Tailoring the recommendation of tourist information to heterogeneous user groups

L. Ardissono, A. Goy, G. Petrone, M. Segnan and P. Torasso

Dip. Informatica, Università di Torino, Corso Svizzera 185, I-10149 Torino, Italy
e-mail: {liliana,goy,giovanna,marino,torasso}@di.unito.it

Abstract. This paper describes the recommendation techniques exploited in INTRIGUE (INteractive TouRist Information GUidE), an adaptive recommender system that supports the organization of guided tours. This system recommends the places to visit by taking into account the characteristics of the group of participants and addressing the possibly conflicting preferences within the group. A group model is exploited to separately manage the preferences of heterogeneous subgroups of people and combine them, in order to identify solutions satisfactory for the group as a whole.

1 Introduction

Web-based information systems have become very popular tools to retrieve specialized information. In particular, the provision of tourist information is extremely appealing, as it supports the search for up-to-date information about services and attractions, without relying on books or travel agencies. The development of Web-based tourist guides is however challenged by the variety of user needs to be satisfied during the presentation of the information. Users are typically interested in different types of attractions (pieces of art, scientific attractions, natural parks, etc.). Moreover, most people do not travel alone, so that possibly conflicting requirements have to be taken into account when recommending the places to visit. Therefore, the presentation of tourist information requires personalized travel guides, satisfying individual information needs; e.g., see [13, 18, 19, 16].

In this paper, we present INTRIGUE (INteractive TouRist Information GUidE), a Web-based adaptive system that provides information about tourist attractions and services, such as accommodation and food. In the following presentation, we focus on the techniques for the generation of the personalized recommendations tailored to the preferences of the group of participants.

INTRIGUE supports the user in a combined search for tourist attractions, based on orthogonal criteria, such as category-based and geographical search. Moreover, the system dynamically generates multilingual presentations, by exploiting efficient template-based NL generation techniques. Our current prototype presents information about the city of Torino and the surrounding Piedmont area, in Italian and in English. The system assists the user in the organization of a tour by providing personalized recommendations about attractions and services. Moreover, it offers an interactive agenda that supports the scheduling of the tour, by considering both the user's needs and the constraints concerning the opening hours, the average visit time for the selected places, and so forth.

As far as the recommendation functionality is concerned, INTRIGUE deals with a structured model of the group of people traveling together, to manage the possibly conflicting preferences of the subgroups. This approach supports alternative recommendation criteria, which the user can select to receive suggestions customized according to specific viewpoints. For instance, the system supports the suggestion of a solution that satisfies all the participants in a more or less a uniform way. However, also the search for solutions focused on the preferences of specific subgroups is available (in our prototype, we consider children and disabled people). In this way, the user can ask for suggestions focused on particular perspectives. As noticed in [7], an essential feature for an intelligent system is the explanation of the reasons for its own suggestions. This aspect becomes even more crucial for group recommendations, where there is no immediate correspondence between the user's preferences and the system's decisions. For this reason, we have developed an explanation technique which supports a clarification of the evaluation strategies adopted by the system in the recommendation. This technique enables the system to specify which properties are most suitable for the characteristics of the various subgroups, therefore helping the user to select the items to include into the agenda in a very informed way.

The paper is organized as follows: Section 2 describes the structure and the management of the group model; section 3 describes the recommendation criteria used in the system to provide users with personalized suggestions tailored to a possibly heterogeneous tourist group. Section 4 describes some related work and section 5 closes the paper, outlining some future work.

2 Management of group models in INTRIGUE

2.1 User groups

The management of heterogeneous groups having different and potentially conflicting preferences is essential to enhance the recommendation task in several application domains, other than the tourism one. For instance, consider the management of personalized television services and the suggestion of items to purchase in a group environment.

In some cases, group preferences have been managed by exploiting stereotypical models that describe the preferences of the group as a whole. For instance, in the TV domain, family models have been used to customize Electronic Program Guides to the standard preferences of groups formed by adults and children. Although this approach suits the applications having a small number of typical user groups to be considered, it is not flexible enough to manage the cases where the groups can be highly heterogeneous and there are many possible combinations of user classes. Tourist services represent an interesting example of this case, because people having very different preferences and requirements may join the same tour, which should be organized by taking into account the interests of all the participants. To this extent, a different group model is needed, in order to avoid a combinatorial explosion in the number of models to be described. In particular, the management of the user group should be considered as orthogonal with respect to the management of different user classes.

Ideally, a group model could be managed as the integration of one individual user model for each member of the group. However, this approach would not scale up when large groups are considered. Moreover, to address individual preferences, the system would need individual descriptions of the group members, therefore imposing strong overhead on the user interacting with the system.

We adopt a different approach, which overcomes the drawbacks of the previous ones. We consider a group as a set of people that can be partitioned into a limited number of subgroups.[1] Each subgroup is modeled as a class of users having similar preferences and the preferences of the whole group can be inferred by combining the preferences of the its subgroups. Moreover, an influence on the system's recommendation is evaluated for each subgroup, in order to support recommendation strategies tailored to the preferences of specific subgroups. A subgroup could be particularly influent either because it represents a very significative portion of the tourist group, or because its members belong to a class having special needs (e.g., children and disabled people).

The types of (sub)group to be considered and their relevance are domain dependent and have to be defined on the basis of the personalization requirements to be addressed. For instance, in a tourism domain, groups with special needs, such as children and disabled, could be given maximum relevance in order to take their needs in particular account. However, other groups could be considered; for example, animals could be "members" of the group, with specific preferences, e.g., the one for places where they are accepted. In the next sections, we describe the management of group models adopted in our system.

2.2 Structure of the group model

INTRIGUE exploits a structured group model, where uniform subgroups are represented as distinct entities. A subgroup model is associated to each *homogeneous subgroup* of people planning the tour together. Each subgroup model is structured in three portions:

- The "Characteristics" section provides information about the characteristics of the participants, acquired by the system by questioning the user via a registration form. For instance, Figure 1 represents a subgroup of people aged between 46 and 55, with a human science background, full mobility capabilities, partial vision capabilities and interested in arts and history.
- The "Preferences" portion specifies the system's predictions for the subgroup preferences. Each preference is represented as a slot and includes:

[1] In the simplest case, the group is formed by homogeneous people and consists of a single subgroup.

Characteristics:
  Age: 46-55;
  Background: human_science;
  Mobility: full;
  Vision: partial;
  Interest_for_arts: yes;
  Interest_for_history: yes;
  Interest_for_science: no;
Preferences:
  Special_transportation_systems:
    Importance: 0; *Values*: missing: 0.3; some: 0.4; present: 0.3;
  Special_facilities_for_vision:
    Importance: 0.8; *Values*: missing: 0.1; some: 0.4; present: 0.5;
  Historical_value:
    Importance: 1; *Values*: low: 0.05; medium: 0.1; high: 0.85;
  Artistical_value:
    Importance: 1; *Values*: low: 0.05; medium: 0.35; high: 0.6;
  Scientific_value:
    Importance: 0.5; *Values*: low: 0.3; medium: 0.4; high: 0.3;
  ...
Group Information:
  Cardinality: 5;
  Relevance: 0.4

Fig. 1. An example subgroup model for the tourism domain.

  - An "Importance" facet, which specifies the importance of the preference to the subgroup. For instance, the travelers represented by the model shown in Figure 1 have no interest for special transportation systems (the importance is 0). In contrast, their interest for the historical value of tourist attractions is extremely strong (importance = 1).
  - A probability distribution over the values of the preference. For instance, the described tourists very likely prefer attractions having high historical value ("high" is dominant in the distribution), while the probability that they prefer attractions with low historical value is almost null.
- The "Group Information" section stores general information about the subgroup. Each subgroup has a *cardinality*, specifying the number of people forming it, and an *relevance*, representing an estimate of the weight that the preferences of a prototypical member of the subgroup should have on the selection of tourist attractions to be recommended. The relevance ranges from 0 (null relevance), to 1 (maximum one). In our example, the subgroup is formed by 5 people and has a medium relevance (0.4).

2.3 Knowledge about user classes

We manage the presence of users characterized by different preferences and requirements by exploiting stereotypical information that describes the characteristics of the various user classes. The tourist population can be segmented according to different perspectives; for instance, we can consider their interests for tourist attractions, their knowledge about arts and other topics, or their mobility and vision capabilities. These perspectives have influence on different sets of preferences, which the system can exploit to evaluate a tourist attraction. For our prototype, we have specified the following perspectives, each one including a set of stereotypes that represent the tourist classes described from that viewpoint:

- "Age_range" perspective: the traveler population is segmented by ranges of age. The main goal is to distinguish children from adults and to model interests for special activities (such as playing) accordingly. The only relevant characteristic is "age", but the stereotypes predict specific interests for activities and types of documentation about tourist places.
- "Interests" perspective: in this cluster of stereotypes, the tourists' interests are modeled. People are segmented into groups characterized by different educational backgrounds (e.g., historical, technical, etc.), also depending on the explicit interests declared in the registration form. This cluster makes predictions on the preferences for different types of tourist attractions: e.g., some places may have a noticeable value from the historical point of view, others may excel in scientific or technological aspects.
- "Mobility_capabilities" perspective: the population is segmented to characterize different mobility capabilities. The preferences concern the reachability of places and the availability of special transportation systems.
- "Vision_capabilities" perspective: this segmentation concerns the travelers' sight and makes it possible to describe the preferences of people having full, partial, or null vision capabilities.

2.4 Representation of user classes

The stereotypical information is stored in a knowledge base, where it is organized as a set of clusters, each one associated to a different perspective: e.g., age, mobility and vision capabilities. A cluster contains a list of stereotypes, representing the classes of tourists forming the partition: e.g., the "vision_capabilities" cluster includes three stereotypes, representing the people having complete, partial and null vision capabilities. Similar to the representation defined for the SeTA system [3], a stereotype includes a set of classification data, describing characteristics of travelers belonging to the represented class, and a set of preferences, describing the typical requirements of such people for properties of the tourist attractions. For instance, the stereotype describing people with null vision capabilities has only one significant classification data, i.e., the "vision", which is a trigger for the stereotype. Moreover, the main predicted preference concerns the availability of vocal presentation devices.

As clusters represent different viewpoints for describing people, their stereotypes may be based on different classification data, although some data may be exploited by more than one cluster. Moreover, the stereotypes belonging to different clusters make predictions on distinct sets of preferences (those significant

PRIMARY-SCHOOL
Classification data:
    Age: up_to_5: 0.0; 6-11: 0.1; 12-14: 0.0; ...; 46-60: 0.0; more_than_60: 0.0;
Preferences:
    Play_activities:
        Importance: 1; *Values*: null: 0.0; low: 0.05; medium: 0.45; high: 0.5;
    Reading_material:
        Importance: 0.8; *Values*: null: 0.05; introductory: 0.2; specialized: 0.05;
            scholastical: 0.7;
    Length_of_visit:
        Importance: 1; *Values*: short: 0.6; medium: 0.35; long: 0.05;
    Background_knowledge:
        Importance: 1; *Values*: low: 0.9; medium: 0.1; high: 0.0;
    ...
Relevance: 1;

Fig. 2. Stereotype describing primary-school children.

from the described viewpoint). For instance, the stereotype described in Figure 2 belongs to the Age_range cluster and describes children aged from 6 to 11, i.e., studying at primary school.

The stereotypes also specify the relevance of the represented group ("Relevance" slot). The relevance predicted by a stereotype S represents the weight that the preferences of a prototypical tourist belonging to S should have on the selection of tourist attractions. This parameter ranges in [0..1], where 0 denotes null relevance and 1 represents the maximum relevance. In order to take into account the fact that some subgroups, such as children, have strong requirements on the organization of a tour, the related stereotypes predict a relevance equal to 1 (e.g., see the "PRIMARY-SCHOOL" stereotype in Figure 2). Instead, most of the other stereotypes have an medium or low relevance.

2.5 Management of subgroups within a tourist group

At the beginning of the interaction, the user visiting the Web site is asked how many people are going to travel together. Then, the system asks her to distribute such people into relevant subgroups, on the basis of a set of pre-defined features. Currently, the main user features which we have considered concern the range of age (to deal with children and elderly) and the mobility and vision capabilities; however, the system can be configured to take into account other features, such as social and cultural aspects. For each subgroup, a registration form is displayed, in order to provide the system with information about the interests of the related travelers, the cardinality of the subgroup, and other similar information. The fields of the forms are not mandatory, but the system's suggestions can be more focused, if more information about each subgroup is provided. Each subgroup model is initialized with the preferences of a very generic traveler, corresponding to an adult with average interests and without special requirements.

This initialization enables the system to have a basic description, in case the direct user does not appropriately fill in the forms.

The system initializes each subgroup model by exploiting stereotypical information about tourists, according to the techniques developed for the SeTA system: the subgroup characteristics (provided by the user in a registration form) are matched against the stereotypical information. The stereotypes best matching the characteristics of the subgroup are then used to make predictions on the subgroup preferences and on its importance. Details about these techniques can be found in [3].

3 Generation of recommendations for a tourist group

3.1 Evaluation of tourist attractions

The evaluation of items for a heterogeneous tourist group is achieved in two steps. First, items are separately evaluated and ranked with respect to each subgroup. Then, the subgroup-related rankings are combined to obtain the overall ranking, from the viewpoint of the whole group. In the following, we will focus on the subgroup-related evaluation of items. Then, in section 3.2, we will discuss how the separate rankings are combined to generate the system's recommendations.

Given the preferences of a homogeneous subgroup, items are ranked by exploiting the same techniques used in the SeTA system for the recommendation of products. As such techniques are extensively described in [3], we only sketch them in the following.

Tourist attractions are represented as entities described by features providing different types of information: e.g., geographical information, category, logistic information, and so forth. In the evaluation of a tourist attraction, the system exploits the *properties* og the item. Properties are features that provide a qualitative evaluation of the attraction. For instance, we take into account the historical or artistic value of an attraction, how much background knowledge is required to appreciate it, or whether the place offers play areas for children ("play_activities" in Figure 2).

The degree of matching (henceforth, *satisfaction score*) between an item and a subgroup model is evaluated by analyzing the preferences of the subgroup towards the properties of the item, stored in the subgroup model. Each property is matched against the related preference to establish an individual score.[2] The overall satisfaction score of the item results from the merge of the individual scores of its properties. Two individual scores, X and Y, are combined by exploiting the following formula:

$$SATISFACTION\_SCORE(score_X, score_Y) =$$
$$score_X * score_Y / (score_X + score_Y - score_X * score_Y)$$

[2] This score is a decimal value in [0..1], where 1 represents perfect compatibility with the preference, while 0 represents total incompatibility. In this evaluation, the importance of the preference in the subgroup model is used to tune the influence of less relevant properties, when they are not compatible with the subgroup preferences.

This formula, described in [2], is additive and therefore supports and incremental evaluation of the overall satisfaction score of an item. The formula takes values in the [0..1] range; moreover, it is particularly selective: being based on the product operator, it returns a 0 satisfaction score for any item having at least one null individual score. This selective power is essential in the tourism domain, as it enables the system to dramatically downgrade the evaluation of items incompatible with basic requirements of the traveler subgroups. For instance, a tourist place without a transportation system suitable for disabled people cannot be recommended as a good solution to a group of tourists with mobility problems.

3.2 Recommendation criteria

INTRIGUE provides the user with alternative recommendation criteria to support her in the selection of the attractions for the tour. The reason for providing different recommendation criteria is that no specific recommendation method can satisfy all the possible requirements. For instance, the user may want to see separate recommendations for each subgroup and compare the lists by themselves. Alternatively, she may prefer to be provided with a single recommendation list, representing a synthesis of the suggestions, in order to avoid the analysis of multiple and possibly long recommendation lists. However, even in this case, different criteria could be applied to generate the list. For instance, items unsuited for at least one subgroup might need to be ignored, although they are interesting options for other groups. Moreover, the recommendations could be fair, trying to satisfy all the subgroups in a uniform way, or they could be biased towards the preferences of the most influent subgroups.

In our system, we have included three recommendation modalities, which the user can choose from by clicking on buttons available in the user interface. See the buttons at the top of Figure 3: “Separate listing by groups”, “Unique listing (method 1)”, “Unique listing (method 2)”. We describe these recommendation criteria referring to a scenario where a user similar to the one described in Figure 1 inspects the civil buildings in Torino. The user is organizing a tour with some children and impaired people. In this case, the system generates three subgroup models: one for the subgroup including the direct user, the others for the two homogeneous subgroups traveling with her.

Separate listing by group. In this modality, the system shows separate lists, one for each subgroup. with items sorted on the basis of the rankings previously evaluated for each subgroup. The best ranked elements for each subgroup are at the top of the lists, while the worst ones are at the bottom.

Figure 3 shows the system’s recommendations, reporting the suggestions for the subgroup including the direct user in the first column, for the children in the second one and for the impaired people in the third one. Notice that each item is associated with an icon (stars), which represents the satisfaction score obtained by the item and supports an easy identification of the best items for each group.

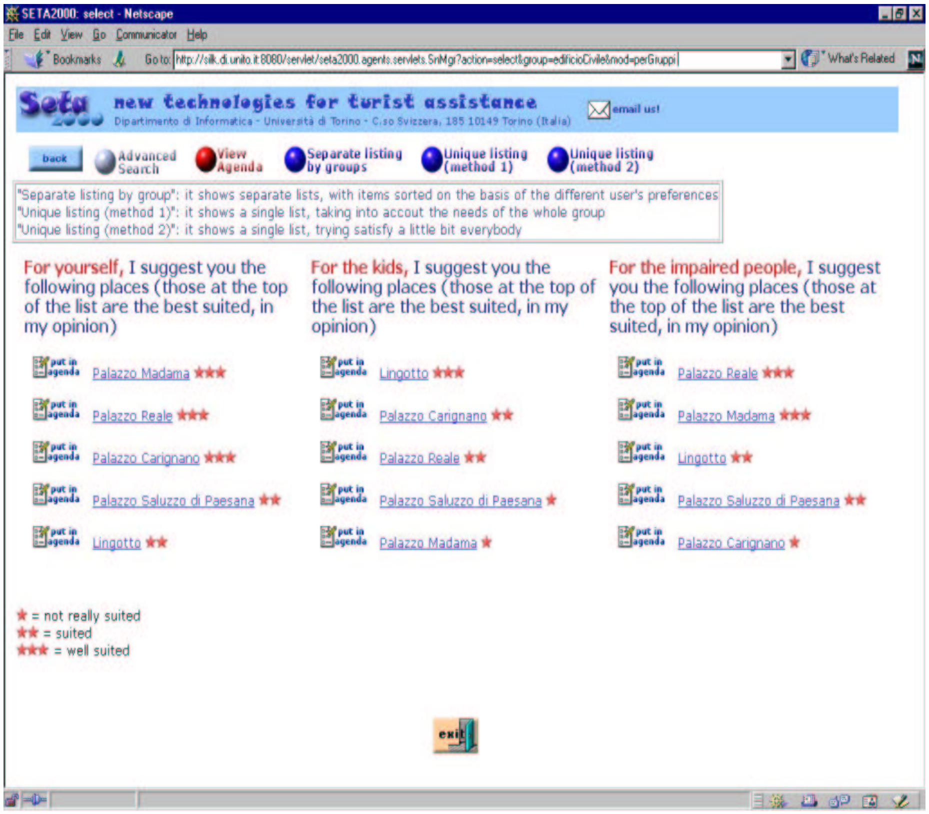

Fig. 3. Separate listing of tourist attractions

This type of recommendation explicitly presents the suggestions for each subgroup and enables the direct user to select the interesting attractions in an informed way, taking into account the ranking of each item in each subgroup-related list. However, it may be confusing if there are too many subgroups to be considered, because the user should compare several permutations of the items, which may result in a certain overhead.

Preferential satisfaction listing (method 1). In this modality, the system displays a single, sorted list of items representing the suggestion for the whole group. The overall ranking of items is obtained by merging all the subgroup-related rankings in a weighted way, depending on the cardinality and the relevance of the various subgroups. Thus, the system's suggestions take into account the presence of large homogeneous subgroups, and that of subgroups with special needs. The overall score S of an item is evaluated by combining the satisfaction scores S_{UM} associated to the item for each subgroup. In our example:

$$S = inf_{direct\_user} * S_{direct\_user} + inf_{children} * S_{children} + inf_{impaired} * S_{impaired}$$

Fig. 4. Unique listing of tourist attractions

For each subgroup, the associated weight (inf_{UM}) represents the influence of the subgroup's preferences within the whole group and is evaluated according to the following formula:

$$inf_{UM} = relevance_{UM} * cardinality_{UM} \; / \; total\_cardinality$$

where $relevance_{UM}$ is the relevance of the subgroup, $cardinality_{UM}$ is the number of its members[3] and $total\_cardinality$ is the total number of tourists forming the overall group. In this way, special subgroups can be privileged in a flexible way, depending on the portion of the overall group they represent.

Figure 4 shows the preferential listing recommendation for the same group of tourists considered in Figure 3. The tourist attractions are sorted according to the overall ranking, resulting from the weighted merge of the individual group rankings. In this case, the stars next to the items represent their overall ranking, related to the whole group of tourists.

This type of recommendation is useful to provide the user with suggestions supporting her in the selection of the items to add to the agenda. In fact, the first

[3] Relevance and cardinality are retrieved from the subgroup model.

items in the list are the best ones for the group considered as a whole, possibly with special attention to the preferences of the most influent subgroups. However, other strategies could be exploited to weight the influence of the subgroups; for example, the group relevance alone could be used, without considering the cardinality, in order to maximize the influence of special subgroups (or, in alternative, the cardinality alone could be used). We are investigating the possibility of providing the direct user with methods to select the parameters to be considered for the evaluation, therefore supporting further personalization forms in the recommendations.

Uniform satisfaction listing (method 2). As an alternative to the preferential satisfaction listing, which sorts tourist attractions in a biased way, the user may be interested in receiving a fair recommendation, where the suggestions uniformly satisfy all the subgroups. To this extent, we have introduced a third modality, where the system shows a single recommendation list whose items are sorted to achieve uniform satisfaction for each subgroup. In this case, the satisfaction score of an item is explicitly used as a *degree of satisfaction* for the related subgroup of tourists: an item with a satisfaction score equal to 1 increases the satisfaction of a subgroup in a maximal way, while an item with a 0 satisfaction score does not modify the satisfaction degree at all. The recommendation list is generated in order to achieve uniform satisfaction for all the subgroups. We do not describe this method in detail, but the idea is the following: at each step, the system selects, out of the set of items to be sorted, the one maximally increasing the satisfaction degree of the subgroup that has received minimum satisfaction in the previous selections. The aim is to raise the group satisfaction degree as much as possible.

3.3 Explanation of the system's suggestions

The explanation capability is desirable for any interactive system. In our case, two types of explanation are particularly important: the first one is the specification of the target of the recommendation. The second one concerns the reasons for suggesting the various items. In INTRIGUE, we have addressed both types of explanation, by supplementing the recommendations with textual descriptions, which explain the main reasons for suggesting an item to a subgroup.

As shown in Figure 4, in the unique listing modalities each item is coupled with a sentence specifying, for each subgroup, the most important properties determining the suggestion of the item.[4] For instance, "Palazzo Reale" (Royal Palace) is a good suggestion for the subgroup including the direct user since it is much eye-catching and has high historical value; moreover it is also good for children because its visit requires low background knowledge. Finally, it suits disabled people because it has no architectural barriers.

[4] The presentation of the most suited properties is not displayed in the separate listing by group due to space constraints on the screen. However, information about items and their properties can be retrieved by asking for their detailed presentation pages.

Notice that the explanations generated by the system for a subgroup do not include the whole list of properties satisfying the preferences in the subgroup model: in fact, they only report the most relevant preferences of the class characterizing the subgroup, in order to produce maximally useful information for the user. For instance, the recommendations for children in Figure 4 are focused on properties such as the background knowledge required to appreciate the attraction, the length of the visit, and so forth. In contrast, they do not mention any property such as being eye-catching because most people (and also children) like eye-catching places, and so this property is not particularly relevant for the specific tourist class addressed in the explanation.

The contextually relevant properties for a subgroup are selected by exploiting the stereotypical information: given the set of properties satisfying the preferences of a subgroup model, the properties to be mentioned are selected by taking into account the importance of the related preferences in the stereotype describing the tourist class characterizing the subgroup. For instance, as shown in Figure 2, children are mostly interested in properties such as the availability of play areas, length of the visit and required background knowledge.

The linguistic form of the explanations are automatically generated by exploiting template-based Natural Language Generation techniques (see [2]), on the basis of a language independent internal representation of the item properties. In this way multilinguality is supported.

4 Related work

The typical tourist information Web sites are static hypertexts and provide non-personalized information about attractions and services available in a town, or in a region. These systems suffer from two major drawbacks: first, they cannot provide users with information focused on specific interests (the only way to search for specific information is typically provided as an embedded search engine). Second, they rely on static descriptions, which become obsolete in a short time and have to be manually revised by the site administrators. Another type of site are the e-travel agencies, like, for instance, Expedia [11]; their main goal is to offer discount airfare, flight, hotel, cars, vacation packages reservations. However, these sites only help the user to gather information or to make reservations and do not help her to organize a trip or a tour of a city.

Some dynamic hypermedia systems have been designed to generate the presentations "on the fly", possibly tailoring contents and styles on the basis of the application of personalization strategies; e.g., see [10, 17] and [14] for an overview. For instance, AVANTI [13] was designed as a kiosk system which generates customized presentations of the services and tourist attractions available in a town. The goal was to support alternative interaction media and personalization strategies were exploited to tailor the presentations to the individual user's interests. More recently, intelligent virtual guides have been designed to personalize the visit of a museum, taking into account several factors such as the user's interests, domain expertise, the fact that the user was visiting the place for the first time, and also the (physical) navigation style within the museum [18, 16]. In particular, there is a strong interest in the development of context-aware applications, supporting a selective presentation of information, based on the physical location of the user [15, 9].

From a related perspective, special attention has been paied to the development of systems supporting the individual user's search for information with personalized recommendations. For instance, see [8, 7, 12]. Finally, some researchers have defined techniques to support the user in the definition of her own search criteria, therefore leading to the configuration of her own information service [6].

5 Discussion

Our work is focused on the provision of multilingual, personalized recommendations for groups of people planning a visit to a given geographical area. Different from on-site kiosks and context-aware applications, the role of our system is in assisting the user to schedule the tour, not in guiding the group during the visit. Therefore, physical context has a marginal role and is exploited only when the schedule of the trip is considered, to estimate, for instance, the appropriate transfer times from one tourist attraction to another.

The main contribution of this paper concerns the management of a *group model*, where the characteristics, interests and preferences of the various components of the group are taken into account to tailor the recommendations in a suitable way. The management of a group model and the group-oriented personalization distinguishes our system from the other recommender systems, which tailor the suggestions to the individual user. In the case of a single user, her preferences may be more or less articulated, but are unique. In contrast, a group of people traveling together may have conflicting preferences (and needs) and the generation of a recommendation which addresses the requirements of all such people is much more complex. In order to address this issue, we have considered the group as composed of subgroups, having homogeneous preferences and needs. Moreover, we have designed different recommendation criteria, which the user can select to get different types of suggestion from the system. These criteria include the separate ranking for subgroups, a uniform merge of the preferences of all the homogeneous subgroups, and a weighted recommendation, where the subgroups have different influence on the system's rankings. Another important aspect is the capability to explain why a recommendation has been made: when the system presents the lists of tourist places to be visited, a sentence specifying the most relevant properties determining the suggestion is generated.

All the Web pages of the INTRIGUE user interface are dynamically generated, accordingly to the following steps: first, the information to be displayed is selected; then, the linguistic descriptions are generated, by exploiting a template-based Natural Language Generator (see [2]) and an XML object, representing the "content" of the page, is produced (see [5]). Finally, the XML object is transformed into a HTML page to be interpreted by a standard browser, by exploiting XSL transformations. Our goal is to have a representation of the personalized content of each page independent from the actual user interface implemented by the system. For instance, the XML object could be fed to a different module, that generates a user interface for a different medium (e.g. a mobile), or stores the personalized content in a database for further processing.

We have not yet focused on the adaptive presentation of tourist attractions. However, the generation module could be easily extended with the techniques developed for the SeTA system, in order select the features of an attraction to be presented, on the basis of the user's interests [2, 3]. We are also working on improving the scheduling facility to realize a virtual "travel-agent" able to help the user to organize her day tour, including hotel and meal arrangements.

INTRIGUE is based on the multi-agent architecture of SeTA, instantiated on the tourism domain and updated with the introduction of the interactive agenda agent. Details about this architecture can be found in [1] and [4].

As our system is still under development, no evaluation has been carried on up to now. However, we plan to test the recommendation and explanation functionalities of the system with users in the next future. In particular, we would like to evaluate the system on some typical scenarios, such as a family visiting a town, a tour to be organized for the a group of students or the organization of a spare day for some people visiting a town for business reasons. In all these cases, we will evaluate the recommendation capabilities of the system by comparing the tourist attractions selected by the user with those suggested by the system. Moreover, we will collect, from a selected number of users, individual feedback on the usefulness and effectiveness of the system's alternative recommendation criteria and explanation capabilities.

References

1. L. Ardissono, C. Barbero, A. Goy, and G. Petrone. An agent architecture for personalized web stores. In *Proc. 3rd Int. Conf. on Autonomous Agents (Agents '99)*, pages 182–189, Seattle, WA, 1999.
2. L. Ardissono and A. Goy. Dynamic generation of adaptive Web catalogs. In *Lecture Notes in Computer Science n. 1892: Adaptive Hypermedia and Adaptive Web-Based Systems*, pages 5–16. Springer Verlag, Berlin, 2000.
3. L. Ardissono and A. Goy. Tailoring the interaction with users in web stores. *User Modeling and User-Adapted Interaction*, 10(4):251–303, 2000.
4. L. Ardissono, A. Goy, G. Petrone, and M. Segnan. A software architecture for dynamically generated adaptive web stores. In *Proc. 17th IJCAI*, page to appear, Seattle, WA, 2001.
5. L. Ardissono, A. Goy, G. Petrone, M. Segnan, and P. Torasso. Dynamic generation of personalized tourist information on the web. In *Proc. of 12th ACM Conference on Hypertext and Hypermedia (PERTEXT 2001)*, page to appear, Aarhus, Denmark, 2001.
6. M. Bauer and D. Dengler. InfoBeans - configuration of personalized information assistants. In *Proc. 1999 Int. Conf. on Intelligent User Interfaces (IUI'99)*, pages 153–156, Los Angeles, 1999.
7. D. Billsus and M. Pazzani. A personal news agent that talks, learns and explains. In *Proc. 3rd Int. Conf. on Autonomous Agents (Agents '99)*, pages 268–275, Seattle, WA, 1999.

8. R.D. Burke, K.J. Hammond, and B.C. Young. The FindMe approach to assisted browsing. *IEEE Expert*, pages 32–39, 1997.
9. K. Cheverest, N. Davies, K. Mitchell, and P. Smith. Providing tailored (context-aware) information to city visitors. In *Proc. International Conference on Adaptive Hypermedia and Adaptive Web-based Systems (AH 2000)*, pages 73–85, Trento, Italy, 2000.
10. R. Dale, S.J. Green, M. Milosavljevic, and C. Paris. Dynamic document delivery: Generating natural language texts on demand. In *Proc. 9th Int. Conf. and Workshop on Database and Expert Systems Applications (DEXA'98)*, Vienna, 1998.
11. Expedia.com. Expedia. http://www.expedia.com, 2001.
12. J. Fink and A. Kobsa. A review and analysis of commercial user modeling servers for personalization on the World Wide Web. *User Modeling and User-Adapted Interaction, Special Issue on Deployed User Modeling*, 10(2-3):209–249, 2000.
13. J. Fink, A. Kobsa, and A. Nill. Adaptable and adaptive information for all users, including disabled and elderly people. *New review of Hypermedia and Multimedia*, 4:163–188, 1999.
14. A. Kobsa, J. Koenemann, and W. Pohl. Personalized hypermedia presentation techniques for improving online customer relationships. *The Knowledge Engineering Review*, page to appear, 2001.
15. P. Marti, A. Rizzo, L. Petroni, G. Tozzi, and M. Diligenti. Adapting the museum: a non-intrusive user modeling approach. In *Proc. 7th Int. Conf. on User Modeling*, pages 311–313, Banff, Canada, 1999.
16. L. Marucci and F. Paternò. Logical dimensions for the information provided by a virtual guide. In *Proc. International Conference on Adaptive Hypermedia and Adaptive Web-based Systems (AH 2000)*, pages 359–362, Trento, Italy, 2000.
17. M. Milosavljevic, R. Dale, S.J. Green, C. Paris, and S. Williams. Virtual museums on the information superhighway: Prospects and potholes. In *Proc. In Proc. Annual Conference of the International Committee for Documentation of the International Council of Museums (CIDOC'98)*, Melbourne, 1998.
18. D. Petrelli, A.De Angeli, and G. Convertino. A user centered approach to user modelling. In *Proc. 7th Int. Conf. on User Modeling*, pages 255–264, Banff, Canada, 1999.
19. R. Wilkinson, S. Lu, F. Paradis, C. Paris, S. Wan, and Mi Wu. Generating personal travel guides from discourse plans. In *Lecture Notes in Computer Science n. 1892: Adaptive Hypermedia and Adaptive Web-Based Systems*, pages 392–85. Springer Verlag, Berlin, 2000.