proposing', in *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence, IJCAI 2001*, ed., Bernhard Nebel, pp. 1443–1448, Seattle, Washington, USA, (August 4-10 2001). Morgan Kaufmann.

[14] B. Smyth and P. McClave, 'Similarity vs diversity', in *Proceedings of the 4th International Conference on Case-Based Reasoning*, Springer-Verlag, (2001).

[15] Hannes Werthner and Stefan Klein, *Information Technology and Tourism - A Challenging Relationship*, Springer, 1999.

[16] Andreas H. Zins and Ulrike Bauernfeind, 'Evaluating travel recommender systems: A case study of dietorecs', in *Destination Recommendation Systems: Behavioural Foundations and Applications*, eds., D. R. Fesenmaier, H. Werthner, and K. W. Wober, CABI Publishig, (2006).