Personalised Cloud-Based Recommendation Services for Creative Tourism

Christos Emmanouilidis
ATHENA Research and Innovation Centre, Greece, chrisem@ceti.athena-innovation.gr
Remous-Aris Koutsiamanis
ATHENA Research and Innovation Centre & Democritus University of Thrace, Greece, ariskou@ee.duth.gr
Aimilia Tasidou
ATHENA Research and Innovation Centre & Democritus University of Thrace, Greece, atasidou@gmail.com
Pavlos S Efraimidis
Democritus University of Thrace, Greece, pefraimi@ee.duth.gr
Ioannis Refanidis
University of Macedonia, Greece, yrefanid@uom.gr

Abstract

Mass tourism increasingly gives way to creative and niche tourism. The creative visitor is a more informed tourist, better integrating with the place of a visit, its local culture, history and people. The best ambassadors of a region are local stakeholders who are able to include their offerings within the wider cultural context of the region, as well asvisitors who enjoyed a personalised experience, tailored to their wishes, needs and interests. Recent advances in computer-assisted tourism services are well-suited to support the personalisation of the touristic product. Specifically, recommender systems make it possible to support a personalised experience by tailoring offered information and services to the specific context of an individual visit. MYVISITPLANNER is a cloud-based service employing a recommender engine to offer personalised suggestions of cultural activities and a planning tool to create valid tour plans for visitors and residents of an area. In doing so, it also addresses the need of local authorities and regional stakeholders to raise the visibility of their offerings, by providing an easy to use platform for their recording and inclusion in individual tour planning. This paper offers an overview of the underlying logic and implementation for a personalised activities recommender system, developed as part of the intelligent itineraries planning system MYVISITPLANNER. The recommender engine takes into account both a taxonomy of possible activities, as well as user preferences and past recommendations, making it capable of offering recommendations tailored to the overall visit profile.

Keywords: Recommender systems; personalised itineraries; cloud services; privacy; creative tourism.

Introduction

Mass tourism has been the dominant force in tourism business for many decades. Increasingly, there is in parallel a strong surge of individual travellers or traveller groups, wishing to create their own personalised trips and experience offerings and activities that best much their interests and wishes, as well as their constraints. The term 'creative tourism' is often employed to denote a stronger participatory role for the traveller in shaping up the visit experience, boosted by enabling information and communication technologies (ICT). While creativity is mostly associated with urban tourism (Richards, 2014) and smart cities concepts(Lamsfus, et al., 2014), in principle it is not restricted within a single city context but may be applicable to wider area visits.

The current landscape in the tourism industry has been hugely influenced and largely shaped up by the advent of a range of emerging technological enablers, primarily related to advances in ICT, namely web-based and semantic computing, ubiquitous networking, media technologies, smart interfaces and highly capable portable devices and smartphones. A typical tourist may now be considered as a highly active actor having access to a multitude of data sources, dynamically interacting with the environment and offered electronic services and making fast decisions on the basis of the most relevant information, according to individual (or group-wise) interests, desires and constraints. A key enabler to delivering this level of personalisation in the touristic product experience is the emergence of ICT applications able to tailor their offered information and services according to the apparent context of a service request (Emmanouilidis, et al., 2013). The ability to offer sophisticated recommendation services is essentially a reflection of the technical ability to close the information loop between services providers, related stakeholders and end users in the tourist industry and its wider value chain.

This paper presents the development and the offered functionality of a new cloud-based service, named MYVISITPLANNER (www.myvisitplanner.com) (Refanidis, et al., 2014). Instead of relying on static information that can locally be stored, the cloud-based service offers the technological means to integrate large amounts of data but make available to the end user only the contextually relevant ones. The cloud-based nature of the service brings additional flexibility in terms of the area range coverage, which is practically unlimited, as well as to the range of stakeholders and end users that it may serve. The offered functionality enables a prospective tourist to create and receive personalised recommendations when planning to visit a geographic area of interest up to the point of creating a valid itinerary plan for the visit, while respecting other commitment constraints, such as pre-set appointments or other activities. One of the core system functionalities is to tailor the recommended visit activities to personal preferences, in a privacy-respectful manner, by exploiting an underlying dedicated ontology of activities, default, edited or inferred personalised profiles, as well as past visit historical data. The hybrid manner of operation of the involved recommender engine enables reaching balanced recommendation suggestions, even in the presence of minimal information about the user or past users.

The paper focuses on describing the recommender engine functionality and is structured as follows. Section 2 discusses the role of recommender systems in tourism, especially focusing on itinerary planning. Section 3 provides a high level overview of the overall system, which is described in more detail in (Refanidis, et al., 2014). Section 4 presents the design and implementation of the recommender system. Section 5 illustrates how the system can be used to derive activities recommendations. Finally, section 6 concludes the paper and sets future directions.

Recommender Systems in Tourism

The lifecycle of a visit experience essentially contains three phases, namely inception& planning, realisation and post-visit experience. Whereas in the past the time lag for any meaningful information flows between these phases was slow, current advances in ICT have made it possible to streamline data exchanges and therefore make a real impact on how the actual touristic product can be massively customised to serve individual needs at virtually real time pace (Figure 1). Having a closer look at each individual tour activity phase brings about a wealth of opportunities to enhance offered services through smart information processing and better exploitation of visit lifecycle data.

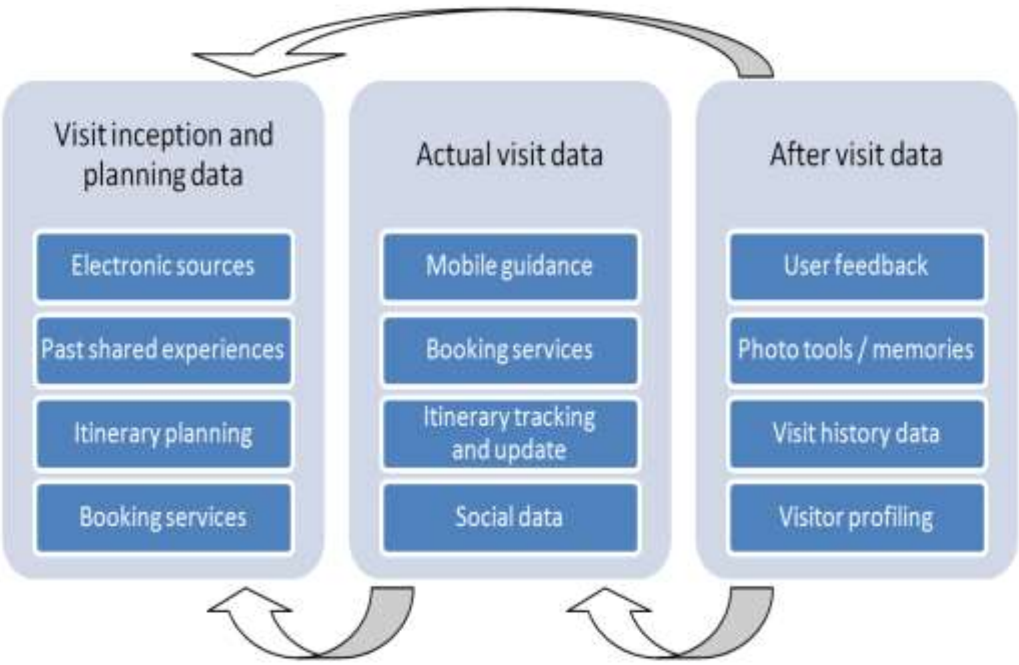

Figure 1. Information availability during visit lifecycle phases

Travellers may consult electronic information sources and shared spaces of past user experience to learn details about their prospective places to visit, determine what to include in their itinerary and use e-booking services to proceed with their planning. During the visit they may use their smartphones to get on-the-move information and services support, share data via social tools and produce a visit track, possibly updating their initial itinerary and make additional bookings if needed. After the visit they may leave available their own visit data, either in the form of explicit feedback and shared text and media, or indirectly through their activities tracking history, enabling visitor profiling. Smart mobile guides are increasingly employed to enhance the visit experience, vastly shortening the time between information acquisition and action, adding flexibility to the actual visit and enabling tourists to customise their actual path on the go, according to their wishes, interests and constraints (D' Amico, et al., 2013; Schaller, 2014).

A key enabling factor for the successful personalisation of services is the ability to capture context and tailor the offered information and services to the apparent context in each case (Emmanouilidis, et al., 2013). While context has been an attribute largely associated in the past with computational linguistics, it is particularly useful when providing mobile services, enabling the user to become a mobile actor interacting with the surrounding environment. While context is not actually uniquely defined, it can be modelled for specific application domains and in the case of a tourist, acting as a mobile user, it may be classified to belong into five broader categories, namely user, environment, social, service and system context (Emmanouilidis, et al., 2013).

The ability to exploit relevant information to offer recommendations for activities and points of interests to include in an itinerary, has given rise to a body of work on itinerary recommender systems(Agarwal, et al., 2013; Cha, 2014; Chen, et al., 2014; Chen, et al., 2013; D' Amico, et al., 2013; Gavalas, et al., 2014; Herzog and Wörndl, 2014; Koceski and Petrevska, 2013; Lucas, et al., 2013; Moreno, et al., 2013; Schaller and Elsweiler, 2014; Yang, 2013). These follow the typical patterns of design for generic recommender systems (Bobadilla, et al., 2013), but are customised for the tourism application domain. In general, recommender systems seek to combine large quantities of often heterogeneous data, compare it with a certain query or situation and offer information or service recommendations. Important factors for the success of such recommender systems are the way the contextualisation of such processes is modelled and implemented, as well as the way they exploit past information and experience, by employing methods such as Collaborative Filtering (CF) and Clustering. Yet, significant challenges still exist. The quantity of data may vary greatly, ranging from little to no past history data all the way to huge quantities of data, especially when a visit is not confined to a city but to a much larger area. Users may act in very different ways, not always following precise patterns. A single visitor might even wish to follow different visit patters, depending on time, budget or other constraints and preferences. Furthermore, the validity and quality of the offered data may not always be guaranteed and there is a need to promote stakeholders's engagement to volunteer such data, a service that they are only likely to offer if they can see a direct business or other benefit in the process.

The MYVISITPLANNER Cloud-Based Recommender and Itinerary Planning System

This section describes how an enhanced recommendation functionality can be achieved by bringing together technological enablers such as web-based and semantic computing, cloud implementation and services, machine learning and privacy preserving computing. It forms part of the MYVISITPLANNER cloud-based personalised itinerary planning system, originally presented in (Refanidis, et al., 2014). The main services offered by MYVISITPLANNER for visitors are (a) recommendations of activities to be performed during a visit, taking into account the area of interest and available activities in the area, according to defined or pre-selected visit profiles and preferences, past visit data evaluations and information content as defined in a domain ontology of visit activity types, (b) itinerary plans, provided by a scheduling engine, taking into account the recommended activities by the recommendation engine, individual visitor calendar constraints and scheduling preferences, as well as available time for the visit. For activity providers MYVISITPLANNER offers service for enlisting, describing in a structured way and classifying theoffered activities within the MYVISITPLANNER system, so as to be incorporated in the personalised recommendations and itineraries.

While the above constitutes a high-level description of the offered services, more fine-grained services are further included in the system to aid both prospective visitors and area activity providers. The overall system architecture and functionality is described in (Refanidis, et al., 2014), whereas the scheduling and planning mechanisms were adapted from (Refanidis and Alexiadis, 2011; Refanidis, et al., 2011) and further elaborated in (Refanidis, et al., 2014), (Alexiadis and Refanidis, 2013) and (Alexiadis and Refanidis, 2013).

A broad range of activities that might be of interest are considered, such as visiting museums, archaeological sites, monuments, exhibition spaces, attending events, walks through interesting urban and rural paths, outdoors activities, such as mountaineering, rafting or swimming, and many others. Personalization is offered both in terms of deriving recommended activities, as well as devising a personalised visit plan, taking into account the visitors’ interest; preferences with respect to the scheduling of activities; and, finally, constraints imposed by other tasks already scheduled within the visitors' calendar. The cloud-based architecture of the implemented solutions makes it possible not to needlessly limit the offered services to a single city or narrow area of interest, as is typically offered by other itinerary planning systems, to better exploit past experience and user evaluations, by means of a hybrid recommendation approach, that incorporates customised collaborative filtering, and to easily scale-up, expand and update the offered services, without the need for users, i.e. activity providers, or visitors, to update software on their own devices (Refanidis, et al., 2014).Sample tool screens are shown in Figure 2.

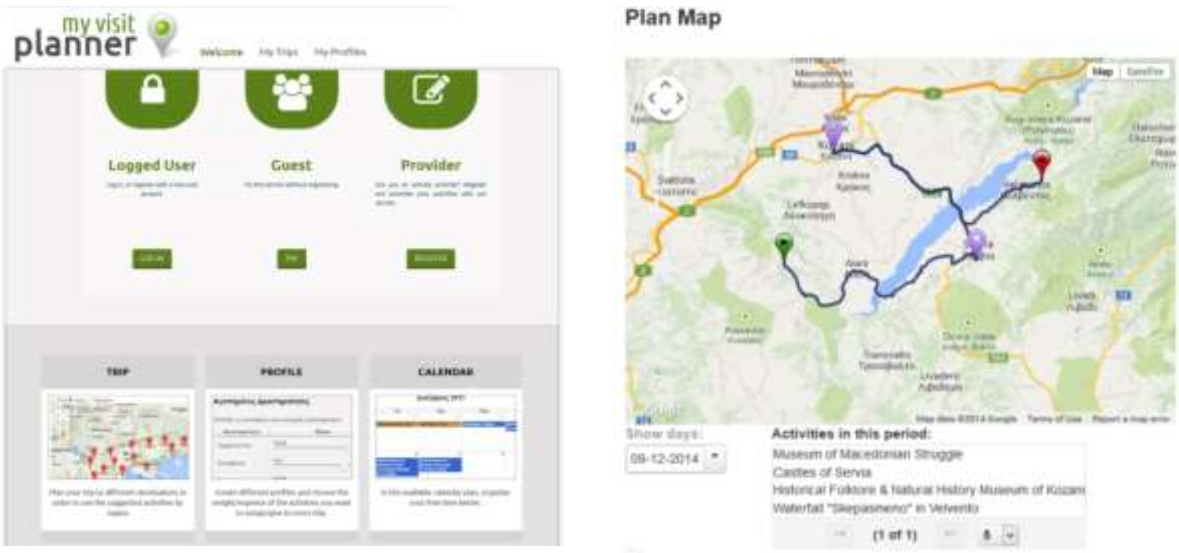

Figure 2. (left) MYVISITPLANNER main interface (right) example of itinerary map-view

MYVISITPLANNER Recommender

The MYVISITPLANNER recommendation approach has been designed so as to offer recommendations exploiting both the similarity of intended visit profiles with existing profiles and activity types, as well as historical evaluation ratings data. In order to achieve this a dedicated activities ontology has been defined, while similarity is assessed by means of appropriately defined distance functions and a fusion of recommendation mechanisms. These are described next.

The activities ontology

The MYVISITPLANNER system ontology design requirements were driven by its intended end-use, that of facilitating the recording of a potential tour visit activities and producing activities recommendations during an area visit. Therefore, it was necessary to design a dedicated ontology to describe activity types in a structured manner. It was clearly not aimed at describing or recording cultural heritage as such an objective would complicate rather than facilitate the end result. This is an important design requirement, since activities offerings were intended to be directly handled by activity providers to input activity descriptions. Such users are not expected to be generally familiar with formal ontological descriptions. Therefore, rather than defining a formal cultural activities ontology, the choice was to define a simple tour activities structure employing commonly perceived terms. A representative subset of the employed activity ontology is presented in Figure 3.

The main hierarchy contains available activity types, such as “Monument” or “Archaeological Site”. The objective is to record the corresponding activities, that is *visiting* a “Monument” or “Archaeological Site”, but not the recording of the actual "Monument" or the "Archaeological" site per se. The activity types are further analyzed at deeper hierarchy levels. An activity provider, thus, has the flexibility to stay at the more abstract hierarchical level. This may be sufficient in many cases. However, it is also possible to proceed at a more detailed hierarchical level when describing offered activities. This may be helpful for users wishing to offer a more elaborate visit profile in order to receive more focused recommendations. Additionally, auxiliary cross-cutting categorizations of the main activity type hierarchy help mitigating a potential combinatorial explosion of activity types that would have otherwise been introduced by a categorization of very fine granularity. More specifically, the theme hierarchy allows the expression of the thematic category of the activity; the historical era (epoch) hierarchy enables a categorization according to the historical period of interest; and the target group hierarchy assists in linking activities with different target groups.

A key usage target for the defined ontology is to support the description of offered activities, as well as prospective visitor preferences for activity types. Therefore, the description allows flexible sets of ontology entries to be specified. For example, in describing a castle on the shore of a lake, the set {Castle, Lake} can be specified. When considering visitor preferences, composite weighted sets of ontology entries are applicable. For example, a visitor may have a specific interest in sites of speleological interest, with a strong dislike for visiting castles, while remaining indifferent to seeing bridges. The case is just an example of the type of preferences that can be served and can be specified by the set {(Cave, 1.0), (Castle, 0.0), (Bridge, 0.5)}. By defining preferences over the activity types rather than the auxiliary hierarchies, the profiling is characterised by simplicity. Therefore, an evident advantage of the adopted ontological approach is that it combines simplicity towards the user with the ability to handle more complex associations, served by employing a composite similarity metric. In this way the produced recommendations can be finely or coarsely elaborated, better matching the level of detail that an end user may be interested in specifying. The overall process remains simple enough for the average user, while it may involve sufficient detail to satisfy visitors looking for specialised recommendations.

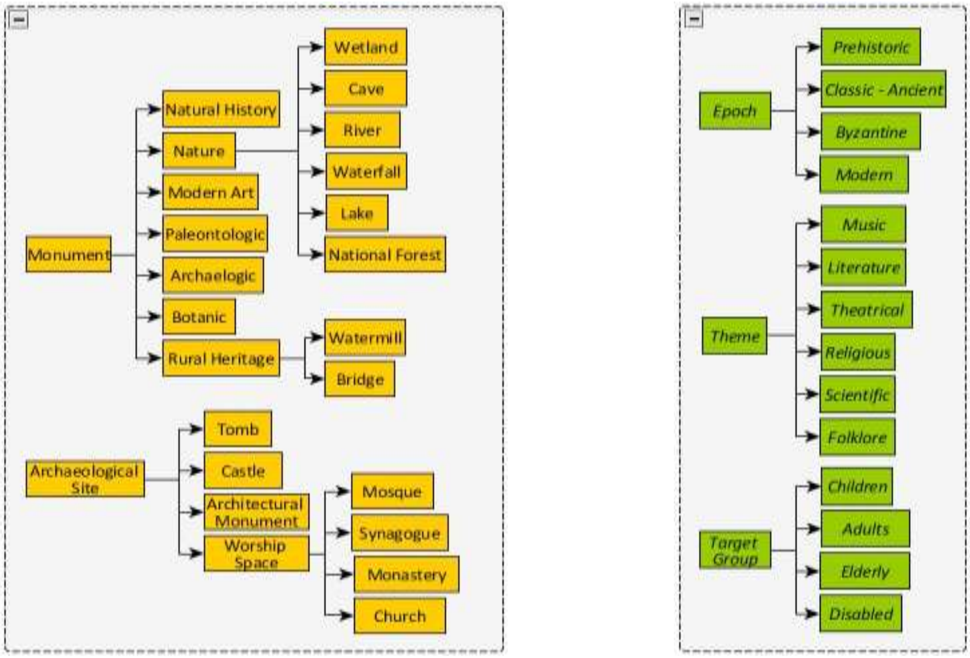

Figure 3. Partial view of the activities ontology (left) activity type hierarchy (right) auxiliary hierarchies

Recommender Architecture

A recommender could be built to offer recommendations based on (a) explicit user-specific defined preferences with respect to activity types, (b) experience from past user-specific interaction with the system to infer eventual user preferences with respect to activity types, and (c) pre-defined generic visit profiles and user assignment to such a visit profile via some form of clustering. While all the aforementioned design options have merits and are supported to a smaller or greater extent in the literature by available recommender systems, each one has specific shortcomings. The first choice is intuitive and practical but the prior matching of end user sense of preferences with activity may fail to capture the eventual preferences, as these might have been captured by past experience from user interaction with the systems. In the second case, historical data can be used to train a system to develop a 'learned' visit profile, thus offering a way forward to take into account visit posterior data, validated by actual choices. However, for such a training to be effective, a dense representation of training data covering different scenario combinations would be needed. This can rarely be the case in practice and it would further limit the ability to offer valid recommendations to cases where sufficient past experience exists. Even if a wealth of such data became available, managing a large set of data comprising combinations of offered activities and users would soon become computationally hard. The clustering of each new visit profile to a pre-existing set of profiles would be a third choice, but this might either become a limiting factor limiting the personalisation of the visit experience to static profiles, or, in cases of dynamically updated visit profiles, would bring about real time performance and complexity issues. In MYVISITPLANNER a balanced solution is implemented (a) by employing a hybrid recommendation engine, exploiting definitions based on a dedicated activity ontology, visit profile clustering and learning from past data to enhance the quality of the recommendation, and (b) by deploying the actual implementation on the cloud, so as to boost the scalability of the solution to work with data from large geographic areas, with a wealth offered activities and potential users. The functionality of the hybrid recommender engine is described next.

MYVISITPLANNER assists users in obtaining relevant activity recommendations by means of a hybrid collaborative filtering scheme system (Ricci, et al., 2010). The hybrid nature of the recommender enables it to combine two individual recommendation engines. Each engine's recommendation is then fed into a fusion function in order to derive a final ordered list of proposed activities. Some aspects of the engine functionality are shown in Figure 4.

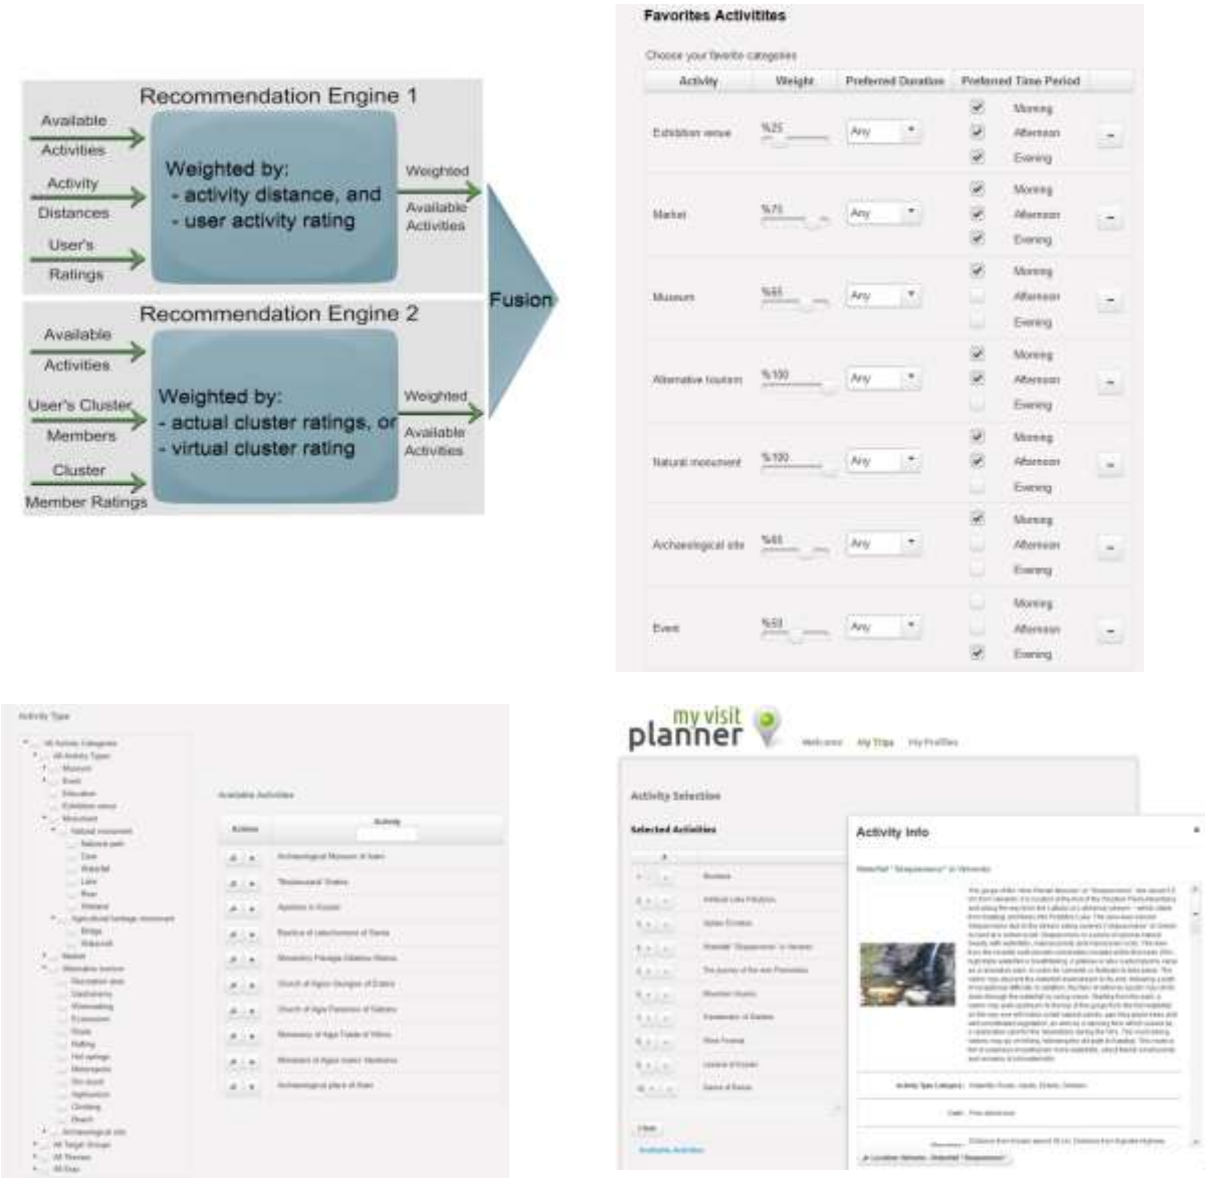

Figure 4. Recommender engine: (top left) the hybrid implementation (top right) visit profiling (bottom left) activity types editing (bottom right) proposed activities

The detailed analysis of the recommendation engine mechanism is beyond the scope of the present paper but a brief summary follows. In general, the recommendation engine mechanism performs recommendations on the basis of functions and relationships defined over the sets of activity types, activities, users, activities evaluations by users and activities preferences by users. Relationships may be more complex than unary, for example users may have more than one preference profiles, activities may belong to more than one activity types and activity types may be defined at multiple levels of abstraction. To facilitate the recommender engine's reasoning mechanism, sets activity types and sets of weighted activity types can be defined, so that a single activity may belong to different activity types (non-weighted) or a user may express preference for different activity types according to a weighting factor. Furthermore, appropriate distance metrics are defined between activity types, as well as between sets of non-weighted/weighted activity types, with the former always considered to be derived as a specific instance of the latter. An important factor that is taken into account in the definition of the distance metric is the information content of activity type nodes or whole sub-hierarchies of activity types, defined via ontology-based information content (Sánchez, et al., 2011). On the basis of such a distance metric, the distance between non-weighted/weighted sets of activity types is then calculated and activities of short distance are retained and included in the recommendation (Pesquita, et al., 2007). Overall, similarity is therefore calculated via a distance metric between the ontological description of each activity, where the description is represented as a non-empty set of tree nodes taken from the activity type ontology.

More specifically, the first recommendation engine performs recommendations on the basis of functions and relationships defined over the sets of activity types, activities, users and activity evaluations by users. Initially, the available activities are collected, consisting of all the activities which conform to the trip's time and location restrictions and the user's language restrictions. Then, the user's past activity ratings are fetched. For each of the available activities, the most similar set of rated-by-the-user activities are estimated. Each of the available activities' recommendation weight is calculated as a function of the distance between itself and the most similar rated activities and the mean rating of the rated activities. One advantage of this approach for generating recommendations is that ratings for activities are not required from other users, since only the user's own ratings are used. Another advantage is that a large part of the calculations can be pre-computed off-line, since the activity descriptions change infrequently and as a result the distance between the activities remains unchanged. The disadvantages are that the user needs to provide ratings for some activities and that the other users' ratings are not taken advantage of. The former is improved by deducing ratings from a similar visit profile. The latter is addressed by the second recommendation engine.

The second engine performs a variation of collaborative filtering recommendation. It suggests user activities by clustering visit profiles via top-down clustering and suggesting activities rated by other cluster members to members of the cluster. The similarity of the users is calculated via the distance between the ontological description of the visit profile preferences and the similarities in age, gender, spoken languages and scheduling preferences. This engine takes advantage of the ontological information available for each visit and profile as well as the activity ratings of other users. As before, the set of available activities is collected. Afterwards, the user's cluster is employed as a proxy for the user's ratings. For each available activity, if the activity has been rated by one or more members of the cluster, the activity's recommendation weight is assigned as the mean of the other members' ratings. If an activity has not been rated by any of the cluster members, the cluster's aggregate preferences are used to rate the activity, behaving as a virtual cluster-average user, but weighted with a factor signifying the diminished confidence in this approach.

Among the advantages of the second engine are the exploitation of other users' ratings and the fact that a large part of the calculations, but not all, can also be pre-computed as clusters, should be relatively stable and the cluster's aggregate preferences need not be frequently updated. Additionally, this engine also takes advantage of user profile preferences, which are updated from their initial values using machine learning techniques on the user provided feedback. The most important, though, is that the prior availability of user ratings is not a prerequisite for the system to make recommendations.

The main disadvantage is the increased computational load, given the need to perform user clustering and that users need to belong to a cluster. However, this is not a major concern, since user clusters are formed and adjusted off-line, by periodically recalculating the clusters, while the prior definition of default representative user profiles enables usage by new users.

In the final merging stage the outputs of each of the two engines are combined. Each engine produces an independent list of (Activity, Weight) tuples. The merging function expresses the confidence in each engine by examining the richness of the information processed by each engine, such as user profile preferences genericity, ratings, cluster size, cluster virtual profile preference genericity, and weighs the two lists accordingly. Finally, the list is returned ordered from the most to the least recommended activity.Some parts of the user model are also used in an auxiliary manner to filter recommended activities out before inputting them into the recommendation engines. Age will filter age-inappropriate activities and spoken languages will filter out activities performed in unfamiliar languages.

One of the problems many systems with explicit user profile preferences have is the lack of user engagement in defining their preferences. Therefore, user profile preferences tend to be generic, neither strongly preferring nor strongly disliking anything. A remedy to this adopted by the present approach is to perform non-intrusive learning of these preferences by logging user choices during system usage, such as selecting, deselecting and viewing activities as a proxy for actual ratings. Obviously, direct user feedback in the form of plan and activity ratings is considered more significant, therefore the information gleamed in this manner is appropriately weighted such that the low confidence in these measurements is appropriately represented.
The recommendation subsystem executes the off-line calculations using the Apache Mahout machine learning library on the Apache Hadoop MapReduce framework.

Privacy Considerations

Storage of large amounts of data concerning user interests, travels, preferences and behaviours is a significant problem for both the user and the service provider who stores this data. The users risk having their private and potentially sensitive data misused. The service provider incentivizes more attacks against itself since more data are to be gained by unlawfully acquiring it, and is also potentially liable for any data theft. At the same time, the recommendation subsystem requires the availability of large amounts of data to be able to function. We have attempted to reach a trade-off, which allows the recommendation subsystem to deliver its intended functionality effectively, while at the same time increasing the users' privacy protection and diminishing the potential for large-scale data exfiltration. The penalty for this decision lies in increased implementation complexity, higher computational overheads and optionally, shifting some of the privacy protection burden to the users.

To enhance data protection, apart from the obvious security measures (e.g. access control, logging, auditing), user data which is deemed sensitive is kept in encrypted form in the database. The data is transparently decrypted whenever the user logs into the system, and is kept decrypted for the duration of the user's session and then re-encrypted automatically. The data in the database is encrypted using a symmetric cipher. The symmetric key is itself encrypted using another cipher, using the KEK (Key Encryption Key) scheme (Landrock, 2005), to allow changing user encryption keys without needing to decrypt the data and re-encrypt with the new key. The data which is considered sensitive and thus protected by the privacy mechanism in myVisitPlanner is shown in Table 1, against the main processes where it is accessed and the entities that need access to the data. At this stage the system allows access to the user data to all entities, when the user is logged in. An additional protective measure could be to limit the access of each entity to the data needed for the processes they perform.

Table 1.Data usage in myVisitPlanner processes

Entity	**User**	**Recommendation**			**Scheduler**
Scope / **Data Item**	**Profile Editing (UI)**	**Activity Similarity Based Recommendation**	**User Clustering**	**User Cluster Based Recommendation**	**Scheduling**
Demographic Data	■		■		
Activity Type Preferences (in User Profile)	■		■		
System Preferences (in User Profile)	■				■
Detailed User Interaction Log	■				
Activity Ratings	■	■	■		

Table 2.System evaluation criteria

General evaluation criteria	
Criteria categories	**Comments**
Response speed	Rating the interfaces response time; rating the response time for obtaining the recommendations and the itinerary plan
User friendliness	Rating user friendliness for different user groups, considering even non-expert users
Robustness	Rating robustness to erroneous data entry, work load, etc.
Functionality	Rating of the extent to which the system covers specified functionality (without taking into account the quality of the recommendations and plans)
Maintainability	Rating of the extent to which the system is capable of handling content updates (in the first instance) and prospects for integration with other systems
Specific service-targeted evaluation criteria	
Criteria categories	**Comments**
Quality of offered recommendations	Rating the relevance of the offered activity recommendations to the recommendation request circumstances (e.g. preferences, profiles, etc.)
Quality of itinerary scheduling	Rating the quality of the scheduling of the proposed itinerary plans
Quality of content	Rating the quality of the offered content with respect to available activities

Evaluation Methodology, Next Steps and Conclusion

MYVISITPLANNER is available at www.myvisitplanner.com and is currently undergoing beta testing and piloting by different user category groups, namely activity providers and prospective travellers. Activity providers were selected from the geographic area of Northern Greece. Each activity provider is acting also as a hub for testing the generation of visit plans by prospective travellers. While testing the system, activity providers and prospective visitors are requested to offer their evaluations of the system. To this end, different evaluation questionnaires have been prepared, one for each end user group category. These include questions to assess the offered services according to the criteria shown in Table 2.

Once beta testing, piloting and evaluation is complete and the results are analysed, the MYVISITPLANNER team aims to develop an action plan for moving to the next step, that is the take up of the prototype solution to the level of viable and sustainable service provisioning. To this end great stakeholder engagement will be sought, as it is understood that by the very nature of the MYVISITPLANNER functionality and its placement in the direction of servicing the creative and niche tourism markets, overall success critically depends on the ability to operate in a highly participative, collaborative and eventually crowd-sourced manner. Although the underlying cloud-based recommender and itinerary planning engines are fairly sophisticated, the underlying complexity is hidden from end users, who just have to deal with simple user interaction interfaces. Therefore the system does not essentially require learning new skills apart from basic internet use.

The pilot system currently supports two basic user types, namely visitors and activity providers. While a narrow viewpoint may see the latter group as direct providers of specific visit activities, a broader perspective is to see them as a more abstract category that may engage in the wider value chain of heritage, tourism and hospitality management. Heritage providers may benefit by having their offerings listed in a practical web-based tool that enables other stakeholders and visitors to easily and quickly create and ultimately execute plans that include such offerings within planned visits, thus increasing heritage organisations' reach and impact. Tourism professionals and tour operators are better enabled to offer the right mix of both group and individualised visit options in their services, thus expanding and enriching their portfolio to better match emerging demands for niche and creative tourism. Hospitality professionals may now have the additional advantages offered by enabling prospective visitors to link hospitality options not just with single locations but whole itineraries or even offer additional itinerary planning options for existing visitors. It is interesting that none of the aforementioned user groups need to posses any specific knowledge about how exactly recommender systems actually work, in a way that bears relevance to the way for example customers of Amazon recommender services not required to know anything about the underlying collaborative filtering algorithms running in the background of the recommendation engine. By design the MYVISITPLANNER system may support a wider value chain of stakeholders, that is an extended cluster of heritage, tourism and hospitality professionals. A future system enhancement to link with existing popular social networking tools is a natural and rather straightforward extension. However, the social nature of humans when planning visits can already be indirectly mapped by recording and applying analytics on observed preferences and actions, as well as evolution of activity patterns. The reported work is a step in this direction with many possibilities still open that with further work would push the exploitation potential of cloud-based recommender itinerary planning services.

Acknowledgement

The reported work has been financially supported by the General Secretariat for Research & Technology in Greece (grant 09SYN-62-129). The collaboration with all project partners is gratefully acknowledged and especially with Gnomon Informatics SA, highly involved in the prototyping of the developed solution's front end and various integration aspects, as well as ANKO SA and Ethnological Museum of Thrace for their contribution in activities mapping throughout the MYVISITPLANNER project lifecycle and all involved research staff within the teams of University of Macedonia and ATHENA Research and Innovation Centre in Greece.

References

Agarwal, J., Sharma, N., Kumar, P., Parshav, V., Srivastava, A., and Goudar, R.H. (2013). Intelligent search in e-tourism services using recommendation system: perfect guide for tourist. In *7th International Conference on Intelligent Systems and Control, ISCO 2013*. (IEEE: Tamil Nadu, India). pp. 410-415.

Alexiadis, A., and Refanidis, I. (2013). Generating alternative plans for scheduling personal activities. In *ICAPS 2013 Workshop on Scheduling and Planning Applications (SPARK)* (L.C. Vidal, S. Chien, and R. Rasconi, eds.). (Rome). pp. 35-40.

Alexiadis, A., and Refanidis, I. (2013). Post-optimizing individual activity plans through local search. In *ICAPS 2013 workshop on constraint satisfaction techniques for planning and scheduling problems (COPLAS)* (M.A. Salido, R. Bartak, and F. Rossi, eds.). (Rome). pp. 7-15.

Bobadilla, J., Ortega, F., Hernando, A., and Guitierrez, A. (2013). Recommender systems survey. Knowledge-based systems 46, 109-132.

Cha, K.J. (2014). Experimental evaluation of an expert system for travel recommender systems. International journal of software engineering and its applications 8, 115-128.

Chen, G., Wu, S., Zhou, J., and Tung, A.K.H. (2014). Automatic itinerary planning for traveling services. IEEE transactions on knowledge and data engineering 26, 514-527.

Chen, J.-H., Chao, K.-M., and Shah, N. (2013). Hybrid recommendation system for tourism. In *ICEBE 2013, 10th IEEE international conference on e-business engineering*. (IEEE: Coventry, UK). pp. 156-161.

D' Amico, G., Ercoli, S., and Del Bimbo, A. (2013). A framework for itinerary personalisation in cultural tourism of smart cities. In *AI*HCI 213, Proceedings of the 1st workshop on AI*HCI: intelligent user interfaces*. (Turin, Italy).

Emmanouilidis, C., Koutsiamanis, R.-A., and Tasidou, A. (2013). Mobile guides: taxonomy of architectures, context awareness, technologies and applications. Journal of network and computer applications 36, 103-125.

Gavalas, D., Konstantopoulos, C., Mastakas, K., and Pantziou, G. (2014). Mobile recommender systems in tourism. Journal of network and computer applications 39, 319-333.

Herzog, D., and Wörndl, W. (2014). A travel recommender system for combining multiple travel regions to compose trip. In *Proc. New Trends in Content-based Recommender Systems (CBRecSys) workshop, 8th ACM Conference on Recommender Systems* (T. Bogers, M. Koolen, and I. Cantador, eds.). (ACM: Foster City, CA, USA). pp. 42-48.

Koceski, S., and Petrevska, B. (2013). Advanced tourist trip planning using hybrid recommender. Engineering management reviews 2, 115-123.

Lamsfus, C., Wang, D., Alzua-Sorzabal, A., and Xiang, Z. (2014). Going mobile: defining context for on-the-go travellers. Journal of travel research (in print), 1-11.

Landrock, P. (2005). Key Encryption Key. In Encyclopedia of Cryptography and Security. (Springer), pp. 326-327.

Lucas, J.P., Luz, N., Moreno, M.N., Anacleto, R., and Figueiredo, A.A. (2013). A hybrid recommendation approach for a tourism system. Expert systems with applications 40, 3532-3550.

Moreno, A., Valls, A., Isern, D., Marin, L., and Borras, J. (2013). SigTur/E-Destination: Ontology-based personalized recommendation of tourism and leisure activities. Engineering applications of artificial intelligence 26, 633-651.

Pesquita, C., Faria, D., and Bastos, H. (2007). Evaluating GO-based semantic similarity measures. In *BioOntologies SIG at ISMB/ECCB – 15th Annual International Conference on Intelligent Systems for Molecular Biology (ISMB)*. (Vienna, Austria.

Refanidis, I., and Alexiadis, A. (2011). Deployment and Evaluation of SELFPLANNER, an Automated Individual Task Management System. Computational intelligence 27, 41-59.

Refanidis, I., Alexiadis, A., and Yorke-Smith, N. (2011). Beyond calendar mashups: SelfPlanner 2.0. In *Proceedings of the system demonstrations, ICAPS 2011, 21th international conference on automated planning and scheduling* (P. Bertoli, and M. Do, eds.). (Freiburg, Germany). pp. 66-70.

Refanidis, I., Emmanouilidis, C., Sakelariou, I., Alexiadis, A., Koutsiamanis, R.-A., Tasidou, A., Kokkoras, F., and Efraimidis, P.S. (2014). myVisitPlannerGR: Personalised itinerary planning system for tourism. In Artificial intelligence: methods and applications, A. Likas, K. Blekas, and D. Kalles, eds. (Springer), pp. 615-629.

Ricci, F., Rokach, L., Shapira, B., and Kantor, P.B. (2010). Recommender systems Handbook. (New York: Springer-Verlag).

Richards, G. (2014). Creativity and tourism in the city. Current studies in tourism 17, 119-144.

Sánchez, D., Batet, M., and Isern, D. (2011). Ontology-based information content computation. Knowledge-based Systems 24, 297–303.

Schaller, R. (2014). Mobile tourist guides: bridging the gap between automation and users retaining control of their itineraries. In *IIiX '14, Proceedings of the 5th Information Interaction in Context Symposium*. (ACM: Regensburg, Germany). pp. 320-323.

Schaller, R., and Elsweiler, D. (2014). Itinerary recommenders: how do users customise their routes and what can we learn from them? In *IIiX '14, Proceedings of the 5th Information Interaction in Context Symposium*. (ACM: Regensburg, Germany). pp. 185-194.

Yang, W.-S. (2013). iTravel: a recommender system in mobile peer-to-peer environment. The journal of systems and software 86, 12-20.