

APPLIED ARTIFICIAL INTELLIGENCE: Special Issue on Artificial Intelligence for Cultural Heritage and Digital Libraries 17(8-9), pp. 687-714, 2003.

INTRIGUE: PERSONALIZED RECOMMENDATION OF TOURIST ATTRACTIONS FOR DESKTOP AND HANDSET DEVICES

Liliana Ardissono
Dipartimento di Informatica, Università di Torino
Corso Svizzera 185, 10149 Torino, Italy
e-mail: liliana@di.unito.it

Anna Goy
Dipartimento di Informatica, Università di Torino
Corso Svizzera 185, 10149 Torino, Italy
e-mail: goy@di.unito.it

Giovanna Petrone
Dipartimento di Informatica, Università di Torino
Corso Svizzera 185, 10149 Torino, Italy
e-mail: giovanna@di.unito.it

Marino Segnan
Dipartimento di Informatica, Università di Torino
Corso Svizzera 185, 10149 Torino, Italy
e-mail: marino @di.unito.it

Pietro Torasso
Dipartimento di Informatica, Università di Torino
Corso Svizzera 185, 10149 Torino, Italy
e-mail: torasso@di.unito.it

Abstract

This paper presents INTRIGUE, a prototype tourist information server that presents information about the area around Torino city, on desktop and handset devices. This system recommends sightseeing destinations and itineraries by taking into account the preferences of heterogeneous tourist groups (such as families with children and elderly) and explains the recommendations by addressing the group members' requirements. Moreover the system provides an interactive agenda for scheduling the tour. The services offered by INTRIGUE rely on user modeling and adaptive hypermedia techniques; furthermore, XML-based technologies support the generation of the user interface and its adaptation to Web browsers and WAP minibrowsers.

INTRODUCTION

The development of Web-based tourist guides is challenged by the variety of user needs they have to satisfy. Users are interested in different types of attractions: e.g., artistic landmarks, or natural parks. Moreover, as most people do not travel alone, group preferences (in addition to individual user preferences) have to be modeled and possibly conflicting requirements must be taken into account in the recommendations. Furthermore, tourist guides can use wireless communication to offer any-time availability, by adapting to device-dependent features (Bergman 1999). To satisfy these requirements, the services offered by these systems and their user interfaces have to be tailored to heterogeneous users and devices (Riecken 2000, Kobsa et al. 2001, Brusilovsky and Maybury 2002).

The present paper describes the recommendation and user interface generation facilities offered by INTRIGUE (INteractive TouRist Information GUiDE), a prototype tourist information server providing personalized information about tourist attractions in a restricted geographical area. The system integrates a number of Artificial Intelligence techniques for developing innovative solutions to the following problems:

- Recommendation of items tailored not only to individuals, but also to user groups.
- Explanation of the recommendations.

- Generation of device-dependent layouts and information contents, starting from the same internal representation of information.
- Support for advanced problem solving techniques (for tour scheduling).

INTRIGUE dynamically generates a multilingual tourist catalog and recommends sightseeing destinations and itineraries by taking into account the preferences of heterogeneous tourist groups, such as families with children and elderly. In order to increase the user's trust and to help her choose the preferred destinations, the recommendations are grounded by explanations addressing the possibly conflicting requirements of the group members. The system also offers an interactive agenda that helps the user to schedule a tour complying with her visiting preferences and with other constraints, such as the opening times of attractions. The recommendation activity is based on a declarative representation of the knowledge about tourist attractions and on the application of fuzzy evaluation functions for ranking the items. Moreover, the device-dependent user interfaces are generated by dynamically producing the content to be presented and by applying a standard, XML-based approach for its surface generation.

The rest of this paper is organized as follows: the following section provides an overview of our system. Section "Modeling data about tourist attractions" describes modeling issues concerning information about tourist attractions and services, which is the basis for the dynamic generation of the tourist catalog. The subsequent section describes the representation of tourist groups and Section "Personalized recommendation of attractions" presents the group recommendation and the explanation techniques applied in our system. Section "User interface" outlines the techniques for the generation of the device-dependent user interface. The last two sections discuss some related work and conclude the presentation.

OVERVIEW OF INTRIGUE

Browsing the tourist catalog

Although the primary goal of our system is the generation of personalized recommendations, a significant amount of attention has been devoted to eliciting information about the characteristics of the tourist attractions the user is most interested in. Thus, the system generates a tourist catalog that can be browsed according to different criteria:

- The user can search for tourist attractions and accommodations according to different viewpoints, such as the historical period, artistic current, types of monuments, and so forth.
- A geographical search enables the user to constrain the area of interest for the tour, starting from the whole region (in our prototype, Piedmont - Piemonte), to various sub-areas; for example, Torino city.

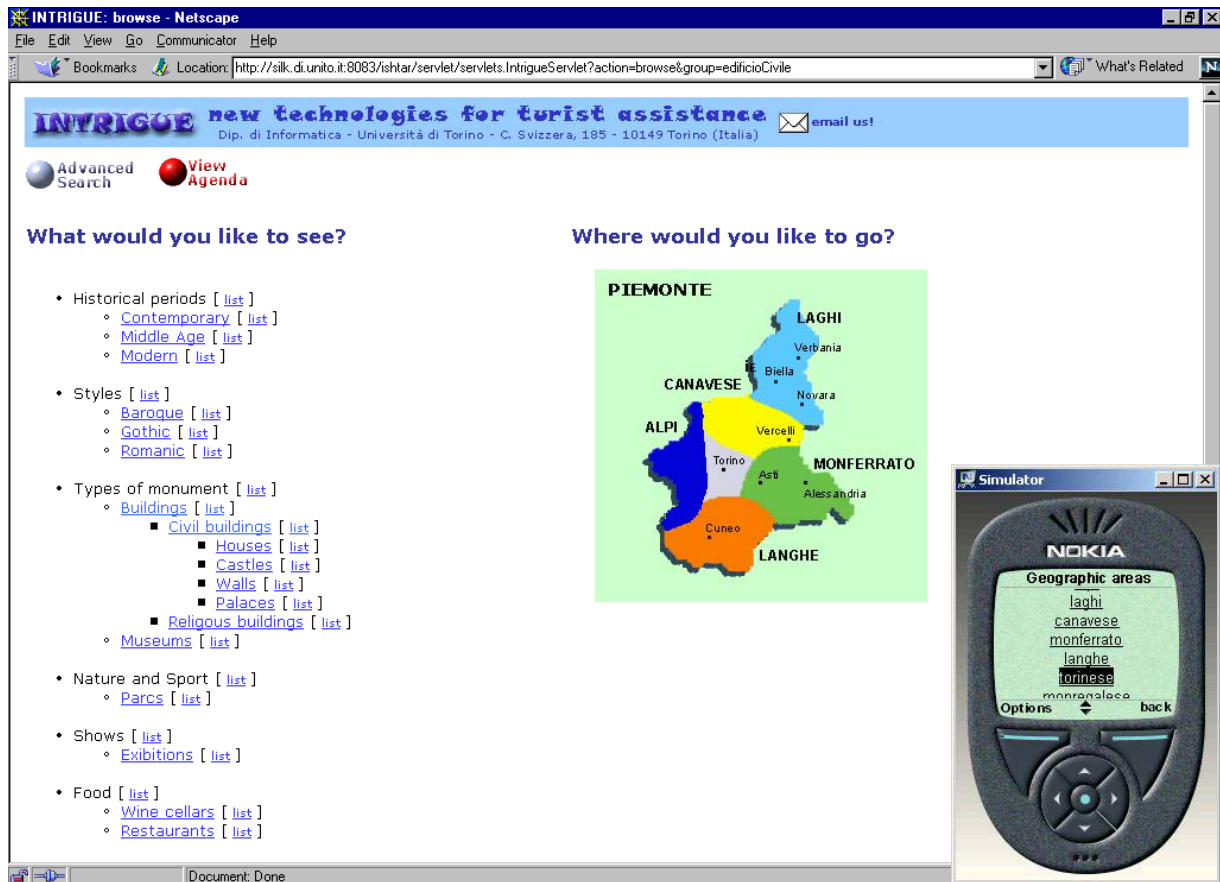


Figure 1: Browsing the tourist information catalog via Web browser and WAP minibrowser.

These search criteria are implemented as pre-compiled queries, which specify conditions on the features of the items associated to the corresponding concepts. As several features are used within the selection criteria, the same tourist attraction can be retrieved by following different browsing paths. The selection criteria are presented in the “What would you like to see?” and “Where would you like to go?” columns of the page shown in Figure 1.¹ The user can browse the tourist catalog by expanding the nodes she wants to inspect: in the figure, the “Buildings” node has been expanded two levels down. Moreover, the user can select the desired geographical area by clicking on a map, or by following hypertextual links, if she is using a handset device, as shown in the right portion of the figure. The two criteria can be combined to narrow down the search: for instance, civil buildings in Torino can be searched for. The items satisfying the search query can be retrieved by following the “list” link besides the desired node, after having selected the geographical area of interest. When the user follows that link, the system combines the query underlying the selected criterion with the one specifying the selected geographical area and retrieves the corresponding records of items from the tourist database.

Presentation of tourist attractions

Tourist attractions are presented in pages dynamically generated out of the information about the items stored in a database. The generation is based on strategies, represented as rules, that rely on metalevel information about the features of itemsⁱⁱ to choose their presentation style. For example, as shown in Figure 2, the features of the Sacra di San Michele are grouped in different sections depending on the type of information they carry: we have introduced a classification of features in information types, such as essential, technical, geographic information, and so forth.



Figure 2: Presentation of a tourist attraction for Web browser and minibrowser.

These strategies also support the generation of device-dependent pages, by taking the size of the screen and other characteristics of the target device into account. In particular, the presentation of attractions for Web browsers is structured in a single page, where information pieces are grouped according to their type. For instance, the top of the page presenting the Sacra di San Michele for Web browser shows a picture of the building and a “description” paragraph based on its properties, that provide a qualitative assessment of the item features (“For Sacra di San Michele, the visit is quite short, it has medium historical value, ...”). Below, beside the opening hours (essential

information), the “information” block displays basic information features, such as the price, followed by a separate list of its main “characteristics”. These characteristics include historical information, such as its artistic current, category and geographical area. In contrast, the presentation of the Sacra for minibrowsers requires more than one page, due to the hard space constraints of their screens, and the main presentation page only includes the essential information. The right portion of Figure 2 shows the main presentation page for minibrowser, which displays the opening hours; the figure also shows the page displaying the description paragraph.

It should be noticed that, thanks to the feature classification and the use of explicit presentation strategies, the same generation module can be used for the production of pages for different devices. Although we have only considered desktop and (black and white) WAP-phone displays, our system can be extended to handle other types of displays, by extending the set of presentation strategies. Moreover, the presentations could be focused on different features of attractions, depending on the user’s interests; see (Ardissono and Goy 2000).

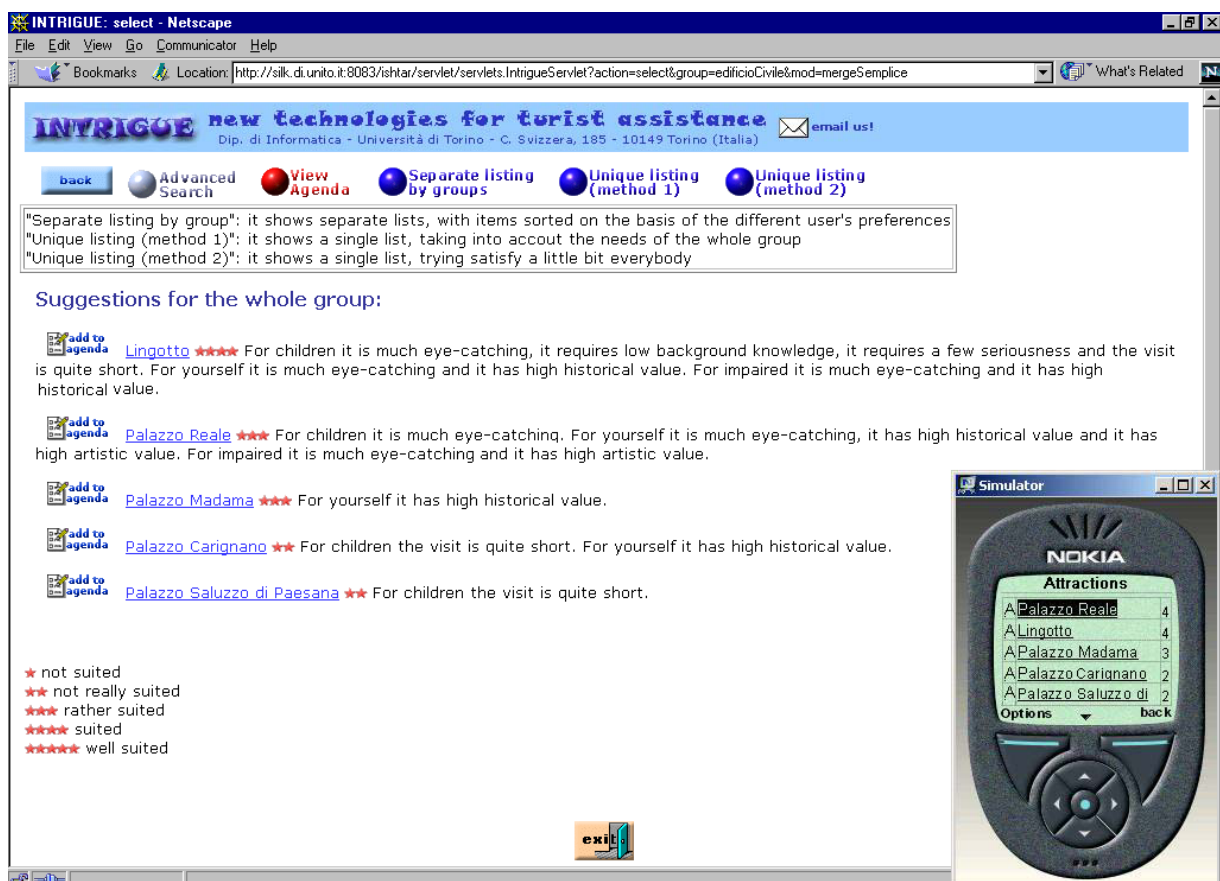


Figure 3: Recommendation of attractions for desktop and handheld interfaces.

Recommendation of attractions

When the user asks for the list of attractions corresponding to a certain selection criterion, the system displays a recommendation list ranked according to the travelers' preferences. The system tailors its suggestions to the possibly conflicting preferences of heterogeneous tourist groups. For instance, suppose that a user having human science background organizes a tour with some children and impaired people. If she searches for the civil buildings in Torino city area, the system displays the page shown in Figure 3, where her preferences are addressed, as well as those of the other participants. The presentation is a bit different for Web browsers and WAP minibrowsers, but each entry in the recommendation list includes: a button to add the item to the agenda storing the places selected for the visit ("add to agenda" / "A"); the name of the item, linked to its presentation page; the evaluation of the item with respect to the group preferences (represented as list of stars, or by a number in [1..5]); a short explanation of the reasons for the system's suggestions (only for Web browser). For example, the following explanation is associated to item Lingotto: "For children, it is much eye-catching, it requires low background knowledge... For yourself it is much eye-catching and it has high historical value. For the impaired ...". The recommendations can be customized according to other criteria, described later on. The user can select the desired recommendation criterion by clicking on suitable buttons; see "Separate listing by groups" and "Unique listing (method 1)" at the top of Figure 3.

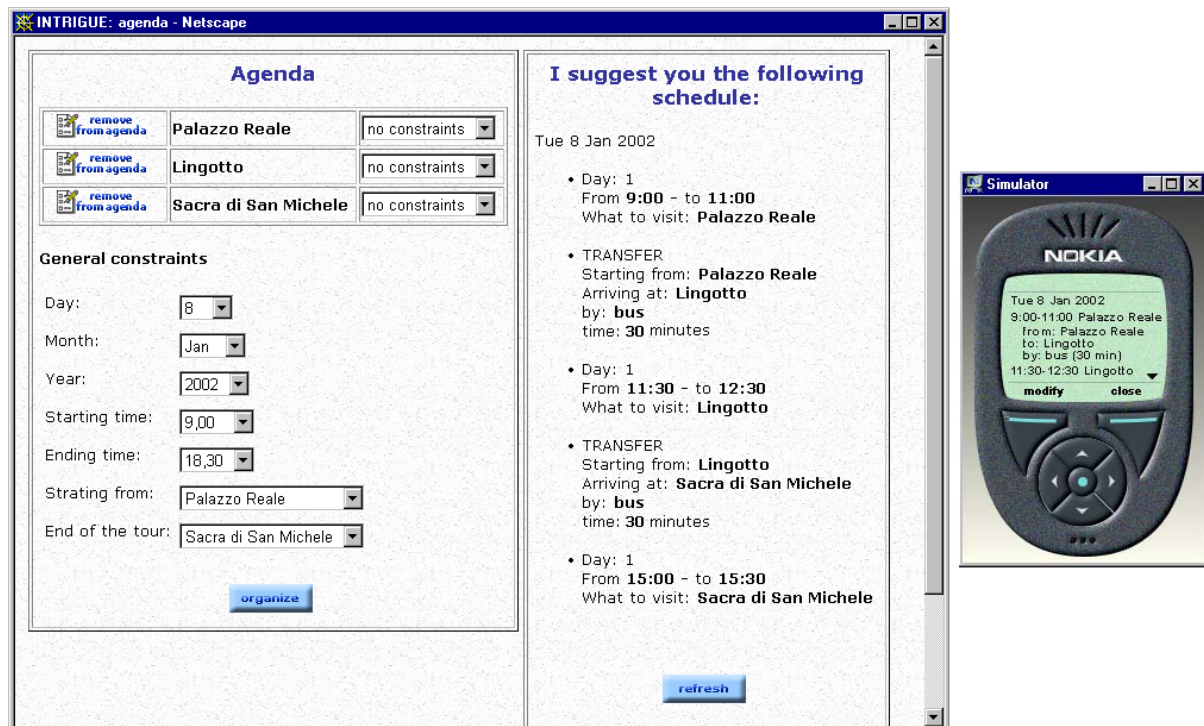


Figure 4: Specification of constraints and tour presentation in the agenda for desktop and handheld interfaces.

Tour scheduling

An interactive agenda facilitates the organization of itineraries and can be opened by clicking on the “View Agenda” button at the top of the recommendation pages; see Figure 3. Given the set of attractions selected by the user, the agenda supports the specification of constraints such as day of visit, arrival/departure time and location, preferred time of visit (morning or afternoon). On the basis of the user’s choices, the system generates a suitable itinerary.

In the presentation of itineraries, the system specifies, for each attraction, the estimated transfer time to reach the attraction from the previous location, the time of the visit and the expected duration. For instance, Figure 4 presents an itinerary for visiting Palazzo Reale, the Lingotto building and the Sacra di San Michele, in a situation where the user asked to start the visit from Palazzo Reale and to end it at the Sacra. The presentation is different for desktop and handheld interfaces: in the former case, the agenda is shown in a single page. In the latter, the information is shown in separate pages, presented in a compact style. If the system fails in scheduling some attractions, it provides information about the violated constraints. The agenda enables the user to save an itinerary for later consideration, supporting the incremental organization of tours and the any-time access to itineraries.

System architecture

The INTRIGUE system has been implemented by exploiting the Seta2000 infrastructure (Ardissono et al. 2001), which supports the development of multi-agent Web-based systems. The Seta2000 architecture is three-tiered; the first tier runs on standard Web browsers and WAP minibrowsers (the handset user interface is currently accessed by using the Nokia WAP-phone simulator). The bulk of the system resides in the middle tier and runs on a Unix environment: the communication with the Web is supported by the Apache HTTP Server and Servlets are used to track the interaction with the user. The third tier includes the databases storing data about products and users.

MODELING DATA ABOUT TOURIST ATTRACTIONS

In INTRIGUE, there is a clear distinction between the information about tourist attractions, which is stored in a database, and the conceptual model of the domain. This distinction plays a critical role since the items in the database are strictly dependent on the specific geographical area the system has been developed for, while the conceptual model is general for tourist attractions and is used to provide “semantic” criteria for accessing the items and for deciding which type of information has to be displayed to the user under different contexts.

<pre>Record 1: Name: Sacra di San Michele; Code: I00113; Category: abbeys; Artistic current: Romanic; Visit lenght: 30; Phone: 011 939130; Opening hours: <Tue, {9.30-12.30, 15.00-16.00}> <Wed, {9.30-12.30, 15.00-16.00}> Price: 8000; ... Background knowledge: medium; Historical value: medium; ...</pre>	<pre>Feature 1: Name: Artistic current; Type: specific characteristic; Domain: {Baroque, Romanic, Gothic}; Importance: 1; Feature 2: Name: Price; Type: basic information; Domain: integer; Importance: 1; Measure unit: Lit; Feature 3: Name: Historical value; Type: property; Domain: {null, low, medium, high}; Importance: 1;</pre>
(a)	(b)

Figure 5: (a) Representation of an item. (b) Representation of some features.

For each individual attraction, a record in the database stores its characteristics and properties as a set of feature/value pairs. Figure 5 (a) sketches the description of the Sacra di San Michele, which is a Romanic abbey (see the “Category” and “ Artistic current” features). The database record reports phone number and opening hours: the Sacra is open all days, except for Mondays, from 9.30 to 12.30 and from 15.00 to 17.00. Moreover, it usually takes 30 minutes for the visit. In addition, some properties are reported: the attraction requires medium background knowledge to be appreciated and has medium historical value.

Since INTRIGUE has to deal with a variety of tourist attractions, described as different categories, the type and the number of features associated to each individual item in the database can vary. The “Category” slot determines which features have to be associated with each item. Moreover, the values of “Category” can be organized into a set of taxonomies as is apparent by looking at “type of monument” selection criterion reported in Figure 1. In order to provide semantic meaning to the description of individual tourist attractions, each feature is described at the metalevel by means of a structured description involving the following slots:

- *Name*: this slot specifies the internal name of the feature.
- *Type*: this slot specifies the type of information provided by a feature (see below).
- *Domain*: this slot represents the set of values that a feature can take (i.e., its value restriction). This set may be a finite domain of qualitative values, an open domain (such as Integer), a range of numerical values, but it may also be structured. For instance, multiple time intervals are needed to describe the opening time of attractions.
- *Importance*: this slot specifies to which degree the feature represents a mandatory piece of information in the description of a category of attractions; see the “salience” of features introduced in (McCoy 1989).

- *Measure unit*: this is the (optional) type of information which the feature values refer to; e.g., a specific currency, minutes, days or years, and so forth.

While the metalevel description of features is used by INTRIGUE for specifying the meaning of the information in the database and for displaying such data in a rational way to the users, some additional knowledge is necessary for specifying the role of each feature within the system. In fact, INTRIGUE has to know which features have to be used as selection criteria in browsing, which ones are related to geographical information and which ones are used for matching preferences and interest of the users. For the sake of flexibility, we have selected an explicit representation of the role of the features, by introducing the “type” slot in their structural description. This slot specifies which of the following classes the feature belongs to: *Geographic features* (such as the country and region), *Essential information* (such as the opening hours), *Basic information* (such as the price), *Specific characteristics* (such as the category a monument belongs to or its historical period) and *Properties* (e.g., historical or artistic value). INTRIGUE uses this type of information for selecting the features to be displayed (for example, the “essential information” features are included in the presentation regardless it is done on a wired or wireless device), for specifying the selection criteria (e.g., features belonging to “geographic information” or to “specific characteristics”), for identifying the features to be used in the evaluation of items with respect to the users’ needs, interests and preferences (the features used for such a role are of type “property”). Figure 5 (b) shows the definition of three sample features: the “Artistic current” is a specific characteristic, takes value in set {Baroque, Gothic, Romanic} and has no measure unit. The “Price” is a basic information feature, takes value in the Integer domain and is measured in Italian Lira. The “Historical value” is a property and takes value in the {null, low, medium, high} set.

Characteristics:

Age: from_46_to_55;
 Background: human_science
 Interest_for_arts: yes;
 Interest_for_history: yes;
 Interest_for_science: no;

Preferences:

Being eye_catching:
Importance: 1; *Values*: null; 0.0; low: 0.1; medium: 0.3; high: 0.6;
 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*: null: 0.0; low: 0.05; medium: 0.1; high: 0.85;
 Artistic_value:
Importance: 1; *Values*: null: 0.0; low: 0.05; medium: 0.35; high: 0.6;
 Scientific_value:
Importance: 0.0; *Values*: null: 0.0; low: 0.1; medium: 0.3; high: 0.6;

Documentation:
Importance: 0.0; *Values*: missing: 0.2; scholastic: 0.2; generic: 0.2; specialized: 0.4;
 ...
 Group Information:
 Cardinality: 5;
 Relevance: 0.4;

Figure 6: An example subgroup model.

MODELING HETEROGENEOUS TOURIST GROUPS

We aim at achieving a compact representation of a heterogeneous tourist group, where individual participants are not described one by one but their possibly conflicting preferences are separately represented. Thus, we model the group as a set partitioned into a number of homogeneous subgroups, whose members have similar characteristics and preferences.ⁱⁱⁱ This type of representation preserves the peculiarities of the subgroups and supports an efficient evaluation of the overall group preferences, which usually does not need to be performed at the granularity level of the individual members. Moreover, to support a flexible generation of recommendations, we model subgroups that may have different degrees of influence on the estimation of the group preferences. For instance, a subgroup could be particularly influent either because it represents a large portion of the tourist group, or because its members belong to a relevant tourist class: e.g., they have special needs, such as children and disabled people.

In the following, we briefly describe the structure of the subgroup models and the type of information used to manage them. INTRIGUE exploits a structured group model, where subgroups are represented as distinct entities. A subgroup model is associated to each *homogeneous subgroup* of people participating to the tour. Each subgroup model is structured in three portions:

- The *Characteristics* store the socio-demographic information about the participants; e.g., Figure 6 represents a subgroup of people aged between 46 and 55, with human science background and interested in arts and history.
- The *Preferences* specify the system's estimates of the subgroup requirements on the properties of tourist attractions. It is worth noting that there is a one-to-one mapping between preferences and properties used for describing tourist attractions, so that for each preference P , we can associate the corresponding property P' . For the sake of simplicity, we use the same name to identify the preference and the corresponding property. Each preference is represented as a slot and has two facets:
 - The *Importance* specifies the strength of the subgroup preference, i.e., it denotes how strongly the preference influences the subgroup's evaluations of items. For instance, the travelers represented by the

model shown in Figure 6 have no interest in special transportation services (the importance is 0), but they are highly interested in the historical value of attractions (importance = 1).

- A *probability distribution* on the values of the preference specifies the likelihood that the subgroup prefers attractions having the specified property values. Suppose that a set of attractions $A = \{A_1, \dots, A_n\}$ are described in terms of the property P' . Each attraction A_i will assume value v_i for the property P' . If, in the group model, the probability distribution specifies that $p(P=v) = x$, x is the likelihood that the subgroup prefers attraction A_i , whose property P' assumes the value v in A_i . For instance, the tourists described by the group in Figure 6 very likely prefer attractions having high historical value: $p(\text{Historical_value}=\text{high})=0.85$.

The *Group Information* stores a *cardinality*, specifying the number of group members, and a *relevance*, representing an estimate of the weight that the preferences of a member should have on the recommendation of tourist attractions. The relevance ranges from 0 (null relevance), to 1 (maximum one). In our example, the subgroup includes 5 people and has medium relevance (0.4).

Each subgroup model is automatically generated by the system by combining information elicited from the user with stereotypical knowledge about the typical tourist classes considered in the domain (Rich 1989). At the beginning of the interaction, the user is asked how many people are going to travel together. The system also asks her to distribute such people in subgroups, on the basis of their characteristics (age, background, interests, etc.). For each subgroup, a registration form is displayed to provide the system with information about the subgroup members (the fields of the forms are not mandatory).^{iv} On the basis of the pieces of information provided about a subgroup, the subgroup model is initialized by matching its characteristics against the stereotypical information and the best matching stereotypes are used to predict the subgroup preferences and its relevance, obtaining a model similar to the one shown in Figure 6; for details, see (Ardissono and Goy: 2000b). It is worth noting that the stereotypes have been defined by analyzing the tourist population under different perspectives: different stereotypes do not only describe tourists with different interests, but also different classes of capabilities; e.g., visually or physically impaired.

PERSONALIZED RECOMMENDATION OF ATTRACTIONS

The recommendation of items to a user is complex by itself, but we have the additional problem of recommending items to heterogeneous groups. For this reason, the evaluation of items is achieved in two steps. First, items are separately ranked by taking the preferences of each subgroup into account; then, the subgroup-related rankings are

combined to obtain the ranking suitable for the whole group. This two-phase approach supports the adoption of alternative recommendation policies, which can be applied to the same set of subgroup related rankings.

Evaluation of tourist attractions

The suitability of a tourist attraction to the preferences of a subgroup (henceforth, *satisfaction score*) is estimated by matching the properties of the item against the subgroup preferences. Our evaluation model takes the importance of the properties into account to support a flexible ranking behavior: an attraction matching the important preferences is recommended to the subgroup even if it fails to match some less important ones. In contrast, as very important properties represent necessary requirements, an attraction failing to match at least one important preference cannot be recommended. This approach reflects the requirements of the tourist domain, where hard constraints have to be modeled. For instance, a tourist place without a transportation service suitable for disabled people cannot be recommended to a group of tourists with mobility problems, although it satisfies their other interests.

Evaluation of the suitability of an item

In the evaluation of the satisfaction score of an item, each item property is matched against the related preference to establish a property score. This score is a decimal value in $[0..1]$, where 1 represents perfect compatibility with the preference, while 0 represents total incompatibility.

The satisfaction score of the item results from the merge of the property scores. Two property scores, X and Y , are combined by means of the following formula. See (Ardissono and Goy 2000b) for a thorough description:

$$SATISFACTION_SCORE(score_X, score_Y) = score_X * score_Y / (score_X + score_Y - score_X * score_Y) \quad \text{(i)}$$

This formula takes values in $[0..1]$. It is additive and therefore supports an incremental evaluation of the satisfaction score. Moreover, it is particularly selective, as it implements a fuzzy AND of the property scores: being based on the product operator, it returns a 0 satisfaction score for any item having at least one null property score. This selective power enables the system to downgrade the evaluation of items incompatible with the basic requirements of some subgroups; in contrast, as 1 is the neutral value of the product, the properties whose score is 1 do not modify the satisfaction score. As described in the related work section, we selected this formula in the place of alternative ones, such as those deriving from the Multi-Attribute Decision Theory, for its stronger selective power.

Evaluation of a property score

Given an item, the property score of a property Q (e.g., “Historical value”), is evaluated by matching the value of Q in the item description (e.g., “high”) with the corresponding preference in the subgroup model. This match depends

on the probability that the subgroup prefers items suiting that specific value. Moreover, as weak requirements should have weak influence on the subgroup evaluations of attractions, the match also depends on the importance of the preference for the subgroup.^v Property scores range in [0..1], where 1 represents perfect match with the subgroup preferences and 0 denotes total mismatch. This means that the properties whose value is preferred by the subgroup with very high probability should receive a score near 1, moreover, irrelevant properties should receive a similar score, even if their value is preferred by the subgroup with low probability. More specifically, consider an item and its property Q and suppose that the value fitting the item is q_i . Suppose also that the subgroup preference for Q has importance Imp_Q and that the probability of q_i is p_{q_i} . The property score of Q is computed as follows:

$$score_Q = Imp_Q * p_{q_i} + (1 - Imp_Q) \quad \text{(ii)}$$

This function raises the score of the less important properties: e.g., the score of a property related to a preference with 0 importance is 1, for any possible value p_{q_i} . In contrast, the score of the important properties only depends on the probability distribution in the subgroup model (if $Imp_Q=1$, then $score_Q = p_{q_i}$). This behavior is coherent with the strategy adopted to evaluate the satisfaction score of the item. Since that evaluation is based on a fuzzy AND of the property scores concerning the item properties, a property with low importance should not influence too much the satisfaction score. Therefore, its property score should be close to 1 (the neutral value for fuzzy AND).




Figure 7: Evolution of the properties scores on the basis of their importance.

Figure 7 shows the impact of the importance in the evaluation of the property score of two sample properties, $P1$ (represented by the solid line) and $P2$ (dashed line). In this example, the value of $P1$ has probability 0.8 in the subgroup model, while the probability for $P2$ is 0.1.

PROPERTY	Item1: Prop_val	Item1: Prop_score	Item2: Prop_val	Item2: Prop_score
Eye_catchiness	high	0.6	high	0.6
Special_transportation_systems	missing	1.0	some	1.0
Special_facilities_for_vision	present	0.6	missing	0.28

Historical value	high	0.85	high	0.85
Artistic value	medium	0.35	high	0.6
Scientific value	null	1.0	high	1.0
Documentation	scholastic	1.0	missing	1.0

Figure 8: Evaluation of the property scores of two tourist attractions.

Example

Consider the preferences stored in the subgroup model shown in Figure 6 and two sample items (Item1 and Item2), whose property values are reported in Figure 8. The first column of the table shows the properties, the second and third columns show, respectively, the property values fitting the Item1 and the resulting property scores; the last two columns report property values and property scores concerning Item2. As an example, we first compute the property scores for the two items (step 1); then, we compute the satisfaction score of the items (step 2).

1. For each item, the property scores are evaluated by matching the property values fitting the item with the related subgroup preferences. This match depends on the importance of the properties. For instance, although the two items differ in their scientific value (Item1 has no scientific value; Item2 has high scientific value) they get the same property score, 1, because the subgroup does not care about the corresponding preference: the importance of “Scientific_value” is 0. Let’s consider the property score of the “Special_facilities_for_vision” property, for Item2: the importance of the preference is equal to 0.8; moreover, the probability for the “missing” value is 0.1. Therefore, the property score is evaluated as: $0.8 * 0.1 + (1 - 0.8) = 0.08 + 0.2 = 0.28$. Finally, the property score of the properties whose importance is 1 (“Eye_catchiness”, “Historical_value”, etc.) coincides with the probability of the property values in the subgroup model.
2. The satisfaction scores of the two items are evaluated by applying formula (i) to their property scores; e.g., the satisfaction score for Item1 is 0.5397. The application of this formula builds on the following idea: at the beginning of the evaluation, we assume that the item perfectly matches the subgroup model, i.e., the satisfaction score of the item is equal to 1. However, as soon as we encounter a property whose score is lower than 1, the satisfaction score downgrades. As the (i) formula is based on the fuzzy AND, a null property score sets the satisfaction score to 0, so that the item results completely unsuitable for the subgroup.

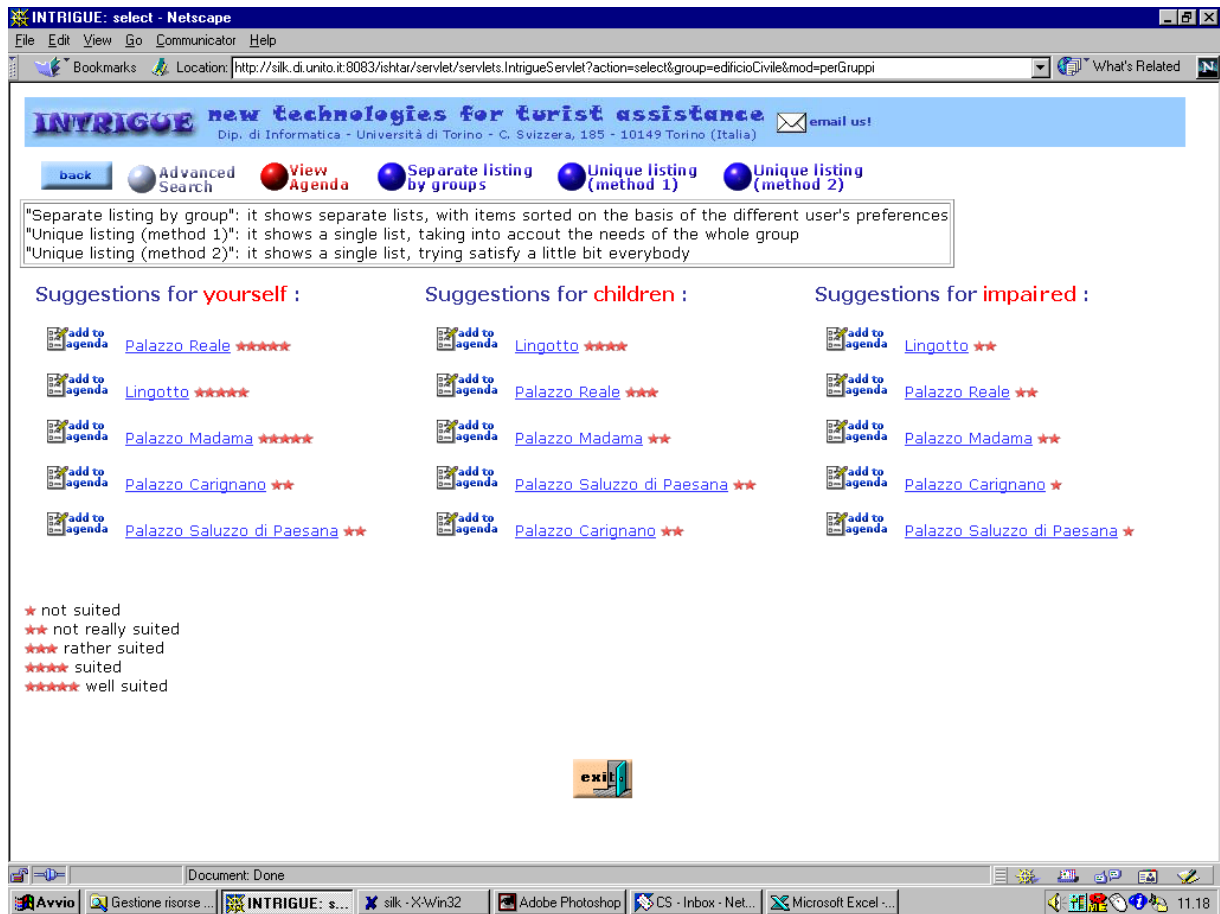


Figure 9: Separate recommendation listings by group.

Recommendation of items

We present two recommendation modalities offered by INTRIGUE: the separate listings and the unique listing.

In the *separate listing by group* modality, the system shows separate recommendation lists, one for each subgroup forming the tourist group. Each list is sorted on the basis of the rankings achieved by the attractions in their evaluation, given the preferences of the corresponding subgroup. Figure 9 shows a separate listing by group recommendation for a user with human science background organizing a tour in Torino city with some children and impaired people. The page shows the suggestions for the subgroup including the direct user in the first column, for the children in the second one and for the impaired in the third one. For each subgroup, the stars associated to an item qualify its satisfaction score for the subgroup. In this way, the best items are easily identified and the user can select the places to visit, by separately considering the subgroup points of view.

In the *unique listing (method 1)* recommendation modality, the system displays a sorted list of items representing the recommendations for the whole group. The overall suitability of an attraction depends on its suitability to the

subgroups and is computed by combining the subgroup-related evaluations. We decided to experiment a flexible evaluation model that supports the introduction of a bias towards the preferences of the most relevant subgroups within the tourist group. This bias is achieved by combining the subgroup-related satisfaction scores in a weighted way, where each weight represents the relevance of the corresponding subgroup. If the tourist group is composed of subgroup models UM_1, \dots, UM_n , the overall score S of an item is evaluated by combining the satisfaction scores S_{UM_i} associated to the item for each subgroup as follows:

$$S = \sum_i w_{UM_i} * S_{UM_i} \quad \text{(iii)}$$

Although different types of bias can be defined, we want to focus the system's recommendations on subgroups with special needs, such as children and disabled, and on large homogeneous subgroups. To this purpose, the weight (w_{UM_i}) of a subgroup is evaluated as follows:

$$w_{UM_i} = \text{relevance}_{UM_i} * \text{cardinality}_{UM_i} / \text{total_cardinality}$$

where relevance_{UM_i} is the subgroup relevance (as stored in the subgroup model), $\text{cardinality}_{UM_i}$ is the subgroup cardinality and total_cardinality is the total number of tourists forming the group. This definition enables the system to privilege special subgroups depending on the portion of the whole group they represent.

The biased recommendation produces as a result the unique listing recommendation shown in Figure 3, where the attractions are sorted according to their overall ranking. In this case, the set of stars behind each attraction represents its suitability for the whole group, with special attention to the preferences of the most influent subgroups. It is worth noting that Figures 9 and 3 show recommendations for the same tourist group.

The biased recommendation of attractions deserves further comments concerning the computation of the suitability of an attraction. Although we selected formula (iii) to merge the viewpoints of the subgroups, other functions could be used. For example, by applying (i), we would compute the final recommendation list as an "intersection" of the subgroup-related recommendations: a tourist attraction would receive a good score only if it satisfies the preferences of all the subgroups, possibly ignoring the less relevant subgroups; moreover, the attraction would receive a very low score if its suitability is low for at least one relevant group. Our prototype uses function (iii) to provide less restrictive group recommendations, as the restrictive viewpoint on each individual attraction is offered by the separate listings, which the user can explore at any time to analyze the suitability of each item for the subgroups.

Also the relevance of a subgroup could be defined in other ways, e.g., to take into account the subgroup relevance, or the cardinality alone. Alternatively, the user might be asked whether she wishes that the system privileges some subgroups (regardless their characteristics) and her answer could be used to set the weights for the weighted sum

“on the fly”. For the moment, we have not experimented such alternative strategies, as an assessment of their plausibility, from the user’s perspective, is needed; for example, concerning the last strategy, we don’t know whether a typical user is able to associate reasonable relevance values to the subgroups.

Explanation of the System’s Suggestions

Although the star-based evaluation of attractions supports a qualitative assessment of the suitability of items, the user should be provided with explicit reasons for the recommendations. The explanation capability is essential when the recommendations are targeted to heterogeneous groups, because conflicting interests can lead to obscure suggestions. In the following, we sketch the explanation facilities offered by our system. When explaining why the system recommends certain items, we aim at synthetically qualifying them, so that the user can recognize the high-level properties making the items suitable to the group, without entering the details of all the features. The qualification of the preferred items has to separately address the subgroup viewpoints, in order to highlight the fact that the same item may fit the group members for different reasons. INTRIGUE supplements the recommendations by specifying the reasons for suggesting an item, given the preferences of the subgroups forming the tourist group. This feature helps the user to identify the aspects of tourist attractions best suiting the participants. As shown in Figure 3, in the unique listing modality each item is coupled with a sentence specifying, for each subgroup, the key properties determining the suggestion of the item. These properties are selected on the basis of two main factors: how important they are to the subgroups, and how well they match their preferences.

- The importance of the properties helps the system focus the explanations on the most meaningful types of information, for each subgroup, while leaving apart other aspects that, according to the system’s estimates, do not influence the subgroup evaluation of the attractions in a strong way.
- Within the set of important properties, the best matching ones are selected for the explanation, in order to highlight the positive aspects determining the recommendation.

For instance, the Lingotto palace (first item in the list) is well suited for the group and the system explains the recommendation by specifying that it is good for the children because it is very eye-catching, it requires low background knowledge and seriousness to be appreciated and the visit is quite short. Moreover, the direct user should appreciate the item because it is very eye-catching and has high historical value.

The provision of positive information about attractions enables the user to compare recommendations on the basis of good reasons for visiting them. This type of information also helps her quickly judge whether the reasons provided

by the system are good enough to accept the suggestion, as an attempt to address the issue of trust in the recommendations. Anyway, other recommendation strategies can be considered; for instance, we are currently extending the system to provide the user with both positive and negative information about attractions, in order to justify low-ranked attractions and to inform the user about potential problems for the various subgroups. In this extension, the trade-off between providing richer explanations and keeping them as short as possible has to be faced.

```

<!ELEMENT scheduler (agenda, itinerary?)>
  <!ELEMENT agenda (userChoices, generalConstraints?, organize?, save?)>
  <!ATTLIST agenda lang (italiano | english) "english">
    <!ELEMENT userChoices (item)*>
      <!ELEMENT item (itemName, delAg, constraints?)>
        <!ELEMENT itemName ( PCDATA)>
        <!ELEMENT delAg (imgDelAg, linkDelAg)>
          <!ELEMENT imgDelAg ( PCDATA)>
          <!ELEMENT linkDelAg ( PCDATA)>
          <!ATTLIST linkDelAg isActive (true | false) "true">
        <!ELEMENT constraints ( PCDATA)>
      <!ELEMENT generalConstraints (day, startTime, endTime, startPlace, endPlace)>
        <!ELEMENT day ( PCDATA)>
        <!ELEMENT startTime ( PCDATA)>
        <!ELEMENT endTime ( PCDATA)>
        <!ELEMENT startPlace ( PCDATA)>
        <!ELEMENT endPlace ( PCDATA)>
      <!ELEMENT organize (imgOrganize, linkOrganize)>
        <!ELEMENT imgOrganize ( PCDATA)>
        <!ELEMENT linkOrganize ( PCDATA)>
        <!ATTLIST linkOrganize isActive (true | false) "true">
      <!ELEMENT save (imgSave, linkSave)>
      [...]
    <!ELEMENT itinerary (user, startDate, activityList, problems?, refresh?)>
    <!ATTLIST itinerary lang (italiano | english) "english">
      <!ELEMENT user ( PCDATA)>
      <!ELEMENT startDate ( PCDATA)>
      <!ELEMENT activityList (visit)*, (transfer)*>
        <!ELEMENT visit (day, startTime, endTime, place)>
          <!ELEMENT place ( PCDATA)>
        <!ELEMENT transfer (day, startPlace, endPlace, time, by)>
          <!ELEMENT time ( PCDATA)>
          <!ELEMENT by ( PCDATA)>
      <!ELEMENT problems ((problem)*)>
      [...]
    <!ELEMENT refresh (imgRefresh, linkRefresh)>
    [...]

```

Figure 10: Portion of the DTD defining “agenda” pages.

USER INTERFACE

The User Interface can be viewed as a sequence of dynamically generated pages, belonging to specific page types.

Figure 1 shows a “browsing” page, while Figure 2 depicts a “show details” page; the pages in Figures 3 and 9 are

“recommendation” pages; finally, Figure 4 shows an “agenda” page. A DTD (Document Type Definition) defines the logical structure of each page type, i.e. the pieces of information it contains, their relations, and the elements enabling the user actions; e.g., buttons and links (W3C 2002). For instance, the DTD defining “agenda” pages in Figure 10 states that a scheduler consist of an agenda and an optional itinerary. The agenda contains a list of user choices and three optional elements: general constraints (day of visit, start/end time, and start/end place); the organize button; the save button. For each item in the list of user choices, the DTD specifies the name, a button to delete the item from the agenda, and local constraints (visiting the attraction in the morning or in the afternoon).

The generation of a page is performed in four steps; the first three of them do not depend on the device the interface is implemented on; the last one generates the appropriate, device-dependent user interface.

3. *Selection of the information to be presented.* This selection is specific for the different types of page. In a recommendation page, like the one shown in Figure 3, the selected information consists of a list of attractions, sorted on the basis of the group preferences. In a page describing an individual attraction, this information includes one or more sets of features, corresponding to the defined feature types. Finally, as far as the agenda is concerned, the information includes the list of attractions selected by the user, the specified constraints and the itinerary produced by the system (with error codes, if any failures occur).
4. *Generation of linguistic descriptions.* The internal representation of the selected information is translated to natural language by exploiting efficient template-based NLG techniques described in (Ardissono and Goy 2000). In the recommendation pages, the NLG module generates the explanations of the system’s suggestions by specifying the properties of items suitable for the group. In the pages presenting attractions, the module generates the descriptions in the style associated to the various sets of features (sentence, item list, etc.).
5. *Generation of the structure of the Web page.* Given the definitions in the corresponding DTD, a XML object representing the personalized content of the page, independent of the device, is produced.
6. *Surface generation.* The user interface is generated by transforming the XML object to a standard HTML, or WML page, depending on the user’s device. During this step, the surface generator decides which subset of the content stored in the XML object should be presented, given the space constraints of the user’s display. The feature classification is the basis for this type of decision. For instance, the page describing a tourist attraction shown in Figure 2 is generated for a Web browser, where there are no serious space constraints. When generating the same type of page for a mobile, the surface generator omits the picture and selects the most

important types of features, leaving the other items available via “more info” links. For instance, both the “description” paragraph, and the “characteristics” one are hidden, although they can be accessed via a link. Moreover, the surface generator provides the layout information, deciding where and how, on the display, the pieces of content should be placed. We exploited XSLT (eXtensible Stylesheet Language Transformations) for the transformation from XML to HTML/WML (W3C 2002). In particular, we defined three stylesheets, representing different graphical layouts (colors, fonts, etc.), for each type of page. These stylesheets are applied to the corresponding XML object to generate the HTML/WML code that is sent to the Web/WAP browser.

RELATED WORK

We compare our approach to the relevant work by considering three topics: the adaptation and explanation facilities, the representations of the information about tourist attractions and the group recommendation techniques.

As far as the first topic is concerned, other adaptive hypermedia systems dynamically generate user-adapted presentations by applying personalization strategies; see (Joerding 1998, Dale et al. 1998, Fink et al. 1999), and (Kobsa et al. 2001, Brusilovsky 1996, Brusilovsky 2001). Moreover, intelligent virtual guides personalize the visit of a museum, taking into account the user’s interests and her (physical) navigation style (Petrelli et al. 1999). There is also a strong interest in context-aware applications, supporting a selective presentation of information, based on the user’s physical location (Bederson 1995, Long et al. 1996, Marti et al. 1999, Abowd and Mynatt 2000, Fischer and Ye 2001). Other interesting examples of services that the user can access both from desktop and WAP phone are provided in (Riecken 2000, De Carolis et al. 2001, Billsus et al. 2000, Cotter and Smyth 2000, Paris et al. 2001).

INTRIGUE differs from these systems in several aspects:

- The first one concerns the design and technical choices for the generation of the user interface. On the one hand, our system supports access from different devices by exploiting standard, XML-based techniques, while most of the other systems use proprietary translation tools for this purpose. On the other hand, INTRIGUE is less flexible than the approaches described in (De Carolis et al. 2001) and (Paris et al. 2001), which take into account the characteristics of the output device both in the content planning phase and in the surface generation one. Our system generates the content to be presented in a unique version and takes the user’s device into account only in the surface generation phase, to split the content in different pages, and possibly omit information. Although this type of filtering is coarse grained, an advantage of our approach is the extensibility

of the system to support different types of user interfaces: this extension only requires the definition of new stylesheets for the transformation of the XML content into the required format.

- INTRIGUE manages heterogeneous group models, in contrast to the homogeneous ones handled in other systems; e.g., see the household models adopted in several TV recommenders. Our approach supports individual and group recommendation, which is particularly relevant to the tourism domain. In fact, people having different preferences and requirements may join the same tour, which should be organized by taking the interests of all the participants into account. Moreover, different from other systems, such as PolyLens (O'Connor et al. 2001), our system handles group heterogeneity without imposing the identification of the individual group members. This approach supports a quicker specification of the heterogeneous group without losing information useful for the recommendation task.
- Our system explains its own recommendations by relating them to the possibly different requirements of heterogeneous groups. The system is informed about both the preferences of the subgroups forming a tourist group and the properties suiting the alternative options. Therefore, it can explicitly address the key properties of items, making the user aware of the rationale behind the recommendations and possibly helping her in the comparison of alternative recommended items. This aspect distinguishes INTRIGUE from other recommenders, which do not explain their own recommendations, or provide justifications based on totally different information. For instance, as described in (Herlocker et al. 2000), collaborative filtering systems typically explain their own recommendations by reporting the ratings of items provided by people with selection histories similar to the user, or they rely on the rate of good recommendations produced in previous interactions with the same user. As another example, NLSA (Chai et al. 2001) only motivates its recommendations by highlighting the features of items satisfying the user's explicit requirements. More interesting, News Dude (Billsus and Pazzani 1999) explains its recommendations in a content-based fashion, by analyzing analogies and differences between the items suggested in previous recommendations and already inspected by the user.

In addition to these aspects, INTRIGUE offers an interactive and persistent agenda, accessible both from desktop and handset devices, that supports the user in scheduling a tour and changing itineraries at any time. Also the GUIDE system supports users in the real-time organization of tours (Cheverest et al. 2000). However, it requires specialized handset devices and hardware located in the town to be visited. Instead, INTRIGUE is based on standard

wireless connections and on the use of mobile phones. Moreover, the interactive agenda maintains persistent itineraries that can be ubiquitously consulted and updated.

As far as the knowledge representation is concerned, other approaches have been applied to the representation of the item features. For example, (Milosavljevic 1999) describes a powerful representation supporting the comparison of items, but requiring a relevant knowledge engineering effort in the specification of the domain knowledge. Alternatively, flat knowledge representations adopted in information filtering systems (Resnick and Varian 1997) do not support the generation of structured presentations. The representation we adopted is less detailed than the one proposed by Milosavljevic and it supports flexible presentation strategies. For instance, in the SeTA adaptive Web store, this representation supported fine-grained decisions about the features of items to be presented (Ardissono and Goy 2000b). Moreover, it was used to apply unobtrusive user modeling techniques for the recognition of the user's interests, on the basis of her navigation behavior in the Web store catalog.

To conclude, we discuss the techniques for ranking tourist attractions on the basis of the preferences of a tourist group. For the estimation of the user's evaluation of tourist attractions, we apply a fuzzy AND function instead of other techniques, such as Multi-Attribute Utility Theory (MAUT, see some applications of this theory in (Jameson et al. 1995, Linden et al. 1997, Carberry et al. 1999)). The reason is that the overall evaluation of each item has to take into account the presence of necessary properties, which a tourist attraction must offer in order to be recommended to the user, and MAUT does not support this selective power in a satisfactory way. More specifically, MAUT is based on an additive evaluation of items and calculates their overall utility for the user as the weighted sum of the individual utilities carried by the attributes of the items (the weights represent the user's interests in the various properties). Therefore, if an item matches all the user's preferences but one, the item very likely receives a good score and is therefore recommended to the user. Unfortunately, this behavior does not fit the requirements of the tourist domain, where the presence of totally incompatible properties must downgrade the evaluation of items, in order to avoid unsuitable recommendations (the properties may represent real needs, not only preferences; e.g., consider the presence of special transportation systems). In our case, a property receiving a null score must have a drastic impact on the overall satisfaction score of the item, bringing it to the minimum level. Although, in MAUT, negative values can be provided, as done in (Carberry et al. 1999), the selection of values supporting this discrimination power is rather difficult. In contrast, the use of fuzzy AND formulae makes the task easier, because they produce scores ranging in a constrained interval ($[0,1]$) and guarantee that, if any contribution to the evaluation process is null, the evaluation of the item is null.

CONCLUSIONS AND FUTURE WORK

This paper has presented the Interactive Tourist Information Guide (INTRIGUE), a prototype tourist information server developed at the Dipartimento di Informatica of the University of Torino for the provision of personalized information about the area around Torino city. The system tailors the recommendations of attractions to heterogeneous tourist groups, such as families with children and groups including disabled people. Moreover, it explains the recommendations with respect to the preferences of the tourist group members and offers an interactive agenda that helps the user to schedule the tour. As the system can be accessed from desktop and handset devices, it provides ubiquitous assistance during the organization of tours: the interactive agenda enables the user to store persistent itineraries and revise them at any time, also during the tour.

We have tested our system with users who helped us to improve its interface and behavior. The users appreciated the system's explanation facilities, both regarding group recommendations and failures in itinerary generation, because they clarify the system's behavior, therefore increasing the trust in its recommendations. Given this positive feedback, in our future work we will improve the generation capabilities to make the explanations more fluent and to provide negative information about items for the explanation of low ranking of items.

There was some negative feedback in the organization of the device-dependent interaction with the user. For instance, the registration form for the specification of heterogeneous tourist groups was considered too time consuming for a minibrowser interaction. To address this issue, we have extended the system to support logging without registering: if the user logs in this way, she receives non personalized recommendations, but she can access other facilities, such as category and geographical search, itinerary scheduling and presentation of tourist attractions. We will also investigate the application of unobtrusive user modeling techniques aimed at the identification of the user's preferences on the basis of the observation of her browsing behavior.

The users appreciated the possibility of saving itineraries and modifying them in subsequent interactions, as this facility would enable them to use the interactive agenda not only at tour scheduling time, but also during the tour.^{vi}

Notice however that, at the current stage of technological development, wireless connections are unreliable and make long-lasting interactions, typically performed in the wired case, problematic. Moreover, in most countries, wireless communications are very expensive. To face these issues, we extended INTRIGUE to support local access to the itinerary scheduling activities: as described in (Ardissono et al. 2002), the interactive agenda can be downloaded on the user's mobile phone and run without exploiting wireless connections.

Finally, our future work will include further experiments, in order to collect usage information and to test the system with real mobile phones: as already specified, the mobile facilities of our system have been tested by exploiting a WAP simulator. Our future testing will also include collecting results on the quality of the system's recommendations, in the case of heterogeneous tourist groups. The basic mechanisms used for ranking individual items for a group of homogeneous users have already received a validation as they are the same as the ones adopted and tested in SeTA; see (Ardissono and Goy 2000b) for a description of the results of those experiments.

ACKNOWLEDGMENTS

This work has been partially funded by MIUR and by the National Research Council CNR (grant number CNRG0015C3). The authors are grateful to Graziano Anselmetti for his contribution to the development of the mobile facilities offered by our system. The authors also thank the anonymous reviewers who helped to improve this paper with very detailed and valuable comments.

REFERENCES

- Abowd, G. D., and E. D. Mynatt. 2000. Charting Past, Present and Future Research in Ubiquitous Computing. *ACM Transactions on Computer-Human Interaction* 7(1):29-58.
- Ardissono, L., and A. Goy (2000). Dynamic Generation of Adaptive Web Catalogs, in *Lecture Notes in Computer Science n. 1892: Adaptive Hypermedia and Adaptive Web-Based Systems*, P. Brusilovsky, O. Stock, and C. Strapparava, eds. Springer Verlag, Berlin, pp. 5-16.
- Ardissono, L., and A. Goy 2000b. Tailoring the Interaction with Users in Web Stores. *User Modeling and User-Adapted Interaction* 10(4):251-303.
- Ardissono, L., A. Goy, G. Petrone, and M. Segnan 2001. A Software Architecture for Dynamically Generated Adaptive Web Stores. *Proc. 17th International Joint Conference on Artificial Intelligence, IJCAI'01*, Seattle, WA, pp. 1109-1114.
- Ardissono, L., A. Goy, G. Petrone, M. Segnan and P. Torasso (2002). Ubiquitous User Assistance in a Tourist Information Server, in *Lecture Notes in Computer Science 2347: Adaptive Hypermedia and Adaptive Web Based Systems*, P. Brusilovsky, P. De Bra, and R. Conejo, eds. Springer Verlag, Berlin, pp. 14-23.
- Bederson, B. B. 1995. Audio Augmented Reality: A Prototype Automated Tour Guide. *Proc. Conf. companion for the ACM Conference. on Human Factors in Computing Systems, CHI'95*, Denver, CO, pp. 210-211.
- Bergman, E. 1999. *Information Appliances and Beyond. Interaction Design for Consumer Products*. Morgan Kaufmann Publishers.

- Billsus, D., and M. Pazzani 1999. A Personal News Agent that Talks, Learns and Explains. *Proc. 3rd International Conference on Autonomous Agents, AGENTS'99*, Seattle, WA, pp. 268-275.
- Billsus, D., M. Pazzani, and J. Chen 2000. A Learning Agent for Wireless News Access. *Proc. 2000 International Conference on Intelligent User Interfaces, IUI'00*, New Orleans, LA, pp. 33-36.
- Brusilovsky, P. 1996. Methods and Techniques of Adaptive Hypermedia. *User Modeling and User-Adapted Interaction* 6(2-3):87-129.
- Brusilovsky, P., 2001. Adaptive Hypermedia. *User Modeling and User-Adapted Interaction* 11(1-2):87-110.
- Brusilovsky, P., and M. Maybury 2002. *Communications of the ACM: Special Issue "The Adaptive Web"* 45.
- Carberry, S., J. Chu-Carroll, and S. Elzer 1999. Constructing and Utilizing a Model of User Preferences in Collaborative Consultation Dialogues. *Computational Intelligence* 15(3):185--217.
- De Carolis, B., F. de Rosis, and S. Pizzutilo 2001. Context-Sensitive Information Presentation. *Proc. UM'2001 Workshop on User Modelling in Context-Aware Applications*, Sonthofen, Germany.
- Chai, J., V. Horvath, N. Nicolov, M. Stys-Budzikowska, N. Kambhatla, and W. Zadrozny 2001. Natural Language Sales Assistant - a Web-based Dialog System for Online Sales. *Proc. 13th Innovative Applications of Artificial Intelligence Conference, IAAI'01*, Seattle, WA, pp. 19-26.
- Cheverest, K., N. Davies, K. Mitchell, A. Friday, and C. Efstathiou 2000. Developing a Context-Aware Electronic Tourist Guide: Some Issues and Experiences. *Proc. Human factors in computing systems, CHI'2000*, The Hague, The Netherlands, pp. 17-24.
- Cotter, P., and B. Smyth 2000. WAPing the Web: Content Personalization for WAP-enabled Devices, in *Lecture Notes in Computer Science n. 1892: Adaptive Hypermedia and Adaptive Web-Based Systems*, P. Brusilovsky, O. Stock, and C. Strapparava, eds. Springer Verlag, Berlin, pp. 98-108.
- Dale, R., S. J. Green, M. Milosavljevic, and C. Paris 1998. Dynamic Document Delivery: Generating Natural Language Texts on Demand. *Proc. 9th International Conference and Workshop on Database and Expert Systems Applications, DEXA '98*, Vienna, Austria.
- Fink, J., A. Kobsa and A. Nill 1999. Adaptable and Adaptive Information for All Users, Including Disabled and Elderly People. *New review of Hypermedia and Multimedia* 4:163-188.
- Fischer, G., and Y. Ye 2001. Exploiting Context to Make Delivered Information Relevant to Tasks and Users. *Proc. UM'2001 Workshop on User Modelling in Context-Aware Applications*, Sonthofen, Germany.
- Herlocker, J. L., J. A. Konstan, and J. Riedl 2000. Explaining Collaborative Filtering Recommendations. *Proc. ACM 2000 Conference on Computer Supported Cooperative Work, ECSCW'00*, Philadelphia, PA.

- Jameson, A., R. Schäfer, J. Simons, and T. Weis 1995. Adaptive Provision of Evaluation-Oriented Information: Tasks and Techniques. *Proc. 14th International Joint Conference on Artificial Intelligence, IJCAI'95*, Montreal, Canada, pp. 1886-1893.
- Joerding, T. 1998. Intelligent Multimedia Presentations in the Web: Fun Without Annoyance. *Proc. 7th World Wide Web Conference, WWW7*, Brisbane, Australia.
- Kobsa, A., J. Koenemann, and W. Pohl 2001. Personalized Hypermedia Presentation Techniques for Improving Online Customer Relationships. *The Knowledge Engineering Review* 16(2):111-155.
- Linden, G., S. Hanks, and N. Lesh 1997. Interactive Assessment of User Preference Models: The Automated Travel Assistant. *Proc. 6th Conference on User Modeling, UM'97*, Chia Laguna, Italy, pp. 67-78.
- Long, S., R. Kooper, G. D. Abowd, and C. G. Atkeson 1996. Rapid Prototyping of Mobile Context-Aware Applications: The Cyberguide Case Study. *Proc. 2nd Annual International Conference on Mobile Computing and Network, MobiCom'96*, New York, NY, 1996.
- Marti, P., A. Rizzo, L. Petroni, G. Tozzi, and M. Diligenti 1999. Adapting the Museum: a Non-intrusive User Modeling Approach. *Proc. 7th International Conference on User Modeling, UM'99*, Banff, Canada, pp. 311--313.
- McCoy, K. F. 1989. Generating Context-Sensitive Responses to Object-Related Misconceptions. *Artificial Intelligence* 41:157-195.
- Greening, D. R. (2000). LikeMinds: Building consumer trust with accurate product recommendations. http://likeminds.com/technology/white_papers/websell
- Milosavljevic, M. 1999. *The Automatic Generation of Comparison in Descriptions of Entities*. Ph.D thesis, Macquarie University, Sydney, Australia.
- O'Connor, M., D. Cosley, J. A. Konstan, and J. Riedl 2001. PolyLens: a Recommender System for Groups of Users. *Proc. European Conference on Computer Supported Cooperative Work, ECSCW'01*, Bonn, Germany.
- Paris, C., S. Wan, R. Wilkinson, and M. Wu 2001. Generating Personal Travel Guides - and Who Wants Them? *Proc. 8th International Conference on User Modeling, UM'01*, Sonthofen, Germany, pp. 251-253.
- Petrelli, D., A. De Angeli, and G. Convertino 1999. A User Centered Approach to User Modelling. *Proc. 7th International Conference on User Modeling, UM'99*, Banff, Canada, pp. 255-264.
- Plua, C., and A. Jameson 2002. Collaborative Preference Elicitation in a Group Travel Recommender System. *Proc. AH2002 Workshop on Recommendation and Personalization in E-Commerce*, Malaga, Spain.
- Resnick, P., and H. R. Varian 1997. *Communications of the ACM: Special Issue on Recommender Systems* 40.
- Rich, E. (1989). Stereotypes and User Modeling, in *User Models in Dialog Systems*, A. Kobsa, and W. Wahlster, eds. Springer Verlag, Berlin, pp. 35-51.
- Riecken, D. 2000. *Communications of the ACM: Special Issue on Personalization* 43.

W3C 2002. World Wide Web Consortium, www.w3.org.

ⁱ The figure also shows the view that the user has when she uses a mobile phone. The topic is discussed in the following section.

ⁱⁱ See section “Modeling data about tourist attractions”.

ⁱⁱⁱ If the group is homogeneous, it consists of a single subgroup. In the most complex case, the group is completely heterogeneous, so each person in the group is represented by a different subgroup.

^{iv} We assume that the user is the tour organizer and is responsible for the description of the participants. See (Plua and Jameson 2002) for an interaction model where group members concur in the specification of preferences.

^v As noticed by a reviewer, we assume that the user has a single preferred value for each property and we model the system’s uncertainty about the value as a probability distribution on the values of the preference. This is a simplification, as the user may prefer more than one value. The assumption might be removed by defining value functions over the possible values of each property; e.g., see (Jameson et al 1995). However, the definition of these functions is complex, as we consider heterogeneous types of properties, requiring the definition of different value functions. For example, not all the scalar properties have the same bias. In some properties, such as the artistic value, or the presence of special transportation systems, the “null” value is the worst one and the “high” one is the best one; other properties, such as the cost and the typical length of the visit, require value functions growing the other way around. Finally, we have to model non-scalar properties, such as the type of documentation about a tourist attraction (“Documentation” in Figure 6) requiring ad hoc value functions. The Seta2000 toolkit, used in the development of INTRIGUE, does not support the definition of these value functions. Its extension might be considered, but the introduction of these functions increases the overhead at Web-store configuration time, due to the specificity of the information to be introduced. Therefore, there is a clear trade-off between the support of finer-grained inferences and the usability of the workbench.

^{vi} Our system was initially designed as a desktop tourist information server; therefore, contextual information such as the user’s physical location was not central for our application. Taking this type of information into account is part of our future work.