PATCH 2013: Personal Access to Cultural Heritage

Liliana Ardissono[1], Lora Aroyo[2], Luciana Bordoni[3], Judy Kay[4], Tsvi Kuflik[5]

1 Dipartimento di Informatica, Università di Torino, Italy
liliana.ardissono@unito.it
2 Web and Media Group, VU University Amsterdam, The Netherlands
l.m.aroyo@cs.vu.nl
3 ENEA, Italy
luciana.bordoni@enea.it
4 School of Information Technology, University of Sydney, Australia
judy.kay@sydney.edu.au
5 Information Systems Department, University of Haifa, Israel
tsvikak@is.haifa.ac.il

The international workshop on Personal Access to Cultural Heritage (PATCH, http://patch2013.wordpress.com/) is a forum for researchers and practitioners who are working on various aspects of cultural heritage and are interested in exploring the potential of state of the art technology to enhance the CH visit experience.

Personalization seems to carry a great potential for cultural heritage (CH), enabling users to cope with ever growing CH material. However, unlike other domains and tasks, where users may have well defined needs and preferences, CH may be a bit challenging as users tend to explore cultural heritage sites as leisure activities, many times with friends and family and are open to new experiences. Hence the question is how to model and support users in this domain, taking into account its specific peculiarities.

PATCH aims at building a research agenda for personalization in CH in order to make the individual CH experience a link in a chain of a lifelong CH experience which builds on past experience, is linked to daily life and provides the foundation for future experiences. The workshop aims to be multi-disciplinary. It is intended for researchers, practitioners, and students of information and communication technologies (ICT), cultural heritage domains (museums, archives, libraries, and more), and personalization.

PATCH 2013 presents novel research ideas, projects and scientific findings focused on the following aspects. The first and largest was personalization of city tours and touristic experiences, bringing lessons learned at the indoors and online museum visitors guide systems to the outdoors – cities and other CH sites and linking them. This is an important new area based on active involvement of a broad range of people to enhance the management of cultural heritage information. In this way the experiences are connect to "daily" lives and the user is not only the information consumer but also the producer. Then another aspect was the introduction of web of things to CH – personalization of visitors' interactions with smart CH objects. A somewhat unusual suggestion for integration of virtual reality into CH related application is also presented and finally, the delicate and challenging matter of how to involve curators in the process of creating personalized information for visitors.

Considering the submissions, it seems that the overall direction is towards integrating the outdoors, indoors and online CH sites via semantic web technologies and internet of things into a continuous personalized CH experience. A key issue is that of participation: to allow the user to participate not only in the consumption of information but also in the interpretation and discussion around the information as well as in the creation of new content. People are no longer passive users or consumers of information but take part in the production, they became "prosumers" of information.

We thank the members of the Program Committee of PATCH 2013 for their support and reviews. Furthermore, we are grateful to all authors who submitted articles and contributed with their works to a fruitful discussion.

PATCH 2013 Program Committee:

Dynamic Personalisation for Digital Cultural Heritage Collections

Cormac Hampson[1], Eoin Bailey[1], Gary Munnelly[1], Séamus Lawless[1] and Owen Conlan[1]

[1] Knowledge and Data Engineering Group, Trinity College Dublin, Ireland
{cormac.hampson, eoin.bailey, munnellg, seamus.lawless, owen.conlan}@scss.tcd.ie

Abstract.
The number of digital collections in the cultural heritage domain is increasing year on year. Improved quality of access to cultural collections, especially those collections which are not exhibited physically is a key objective of the digitisation process. Despite some successes in this area, many digitised collections struggle to attract users or to maintain their interest over a prolonged period. One of the key reasons for this is that users of these archives vary in expertise (from professional researchers to school children) and have different tasks and goals that they are trying to accomplish. This paper describes CULTURA, an FP7 funded project that is addressing this specific issue through its four-phase personalisation approach and accompanying suite of services. By employing such personalisation techniques, CULTURA is helping the exploration of, linking to, and collaboration around cultural heritage collections.

Keywords: CULTURA, Digital Humanities, Personalisation, Adaptation, Exploration; Communities

1 Introduction

Recent years have seen a huge growth in the digitisation of cultural heritage collections. However despite this increased availability of cultural archives, curators can still struggle to instigate and enhance engagement with these collections. Simple "one size fits all" web access is, in many cases, not appropriate in the digital humanities, due to the size and complexity of the artefacts. Furthermore, different types of users need varying levels of support, and every individual user has their own particular interests and priorities. Personalised and adaptive systems are thus important in helping users gain optimum engagement with these new digital humanities assets.

Without improved support for personalisation, digital cultural heritage collections will struggle to reach their full potential. This paper describes the CULTURA project [1][2] which is directly addressing this problem. Specifically, CULTURA employs personalisation to help support exploration of digital collections, the collaboration of users around these collections and to understand a user's interests in external digital archives with similar content. Section 2 outlines some background research on personalisation in digital cultural archives, and describes the aims of the CULTURA project. Section 3 introduces CULTURA's four-phase personalisation approach; section 4 outlines its user model; with section 5 highlighting the use of community adaptation. Finally section 6 discusses how CULTURA can link to external resources; with section 7 summarising the paper.

2 Background and the CULTURA Project

While there have been recent attempts to use Adaptive Hypermedia techniques to support the personalised retrieval, interrogation and presentation of cultural heritage content collections, these have to-date been limited. The MultimediaN N9C Eculture project[1] aims to provide multimedia access to distributed collections of cultural heritage objects. It is an aim of the project to support the generation of various types of personalised and context-dependent presentations of cultural material. However, the current system only provides static semantic search across entities in manually annotated content collections.

The CHIP project[2] aims to provide personalised presentation and navigation of the Rijksmuseum cultural resources. The Artwork Recommender supports the rating of artworks/topics to generate a user profile, which is then used to drive future artwork recommendations. The Tour Wizard is a web-based tool which uses the user profile to semi-automatically generate personalised museum tours. In the MOSAICA[3] project a mobile device-based demonstration is used to engage novice and intermediate users. The system does provide virtual visitors with access to structured descriptions of collections through a search interface, but little adaptivity or personalisation of the experience is used.

The QViz[4] project has some similarities in approach to the CULTURA project in that it makes explicit recognition of the value of users as members of communities, and as contributors to digital cultural heritage collections. The focus of the QViz system is on temporal and spatial search and retrieval of archival content. While QViz is a social semantic application, facilitating user contribution and structured representation of knowledge, it does not have a personalised or adaptive aspect.

Because CULTURA is producing a generalisable solution, it must be able to add value to a wide range of digital cultural heritage collections, of which there are many. One example is the Europeana project[5], which represents metadata from collections across many EU member states. While Europeana does not directly host content, it is a large repository of metadata which could be processed, alongside a specific collection's content, to seed the CULTURA environment. Many projects within the cultural heritage domain, including PATHS[6] and Natural Europe[7], already encompass

---

[1] http://e-culture.multimedian.nl/
[2] http://www.chip-project.org/
[3] http:// www.mosaica-project.eu
[4] http://www.qviz.eu/
[5] http://www.europeana.eu/portal/
[6] http://www.paths-project.eu/

rich metadata from Europeana within their environments.

Improved quality of access to cultural collections, especially those collections which are not exhibited physically, is a key objective of the CULTURA project. Moreover, CULTURA supports a wide spectrum of users, ranging from members of the general public with specific interests, to users who may have a deep engagement with the cultural artefacts, such as professional and trainee researchers. To this end, CULTURA is delivering a corpus agnostic environment, with a suite of services to provide the necessary supports and features required for such a diverse range of users.

Within the CULTURA project, two rich cultural archives are being used to showcase the features offered by the environment. The 1641 Depositions are seventeenth-century manuscripts that comprise over 8,000 witness statements, relating to the Irish rebellion of 1641; and the Imaginum Patavinae Scientiae Archivum (IPSA) collection is a digital archive of illuminated astrological and herbal manuscripts from the 14th century. The IPSA manuscripts have the rare characteristic of containing high quality and very realistic illustrations, and the archive consists of digitised images and related metadata descriptions. As such, from a technical perspective, IPSA represents a very different kind of digital humanities collection to the 1641 Depositions. Importantly, initial evaluations of the work done with both the IPSA [3] and 1641 collections [4] have yielded positive results.

3 Four-Phase Personalisation Approach

The employment of adaptation techniques can help empower experienced researchers, novice researchers and the wider community to discover, interrogate, and analyse cultural heritage resources. Hence, core elements of the CULTURA architecture are its personalisation methods. The techniques employed by CULTURA have been heavily influenced by Adaptive Hypermedia (AH) and Adaptive Web systems research, which are concerned with improving the retrieval and composition of information. This improvement is achieved by creating a more context-sensitive and personalised interaction with digital content, and is often predicated on rich metadata [5].

One reason why novice users struggle to engage with large cultural collections is a lack of guidance when they initially encounter a set of resources. Likewise, more experienced researchers often lack the tools to efficiently search, share, visualise, analyse and correlate data from important cultural collections. To help counteract such issues, CULTURA employs a four-phase personalisation approach (see figure 1). Each of these phases (*guide*, *explore*, *reflect* and *suggest*) are now described in turn.

[7] http://www.natural-europe.eu/

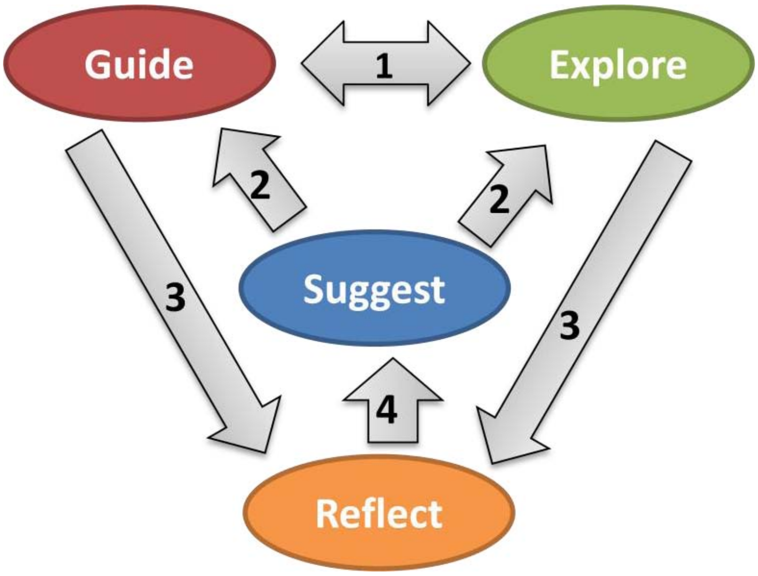

Fig. 1. CULTURA's Four-Phase Personalisation Approach

3.1 Guide

Users with little experience of the underlying resources typically start their investigations within the guide phase. Here CULTURA employs a "narrative" module, which enables resources within the collection to be sequenced on a specific theme e.g. the chain of derivation of illustrations in the IPSA collection, or the reliability of witness statements in the 1641 Depositions. Furthermore, how these resources are rendered to the user (text, visualisation etc.) can also be specified within the narrative metadata (encoded as XML). Guiding users through a collection is facilitated through parameterised launch of services using URLs, and is especially useful in providing users with a path through specific content (see figure 2). Importantly, the narratives can be adjusted in length, either explicitly by the user (by choosing to see more resources on the theme) or implicitly by the narrative module, which analyses the user model for changes in user interest. An authoring tool based on SABer [6], which helps non-technical people to encode domain expertise, is planned to assist the creation of these narratives.

Being on a guided path does not limit a more adventurous user, as they can use these sequenced narratives as a springboard for their own investigations (in fact many narratives explicitly encourage users to do so). Within the four-phase personalisation approach, this involves stepping from the *guide* phase to the *explore* phase (number 1 in figure 1). Importantly, by monitoring the user model metadata (perceived level of interest in particular concepts/entities etc.) as they explore the resources, the narrative path itself can be adapted. This is achieved by selecting documents for their path that most closely match the user's interests, and can result in the path being enriched with further resources and concepts.

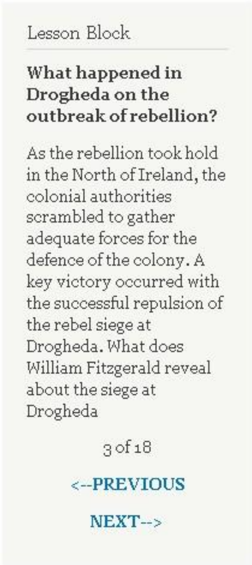

Fig. 2. Guided Lesson Plan within CULTURA

3.2 Explore

In the *explore* phase CULTURA offers tool assistance (e.g. data enriched maps, entity based search, social network analysis) to support exploration and browsing of the underlying resources. At any stage a user can return to where they left their path in the *guide* phase, and users with little prior knowledge of the resources often flick between the *explore* and *guide* phases several times. In contrast, professional researchers with a deep understanding of the collection typically spend the majority of their time within the *explore* phase and may never involve themselves with the *guide* phase.

The types of services that CULTURA offers during the *explore* phase include normalised search [7], entity based search [8] and geographic visualisations (see section 6). Another useful service offered is shown in figure 3, where a visualisation of the social network described within a single document of the 1641 depositions is depicted. Among the entities it depicts are the people (including their role e.g. rebel, victim, landlord etc.) and the crimes that are documented within that deposition. By clicking on an entity in this graph (represented as one of the circles in the visualisation) a user can view all other entities in the collection that are associated with that entity. This allows a user to quickly find connections between resources, which might not be so easily apparent using text searching alone. Significantly, due to the service-based architecture used by CULTURA it means that the suite of services offered within the *explore* phase can be extended iteratively over time, allowing new features to be offered to existing and future collections.

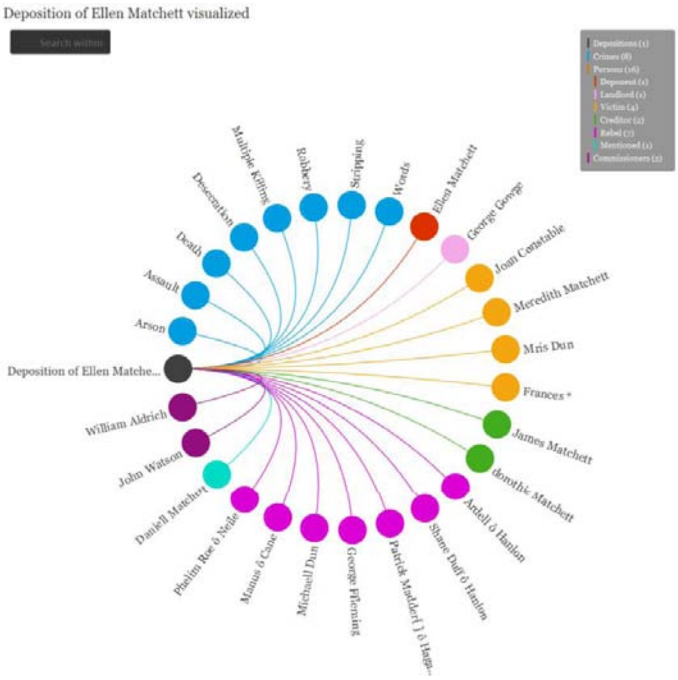

Fig. 3. Social Network Analysis in the "Wheel"

3.3 Suggest

Whether, within the *guide* or *explore* phase, a user will be given personalised suggestions for related content (see figure 4) or tools to view resources in. This process occurs in the *suggest* phase (which works in parallel with the *guide* and *explore* phases), with hints pushed to users for their review (number 2 in figure 1). These hints are influenced by the content the user is currently viewing, as well as the data stored in their user model e.g. search terms, entities commonly viewed, annotations created etc.

Figure 4 shows one example of the recommended content shown to users who browse the 1641 Depositions using CULTURA. When a user views a deposition, entities (people, places etc.) are extracted from the text, and complementary depositions that also mention these entities are located within the collection. For instance, in figure 4, blue text links to depositions related to the entities “Trim”, “Meath” and “Lismore” are displayed. By showing the entity names beside the recommended links, it makes clear to the user why these links are being rendered to them. The recommender box is displayed beside the deposition text and enables users to quickly locate new resources that are relevant. Furthermore, by clicking on one of these links, the user model is updated with the corresponding entity, as it indicates a user interest in said entity. Other services that are triggered in the explore phase include the search module, where the results presented to users are influenced by terms stored in a user’s model.

Fig. 4. An example of the recommended content displayed to users

3.4 Reflect

At any stage within the *guide* or *explore* phases, a user may enter the *reflect* phase (number 3 in figure 1), which involves viewing a subset of their user model (rendered as a tag cloud) and seeing what terms are influencing the recommendations they are receiving. Moreover, they can easily delete terms they deem irrelevant or increase/decrease the relative size of terms depending on how influential they think they should be. New terms will also be allowed to be manually added. Figure 5 shows the current user interface for interacting with a user model. Such scrutability [9] of user models is vital for making the underlying personalisation processes more transparent, and gives users a greater feeling of control over the adaptation features offered by CULTURA. Importantly, any changes made by users during the *reflect* phase directly impacts on the *suggest* phase and the recommendations that eventually filter down into the *guide* and *explore* phases (number 4 in figure 1).

Fig. 5. User Model Tag Cloud

4 User Model

By espousing the four-phase personalisation approach, CULTURA dynamically adapts to users, and renders useful suggestions to them at appropriate times. Moreover, this process provides mechanisms that are appropriate for a range of users with different levels of ability or different interests in the underlying resources. A central component of this process is CULTURA's user model, which is stored in MongoDB[8], a scalable high-performance NoSQL database. One of the main benefits of MongoDB is that it allows the database schema to evolve over time, which is very useful in a service orientated environment such as CULTURA, where new suites of tools are being introduced gradually.

A model of each user is built silently as a user interacts with the system, however, as described in section 3.4, each user will be given the opportunity to scrutinise and change their model. All actions a user performs are recorded in order to build up detailed information on each user. This includes viewing, bookmarking or annotating a deposition, as well as the searches performed, entities commonly viewed and the visualisations they rendered. Further user information is to be incorporated into the model such as the relative expertise of a user, and the communities they are a part of. By analysing the entities that a user has annotated it will also provide better evidence of a user's interest.

Important elements of CULTURA's user model to be developed in the next phase of the project include allowing its partition into different parts (users can choose to reset their user model when tackling different tasks or projects), as well as to have

[8] http://www.mongodb.org/

both short-term and long-tem models of interests. By storing such rich details in the user model, it will enable a better understanding of the user and ultimately better recommendations to be made to them.

5 Community Adaptation

Complementing the user model information stored in MongoDB are the comprehensive logs that are stored in Drupal's MySQL database. These provide further rich metadata (such as time-stamped logins and page views) that can be exploited by the system for personalisation and adaptation purposes. This supplementary information is being combined with the user model and can be displayed to the user so that they have a record of the actions they have performed in the system. There will also be the option to share this data with your community and to view what others have been doing in the system (see figure 6). Such features help encourage community awareness within the CULTURA environment.

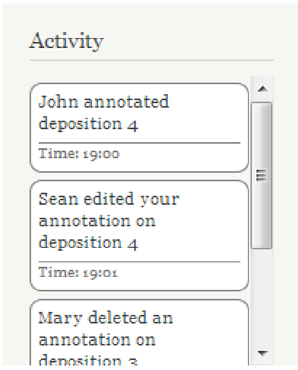

Fig. 6. Prototype Activity Feed in CULTURA

A central component of CULTURA is its annotation service CAT [10][11], which enables the annotation of both text and images within the CULTURA environment. Further to allowing a user to comment on document text and images, the annotations created using CAT allow a user to link their annotations to other resources internal and external to the CULTURA environment. While CAT is beneficial for researchers and educators, it is also being used as an important source of user metadata for CULTURA. For instance, annotations provide an insight into which entities are of interest to a user. If a user is frequently annotating a document, it is likely that this document is of interest to them. Furthermore, by analysing the text being annotated (as well as the text of the annotation itself) using entity extraction; it is possible to discern specific items of interest to the user. Other data that the annotation tool allows to be collected is the interaction of users within a group, and how often a person views annotations by a specific contributor. This data can improve the recommendations given by CULTURA, as well as help foster collaboration within the environment. However, the development of adaptive user communities is a challenging problem. For instance, the different user types require different degrees of sharing and

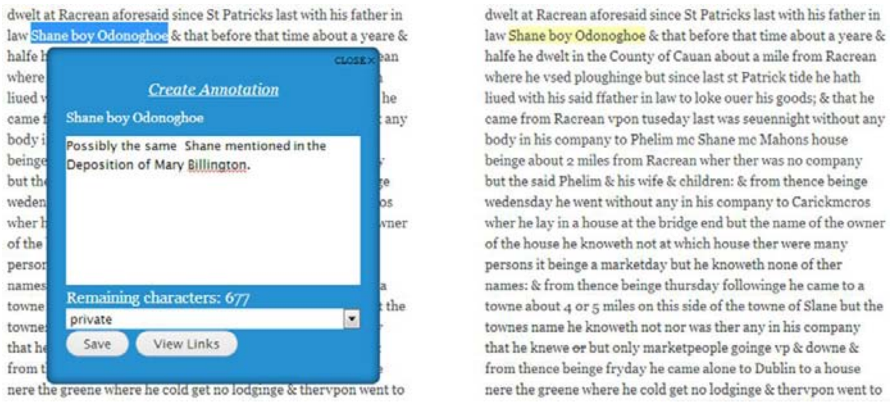

Fig. 7. Screenshot of CAT within CULTURA

grouping, e.g. a professional researcher generally does not want to share their annotations prior to publication of research whereas an interested member of the public often will have no such issues.

With the increased richness of user models and logs, coupled with increased usage of the system, it means that it will also be possible to establish recommenders for similar users in both the 1641 and IPSA collections. This will greatly benefit those wishing to find other users to collaborate with, or relevant communities to join. Importantly, by analysing the members of specific communities, it will be possible to determine common characteristics shared in their individual user modules. This gives a better understanding of the community as a whole. Furthermore, with increased community use of CULTURA, it means that Influence Network Analysis (a form of social network analysis that identifies the most influential nodes within a graph) can be performed over these communities. The metrics used to rank an entity's influence in its community can also be used to adjust the adaptations received. Enhanced community-based adaptations, which account for items such as a group's type (students, professionals, a mixture etc.), size, diversity, activity levels, and length of time in existence, are also a priority.

6 Links to External Resources

If a data collection contains entities that are referenced on the web of data, it is possible to map between these entities so that external information on that entity can be incorporated into the CULTURA environment. For example, the 1641 depositions contain many place names which are linked to the geonames[9] database. An example of what can be achieved with such an approach is shown in figure 8, where depositions from the 1641 collection are mapped to a specific geographical area and visualised on in interactive display. By creating links between CULTURA collections and

[9] http://www.geonames.org/

the wider web, it helps to increase traffic to archives and provide a richer landscape for digital humanities research. This approach of enriching CULTURA with data from open corpus content will be further examined in future.

If the external collections use dereferenceable URIs for entities (as Europeana[10] does), a complementary process can occur, where data stored in CULTURA's user model is used to help deliver more relevant recommendations for resources within these external archives. Furthermore, when new collections are added to CULTURA itself, and they contain similar content and entities as existing archives, a user's model can help to determine relevant resources from the new collection, as well as assisting against the *cold start* problem. Thus, incorporating personalisation into CULTURA's interactions with external collections has great potential, and can greatly help in improving a user's experience with new cultural heritage archives.

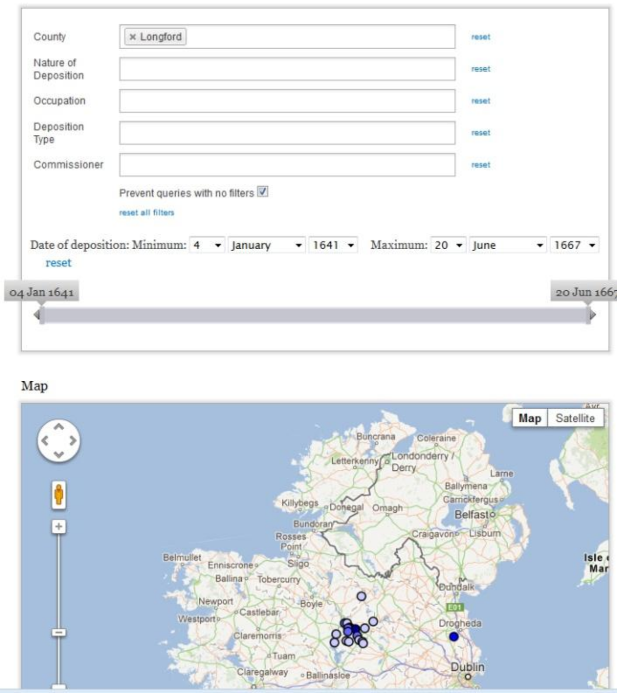

Fig. 8. Visualisation of deposition locations within CULTURA

7 Summary

This paper has described how personalisation is being used in the CULTURA project to help users of different skills and level of interest, to engage more meaningfully

[10] http://pro.europeana.eu/linked-open-data

with cultural heritage collections. CULTURA's four-phase personalisation approach was introduced along with details of how user models and logs are deployed within its environment. CULTURA's focus on community adaptation and links with external data collections were also outlined. Detailed evaluations, which include all relevant stakeholders from the domain are ongoing, and will feedback into the refinement of the personalisation technologies and methodologies employed by CULTURA

Acknowledgments. The work reported has been funded by the Seventh Framework Programme of the European Commission, Area "Digital Libraries and Digital Preservation" (ICT-2009.4.1), grant agreement no. 269973.

References

1. Hampson, C., Lawless, S., Bailey, E., Yogev, S., Zwerdling, N. and Carmel, D. 'CULTURA: A Metadata-Rich Environment to Support the Enhanced Interrogation of Cultural Collections'. In the Proceedings of the 6th Metadata and Semantics Research Conference, Cádiz, Spain, pp. 227-238, Springer, Heidelberg (2012)
2. Hampson, C., Agosti, M., Orio, N., Bailey, E., Lawless, S., Conlan, O., and Wade, V. 'The CULTURA project: supporting next generation interaction with digital cultural heritage collections', In the Proceedings of the 4th International Euromed Conference, Limassol, Cyprus, pp. 668-675, Springer, Heidelberg (2012)
3. Ponchia, C., 2012.Engaging the User: Elaboration and Execution of Trials with a Database of Illuminated Images. Taking Care of Legacy and Cultural Heritage Documents – Proc. of the 8th Italian Research Conference on Digital Libraries, [In Press] (2013).
4. Bailey, E., Lawless, S., O'Connor, A., Sweetnam, S., Conlan, O., Hampson, C. and Wade, V. "CULTURA: Supporting Enhanced Exploration of Cultural Archives through Personalisation". In the Proceedings of the 2nd International Conference on Humanities, Society and Culture, ICHSC 2012, Hong Kong, China. October 27th-28th (2012)
5. Brusilovsky, P., Kobsa, A., and Nejdl, W. (Eds.) 'The adaptive web: methods and strategies of web personalization' Lecture Notes in Computer Science, LNCS, vol. 4321 (2007)
6. Hampson, C. and Conlan, O., Facilitating Casual Users in Exploring Linked Data through Domain Expertise, DEXA 2011, Toulouse, France, pp319-333 (2011)
7. Lawless, S., Hampson, C., Mitankin, P., and Gerdjikov, S.,'Normalisation in Historical Text Collections', In Proceedings of Digital Humanities 2013, Lincoln, Nebraska, USA [In Press] (2013)
8. Yogev, S., Roitman, H., Carmel, D. and Zwerdling, N. 'Towards expressive exploratory search over entity-relationship data'. In Proceedings of the 21st international conference companion on World Wide Web, New York, USA, pp.83-92 (2012)
9. Kay, J. "Scrutable adaptation: Because we can and must." In Adaptive hypermedia and adaptive web-based systems, pp. 11-19. Springer Berlin Heidelberg, (2006.
10. Ferro, N., Munnelly, G., Hampson, C. and Conlan, O., Fostering Interaction with Cultural Heritage Material via Annotations: The FAST-CAT Way, In Proceedings of the 9th Italian Research Conference on Digital Libraries, Rome, Italy [In Press] (2013)
11. Munnelly, G., Hampson, C., Ferro, N., and Conlan, O., 'The FAST-CAT: Empowering Cultural Heritage Annotations', In Proceedings of Digital Humanities 2013, Lincoln, Nebraska, USA [In Press] (2013)

Traditional Sports and Games: A New Opportunity for Personalized Access to Cultural Heritage

Mariate Linaza[1], Kieran Moran[2,3], and Noel E. O'Connor[3]⋆

[1] Dept. of eTourism and Heritage, Fundación Centro de Tecnologas de Interacción Visual y Comunicaciones – Vicomtech-IK4, Spain
[2] Applied Sports Performance Research Group, School of Health and Human Performance, Dublin City University, Ireland.
[3] CLARITY: Centre for Sensor Web Technologies, Dublin City University, Ireland

Abstract. Sport is the most universal of cultural pursuits – it is accessible and of interest to all. Traditional Sports and Games (TSG) are as diverse as our cultures. TSG organisations work tirelessly to promote participation in their sports, but also act as custodians of custom, language and history. However, trends in globalisation have led to a convergence of the majority of spectator interest to just a few mainstream sports with culturally homogonous identities. In this position paper, we present the case for preservation of TSG using both existing state of the art 3D digitisation technology but also highlight the need to develop low-cost personalisable solutions. This dual approach would potentially allow the styles of play of elite sportspersons (national heroes) to be captured with precision for posterity, and amateur sportspersons (local heroes) to be captured with inexpensive setups to ensure personalized accessible solutions.

Keywords: Intangible cultural heritage, motion capture, 3D digitisation

1 Introduction

The history of sport can be traced back to the existence of human civilisation itself. It is a key part of cultural identity, and a mechanism for the protection and promotion of cultural diversity. Thus, retaining knowledge of our traditional sporting practices is vital in terms of preservation and promotion of sport as an expression of *Intangible Cultural Heritage* [1][2]. Worldwide there is a staggering cultural richness of indigenous, traditional, historical, and regional folk sports and games from different nations and ethnic minorities, many of which are fascinating not only for their differences, but also for the similarities of shared common features. The Encyclopaedia of World Sport[4] includes over 3000 traditional

⋆ The research that lead to this paper was supported in part by the European Commission under contract FP7-ICT-601170 (RePlay), `http://www.fp7-replay.eu`

[4] See: http://www.sportencyclopedia.com/

sports and games from all around the world. European examples include Gaelic games (hurling, football, camogie, handball) in Ireland, Basque Pelota (and local variants) in Spain, Longue Paume in France, Tamburello in Italy among many others. The importance of traditional sports and games to our cultural heritage and the need for preservation has been officially recognized by UNESCO[5]. Despite this, this important cultural heritage domain has yet to receive significant attention from the research community.

2 The Threat of Globalization

Many traditional sports and games are already lost and those that have survived are in danger of disappearing, owing to the various tendencies of globalisation, and convergence in the rich diversity of sports heritage toward a small set of culturally homogenous sports. This has occurred because of the globalisation of communication networks by an ever decreasing number of media organisations whose legitimate aim is to gain as large a readership or viewership as possible via minimum expenditure. This is achieved by appealing to the mass market which favours mainstream sports. In this environment, soccer has become the dominant world spectator and participative sport. The "Beautiful Game" however is culturally homogenous and its identity is arguably indistinct.

In contrast, traditional sports almost exclusively have an amateur ethos and their associations lack significant financial resources, albeit with a small number of notable exceptions. In conjunction with this amateur ethos, traditional sport associations are altruistic in their support of other traditional sports. This rich collective heritage is managed by sporting organisations and governing bodies with few salaried employees and relies on volunteers working tirelessly to retain and promote their sport as a social duty. In addition, traditional sports and their organisational structures are a central nurturing hub to retain, promote and pass on more widespread non-sporting aspects of cultural heritage (e.g. language, dance, music, storytelling).

3 National and Local Heroes

Interest and participation in a particular sport or sports club is significantly dependent upon a person's affinity to their sporting heroes. But national sporting heroes are not only a means of promoting participation; they themselves are part of the cultural heritage. They have unique patterns of movement and skill execution, handed down through successive generations, that have taken years to master. Unlike other forms of cultural heritage, there is no tangible artefact to hold (e.g. a sculpture, a relic) but the grace, fluency and precision of their movements are evident. A particular flick of the wrist, a strike of the ball, or a side-step on the field, the synergy between players all constitute unique movement signatures that are as worthy as the brush stroke of a painter. There is a need to

[5] See: http://unesdoc.unesco.org/images/0014/001428/142825e.pdf

capture and preserve these movement signatures. However, as heroes have become more international, the sport itself loses its attachment to a local/national cultural heritage and thus there is a clear need to continue to promote the local heroes within “grass roots” communities.

(a) Full body motion reconstruction using a small number of cheap accelerometers, from [3]. Motion is synthesized by using low-dimensional signals from wearable accelerometers to index into a motion graph created using optical motion capture.

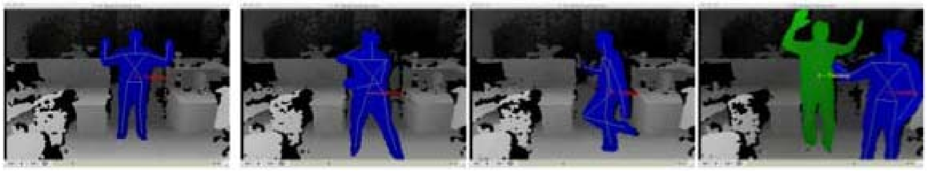

(b) Kinect-based sensing of human motion using publicly available APIs.

(c) Automatic comparison of movements between a dance teacher and student, from [4]. A dancer’s “score” is computed based on the modulus of the quaternionic correlation coefficients for joint positions and velocities combined with 3D motion vector analysis.

Fig. 1. Illustration of various technology components that can be leveraged to perform low-cost motion capture, analysis and visualisation

4 Current Technology: A Partial Solution but a Way Forward

To increase engagement, participation and interest in national, regional and local sports, and their associated cultural heritage, connections need to be forged between the sport and their potential membership. This may be achieved by creating interactive media of local/regional/national heroes that gains a person’s interest, especially children. Cultural heritage would be further enhanced if information on these heroes is retained to span generations. Furthermore, the technology should encompass the principle that children want to **"play like"** and **"play with"** their heroes, strengthening traditional sports' already strong focus on participative sport irrespective of gender.

Whilst capture of national heroes and other aspects of traditional games is non-trivial, the technology exists – high-cost optical motion capture is used to great effect to capture subtleties of human motion, not just in sport, but also in the movie industry[6]. Thus, what is needed here is a considered and comprehensive capture strategy. However, considering their cost, such solutions are not realistic for TSG local heroes within amateur voluntary organizations. Unfortunately, low-cost solutions that can be personalized for local heroes do not yet exist and this is a key challenge for the research community.

To address this, we propose the vision of a low-cost capture and visualization platform suitable for use in TSG, in which the tools and visualization mechanisms can be personalized to an individual user. We propose a virtual immersive environment in which an end user can manifest him/herself via full 3D reconstruction or via an artificial avatar that can be personalized based on his/her preferences. Within this environment, the user performs key movements guided by a 3D representation of the national hero that has been captured using full body scanning and optical motion capture and that is displayed via high-fidelity rendering. In this way, the user *plays with* his/her hero. Of course, in theory the hero need not even be a currently popular athlete, but could be a rendered version of a recognised master practitioner from the history of the game in question captured based on imitating his/her characteristic style of play (a process that would need to be guided by the cultural custodians of the sport). Real-time feedback should be provided based on detailed comparison of the user's movements to their hero indicating how well they *play like* their hero i.e. how close their movements are to those of their hero, where the differences are, how substantial these are, etc.

Recent research results indicate that much of what is needed for the various different components required to realize such a platform is either already available or starting to mature. We outline some of the key technology components required for this in the following and show some illustrations of potential solutions from existing work in Figure 1.

- **Low-cost motion capture using emerging cheap sensor devices, such as MS Kinect and wearable accelerometers**. Significant advances have been made recently in using cheap wearable inertial sensors to sythesize full body motion tailored to an individual and some of these can achieve very low joint angle error [5][3]. Of course, motion capture using the Kinect is currently an extremely hot topic in computer vision.
- **3D photo-realistic visualization of human's performing motion**. The Kinect sensor provides a cheap way not only to track motion but also a way to create full 3D mesh-based reconstructions of humans. Using multiple

[6] E.g. Optical motion capture systems from Vicon, http://www.vicon.com/

Kinects, techniques have been developed that can start to approach the quality of expensive full body scanners [6].

- **Algorithms to compare human motion from low-cost sensors to gold standard**. Two ACM Multimedia Grand Challenges[7] have been run in recent years that call for solutions for comparing dance movements between teachers and students based on an extensive dataset of human motions [7]. A number of different solutions have shown the feasibility of this [4][8].
- **3D scanning of sport implements (e.g. racquets, balls, etc) and apparel**. 3D scanners are becoming more common with even desktop products now available at reasonable cost[8].
- **Avatar authoring tools**. Web-based tools are starting to appear to enable even inexperienced users to create avatars of themselves using a small amount of manual editing. A good example is the RAAT (REVERIE Avatar Authoring Tools) produced by the REVERIE EU project[9].
- **Navigable virtual immersive environments for optimum end-user experience**. For this we can take inspiration from many cultural heritage research projects that have investigated the use of virtual or augmented reality – the Archeoguide project is one of the earliest and most notable successes in this regard [9]. More recently, computer vision advances for 3D imaging and increasing availability of 3D viewing technology makes it feasible to create extremely realistic representations of real-world environments that can be interfaced with computer game engines for ease of navigation, e.g. [10].

5 The Opportunity

In this position paper, we present the case for the urgent need for research effort towards preservation of TSG using existing state of the art 3D digitisation technology whilst also highlighting the need to develop low-cost personalisable solutions. The former would potentially allow the styles of play of elite sportspersons (national heroes) and the evolution of the games and their accoutrements to be captured with precision for posterity using existing high-end 3D capture equipment. The latter, on the other hand, would allow amateur sportspersons (local heroes) to be captured with inexpensive setups to ensure personalized accessible solutions that truly engage the local community.

Recent technological advances mean that it is now potentially possible to digitally capture various aspects of traditional sports although further research is required to produce robust low-cost platforms that integrate these existing techniques in an effective manner. In conclusion, we believe that there is a key opportunity for technology development for capturing a key aspect of our collective cultural heritage that has been overlooked to date. Furthermore, we should not be complacent in this task. Some TSG are under real threat – a variant of

[7] See: http://www.3dlife-noe.eu/3DLife/emc2/grand-challenge/
[8] E.g. NextEngine 3D Laser Scanner, http://www.nextengine.com/
[9] See: http://www.reveriefp7.eu/resources/demos/reverie-avatar-authoring-tool/

Longue Paume died out only in the last 20 years or so – so the time to act is now.

References

1. F. Lenzerini, "Intangible cultural heritage: The living culture of peoples," *European Journal of International Law*, vol. 22, pp. 101–120, 2011.
2. M. L. Stefano, P. Davis, and G. Corsane, Eds., *Safeguarding Intangible Cultural Heritage: Touching the Intangible*. Boydell & Brewer, 2012.
3. P. Kelly, C. Ó Conaire, and N. E. O'Connor, "Human motion reconstruction using wearable accelerometers," in *SCA 2010 – ACM SIGGRAPH Eurographics Symposium on Computer Animation*, Madrid, Spain, July 2010.
4. D. S. Alexiadis, P. Kelly, P. Daras, N. E. O'Connor, T. Boubekeur, and M. B. Moussa, "Evaluating a dancer's performance using Kinect-based skeleton tracking," in *Proceedings of the ACM International Conference on Multimedia*, Scottsdale, Arizona, USA, November 2011, pp. 659–662.
5. J. Tautges, A. Zinke, B. Krüger, J. Baumann, A. Weber, T. Helten, M. Müller, H.-P. Seidel, and B. Eberhardt, "Motion reconstruction using sparse accelerometer data," *ACM Trans. on Graphics*, vol. 30, pp. 18:1–18:12, 2011.
6. D. S. Alexiadis, D. Zarpalas, and P. Daras, "Real-time, full 3-D reconstruction of moving foreground objects from multiple consumer depth cameras," *IEEE Transactions on Multimedia*, vol. 15, pp. 339–358, 2013.
7. S. Essid, X. Lin, M. Gowing, G. Kordelas, A. Aksay, P. Kelly, T. Fillon, Q. Zhang, A. Dielmann, V. Kitanovski, R. Tournemenne, A. Masurelle, E. Izquierdo, N. E. O'Connor, P. Daras, and G. Richard, "A multi-modal dance corpus for research into interaction between humans in virtual environments," *Journal on Multimodal User Interfaces*, vol. 7, pp. 157–170, 2013.
8. M. Gowing, P. Kelly, N. E. O'Connor, C. Concolato, S. Essid, J. L. Feuvre, R. Tournemenne, E. Izquierdo, V. Kitanovski, X. Lin, and Q. Zhang, "Enhanced visualisation of dance performance from automatically synchronised multimodal recordings," in *Proceedings of the ACM International Conference on Multimedia*, Scottsdale, Arizona, USA, November 2011, pp. 667–670.
9. V. Vlahakis, J. Karigiannis, M. Tsotros, M. Gounaris, L. Almeida, D. Stricker, T. Gleue, I. T. Christou, R. Carlucci, and N. Ioannidis, "Archeoguide: first results of an augmented reality, mobile computing system in cultural heritage sites," in *Proceedings of the 2001 Conference on Virtual Reality, Archeology, and Cultural Heritage*, Glyfada, Greece, November 2001, pp. 131–140.
10. D. S. Monaghan, J. O'Sullivan, B. Kelly, N. E. O'Connor, O. Kazmierczak, and L. Comer, "Low-cost creation of a 3D interactive museum exhibition (Best Technical Demo)," in *Proceedings of the ACM International Conference on Multimedia*, Scottsdale, Arizona, USA, November 2011, pp. 823–824.

Personalized Cultural Heritage Experience outside the Museum

Connecting the outside world to the museum experience

Alan J.Wecker[1,2], Tsvika Kuflik[1] and Oliviero Stock[3]
[1]University of Haifa, [2]University of Trento, [3]FBK-irst
ajwecker@gmail.com, kuflik@is.haifa.ac.il, ostock@fbk.eu

Abstract. We propose a new area to focus upon as a research challenge: How to use personalized technology to connect our cultural heritage experiences to our "daily" lives. In particular we want to connect the visitor's museum experiences with outside relevant mobile cultural heritage experiences. We begin to examine what processes and technologies will be necessary to accomplish the when, what and how of this challenge.

Keywords: lifelong user models, museum applications

1 Introduction

Museums, as a cultural heritage site, have long been a primary showground for the exploration of new technologies. This combination of museum and technologies has led to conferences such as Museum and Web and workshops such as this one (PATCH). A plethora of museum guides have been developed and explored. Recent new directions for research in this field, have concerned themselves with 1) expanding the on-site visit with prior and post experiences, normally at a desktop computer at home;1 2) expanding the visit from a onetime experience to an experience that may repeat itself multiple times over a lifetime, including the reuse of the information elicited from experience gained onsite (e.g. a user model) at multiple sites [1]. Personalization is deemed a key factor in the success of such objectives. Our proposed third direction for research, is **examining how to enhance other experiences outside the museum site based on experiences at that site**. By doing this we can connect our cultural heritage experiences to our "daily" lives.

To illustrate such an experience, consider the following scenario: A user, who have shown an interest in bible studies is traveling near Emek HaElah in Israel, the user gets an SMS or short verbal message, that the user is nearing the site of the biblical battle between Goliath and David. While driving, the user can be presented with either classical music or modern Israeli music connected with this theme. When she gets out of her car she can be presented with information (either short videos or pictures) presenting some information on the history of the area and of weaponry from the time of King David which she had seen in a museum exhibition some time ago. Alternatively consider another scenario: A user is reading a text (or alternatively

[1] We are in the process of submitting a paper on this topic for publication. Interested persons can contact the authors for a pre-publication draft.

browsing the web), which contains references to Rome, on an electronic mobile reader (such as Kindle) while auxiliary information, i.e. pictures from past visits to the Coliseum, are presented on a large display (TV, wall display), not only pictures taken by the user, but pictures of historic significance personally related, such as the visitor's parents standing by the Arch of Titus can be shown. In this case the specific text the user is reading is augmented by relevant personal information and additional relevant information that may be available on-line filtered by the user's model.

From our example scenario emerge three processes which we feel are worthy of further examination: 1) **identification** of ***contextual opportunities*** outside the museum site, which are suitable to present cultural heritage information (primarily from museums, that the user has visited or may visit in the future), 2) **selection** of material which may interest the user in this particular contextual opportunity 3) **delivery** of such material given the context (devices, displays) and user preferences (user profile).

2 Related Work

As noted, museums have been the site of personalization research, mainly focusing on new technologies. There has been considerable research done in this area (PEACH, PIL, CHIP)[2-4] however the industry has not yet widely adapted personalization. Primarily in the past, personalization (via user models) has been used for cultural heritage frameworks for single visits. Challenges and a roadmap are discussed in [1].

The field of Pervasive Computing, due to the huge amounts of data involved has developed a number of different methodologies to work with context.[6-10]. Dealing with context can be viewed as a method to obtain personalized services, for example using contextual services based on location. Some thought has also been given to personalization of pervasive applications [10,11] .

User Models have been used as a vehicle for personalization and adaptation. In the field of user modeling a great deal of thought has been placed recently on the reuse of models for a number of applications [12,13]. GUMO, UserML and UbisWorld [14] are examples of efforts to achieve some sort of agreed upon interchange format in this area. Another effort to encourage reuse of user models is the Personis user modeling server that suggest the reuse of data stored centrally by a variety of different applications by applying specific resolvers [15,16]. All of the above point to the issue of user models which evolve over time (lifelong user models), this issue has been examined extensively for the area of lifelong learning models [17]. The importance of context to user models and how to represent and reason with such models is an emerging issue. The milieu of what group the user is visiting with can affect his preferences and hence the model.

In order to accomplish our goals, it is clear, that we will need to make use of nascent standards and methodologies such as: TourML[18] which is a standard to connect digital assets to cultural heritage items, UserML and Personis (discussed above).

3 Proposal

As mentioned above in trying to think about how to connect the museum to the "outside" world we are concerned with three processes 1) opportunity identification 2) content selection 3) contextual delivery. Each of the processes described below heavily depend on technologies of personalization and context, and their interplay as delineated in the Related Works section

- **Opportunity Identification**: In considering when to bring cultural heritage information, from for example a prior museum visit, to enhance experiences outside the cultural heritage site, one should exploit opportunities as they arise. Opportunities can be triggered by explicit requests for information while on the "go" outside the cultural heritage site, or implicitly, when certain conditions are met. Examples of implicit opportunities include: being in certain locations connected with cultural heritage, meeting people associated with cultural heritage, certain key historical dates and browsing of content that can be associated with cultural heritage.
- **Content Selection**: In today's digital world, during the museum visit one may have access to a large number of digital assets. The assets can be made available by the museum, through the use of guides (both stationary and mobile). Not only does the visitor encounter digital assets previously prepared by others, the visitor may be directly responsible for creating new assets (for example: pictures, video, recording, written impressions). These assets can also take the shape and form of music, poetry and can be stored in a variety of multimedia formats ranging from video to audio to text. This content may not necessarily be of an informative nature, but can be of a social nature, which may evoke emotions to bind the visitor to his cultural heritage experience. In addition one may need to filter and order these vast quantities of information in order that user not be overwhelmed and get material that is suitable to their individual preferences (personalization).
- **Contextual Delivery:** These assets can also be utilized to enhance a number of other user experiences, not only during a visit to a cultural heritage site (which should probably be primarily dedicated to experiencing the site itself). However a user, while mobile, will not necessarily have access to devices that can display or play these assets. Hence, the system must make choices using the current context and user model to decide which and how to present these assets given the users mobile context. Alternatively, when reading a text with references to cultural heritage sites, information from a person's past experience (both the user model and the digital assets) which is related can be brought to bear on his present experiences using a variety of secondary displays according to the user's model (preferences).

4 Summary

Given the above, the major research question we wish to propose exploration of is:

How can we apply technology, enhanced with accumulated experiences at cultural heritage sites (e.g., museums), to opportunities for delivering contextual information in settings external to such sites?

This question can be broken down to three phases (opportunity identification, content search and recommendation, contextual delivery) each one with a more specific question: 1) How can cultural heritage opportunities be identified? 2) When an opportunity presents itself; how can relevant content (informative, emotional, social) be identified/found? 3) How can this content (personalized) be presented within the context of the specific user?

5 References

1. L Ardissono, T Kuflik, D Petrelli. Personalization in cultural heritage: the road travelled and the one ahead, User Modeling and User-Adapted Interaction. 22 (2011) 1-27.
2. T Kuflik, O Stock, M Zancanaro, A Gorfinkel, S Jbara, S Kats, et al. Visitor’s Guide in an “Active Museum”: Presentations, Communications, and Reflection. Journal of Computers and Cultural Heritage. (2011).
3. O Stock, M Zancanaro, PEACH: Intelligent Interfaces for Museum Visits, Springer-Verlag New York Inc 2007.
4. L Aroyo, N Stash, Y Wang, P Gorgels, L Rutledge. Chip demonstrator: Semantics-driven recommendations and museum tour generation, The Semantic Web. (2008) 879-886.
5. O Stock, M Zancanaro, P Busetta, C Callaway, A Krüger, M Kruppa, et al. Adaptive, intelligent presentation of information for the museum visitor in PEACH, User Modeling and User-Adapted Interaction. 17 (2007) 257-304.
6. K Sheikh, M Wegdam, M van Sinderen, Middleware support for quality of context in pervasive context-aware systems, (2007) 461-466.
7. J Soldatos, N Dimakis, K Stamatis, L Polymenakos. A breadboard architecture for pervasive context-aware services in smart spaces: middleware components and prototype applications, Personal and Ubiquitous Computing. 11 (2007) 193-212.
8. JE Bardram. The Java Context Awareness Framework (JCAF)–a service infrastructure and programming framework for context-aware applications, Pervasive Computing. (2005) 98-115.
9. N Davies, J Landay, S Hudson, A Schmidt. Guest Editors' Introduction: Rapid Prototyping for Ubiquitous Computing, IEEE Pervasive Computing. (2005) 15-17.
10. K Henricksen, J Indulska, A Rakotonirainy. Modeling context information in pervasive computing systems, Pervasive Computing. (2002) 79-117.
11. A Held, S Buchholz, A Schill. Modeling of context information for pervasive computing applications, Procceding of the World Multiconference on Systemics, Cybernetics and Informatics. (2002).
12. T Kuflik, J Kay, B Kummerfeld. Lifelong Personalized Museum Experiences, Proc.Pervasive User Modeling and Personalization (PUMP'10). (2010).
13. A Kobsa. Generic user modeling systems, The adaptive web. (2007) 136-154.
14. D Heckmann, T Schwartz, B Brandherm, M Schmitz, M von Wilamowitz-Moellendorff. Gumo–the general user model ontology, User modeling 2005. (2005) 149-149.
15. S Gerber, M Fry, J Kay, B Kummerfeld, G Pink, R Wasinger. PersonisJ: mobile, client-side user modelling, User Modeling, Adaptation, and Personalization. (2010) 111-122.
16. M Assad, D Carmichael, J Kay, B Kummerfeld. PersonisAD: Distributed, active, scrutable model framework for context-aware services, Pervasive Computing. (2007) 55-72.
17. J Kay. Lifelong learner modeling for lifelong personalized pervasive learning, Learning Technologies, IEEE Transactions on. 1 (2008) 215-228.
18. R Watson, S Akselsen, E Monod, L Pitt. The Open Tourism Consortium::: Laying The Foundations for the Future of Tourism, European Management Journal. 22 (2004) 315-326.

Curators in the Loop: a Quality Control Process for Personalization for Tangible Interaction in Cultural Heritage

Elena Not[1], Daniela Petrelli[2]

[1]Fondazione Bruno Kessler, Trento, Italy
{not@fbk.eu}
[2]Sheffield Hallam University, Sheffield, United Kingdom
{D.Petrelli@shu.ac.uk}

Abstract. Personalization research for the Cultural Heritage domain has mostly focused on supporting the visitors by automatically creating the narratives and the visit experience. The meSch project shifts the focus of research from visitors to curators and re-designs the personalization process in a context of tangible interaction, to strengthen the role of human-supervised steps in the definition of the adaptive structures for both the content and the interaction. Targeting the work of curators opens the way toward personalization to become deployable in real and complex settings: professionals in control of the platform will use personalization in any context they consider worth it, to adapt both the content and the interaction behavior of smart objects delivering it, and will monitor the quality of adaptive experiences from the design to the actual onsite delivery and subsequent follow-up online exploration. This paper presents the articulated personalization research agenda of the meSch project.

Keywords: personalization architecture, human-supervision, tangible interaction

1 A Story of Research Successes and of Scarce Exploitation

Cultural heritage has been a privileged application domain for personalization technologies for many years, since visitors can highly benefit, during cultural heritage exploration, from individual support that takes into account contextual and personal attributes and visitors' evolving behavior during the visit, given that visitors are highly heterogeneous and require different types of information, at different levels of detail [9][26].

Research efforts have successfully demonstrated how personalization technology can affect various aspects of the interaction with visitors: the information selected for presentation; the organization of the overall presentation; the media used to interact with users; the individual vs. group type of adaptation; the interaction modalities [2]. Starting from the first web-based adaptive hypermedia systems [24][16] in the '90s, personalization researchers have progressively taken advantage of novel technology like portable devices, wireless connectivity, localization methods, and powerful platforms to experiment adaptive approaches also within technology enhanced mobile scenarios [20][25] or virtual reality applications [6][3]. Experiences of combining personalization opportunities with tangible and embodied interaction in the museum setting are instead still limited [12]. Furthermore, the convergence of internet and wireless technology has made the exploration of (digital and material) cultural heritage a continuous process, starting before the visit and ideally never ending, as the user is able to plan the visit online, visit the site, and then "revisit" places of interest online again [14].

However, despite of over 20 years of research in personalization for cultural heritage [2] only a few examples like the CHIP [11], PEACH [25], and PIL [15] projects reached a large scale evaluation with museum visitors. Conversely, personalization services in areas like e-commerce applications, news alerts, support to digital content access, tourist guides [1][5] have left the lab and are being used commercially. The reasons for personalization in cultural heritage to lag behind are many [2], including difficulties in modeling groups, and the challenge of evaluation. Practical considerations, such as cumbersome technological infrastructures (e.g. sensors to be placed indoors for position detection), or an unbalanced cost-benefit [4] have a significant impact. One of the challenges identified by Ardissono et al. [2] is the lack of standardization and the often idiosyncratic data representation that on one hand allows sophisticated personalization but on the other limits the reuse of existing digital resources such as heritage databases. In essence, what hinders the actual adoption of personalization in cultural heritage is the complexity that makes it prohibitive for institutions in terms of time (to prepare the data in the right format) and technical expertise, and the high cost against benefit for just a small part of the audience.

1.1 Curators and Exhibition Designers in the Loop

One additional aspect that it is usually underestimated in analyzing the weaknesses of personalization exploitation is the scarce role that is often reserved for curators and authors. Indeed, curators have long been considered as a precious source of expert information on how to compose effective labels, audio commentaries and texts to derive rules and algorithms to be implemented in personalization systems, in the utopic view that the system, when appropriately instructed by programmers and equipped with properly annotated data sources, can then autonomously decide how to select and dynamically compose the bits of content to deliver all the required forms of personalization in whichever situation. This approach inevitably suffers from problems like lack of portability, reusability, robustness, difficulties in checking the validity of system outputs,… all problems that hinder the deployment in real settings.

In this paper we discuss how the process of defining adaptive structures can be improved, made more flexible and portable, by identifying and separating the strength of automatic mechanisms from that of the human judgment and effectively synchronizing the activities of the two actors to achieve a superior quality of the results, a technique effectively used in other cases of human-intelligent systems collaboration [7][19] and of end-user development [10]. We present here the research agenda of the meSch project, where we are pushing the boundaries of investigation even further, by uncovering the potential of combining (i) the personalized presentation of digital content with (ii) tangible interaction with technology augmented exhibits or spaces and (iii) social interaction. meSch adopts a co-design approach where curators, designers, computer scientists and engineers all work together to explore the possibilities offered by personalization technology. We are currently in the process of concepts generation and sketching-in-hardware what a personalized visit can be and in analyzing the technical requirements that the needed forms of personalization impose on the system architecture. Later in the project we will shift our attention to the authoring stage and the authoring tools that will be again co-designed with heritage professionals.

2 Personalization in a Scenario of Tangible Interaction

In meSch we envisage a cultural space filled with smart objects, each with their own (adaptive) stories embedded therein, that will be revealed if and when conditions are right, e.g. visitors have reached the right time in the storyline, or a group of them is acting in a certain way, or another smart object is close by. The role of curators and exhibition artists is fundamental in conceiving the very first idea of the intended educational/experiential mission of the exhibition, in selecting the most meaningful and evocative objects, in imagining alternative threads of narration, in deciding when and how digital information should be unlocked. Fig. 1 illustrates the overall production cycle for personalized experiences.

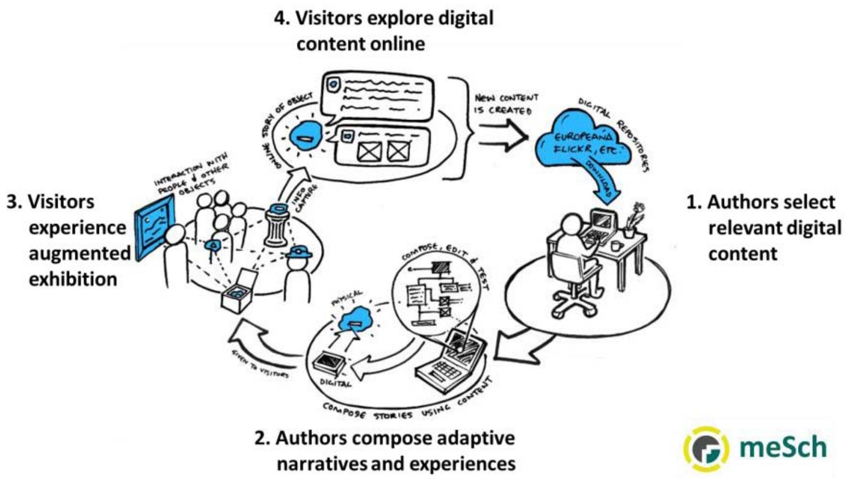

Fig. 1. The meSch production cycle for personalized experiences in a cultural space.

1) Curators conceive a new exhibition, or an adaptation to an existing one, and select pivotal objects (possibly replicas) that have been augmented with digital capabilities through the embedding of miniaturized multisensory integration platforms (e.g. Arduino, .net Gadgeteer, Raspberry Pi) or of more traditional devices (e.g., an mp3

player or a smartphone of which only the screen is made visible)[23][18]. They edit or select from proprietary or public domain multimedia databases (like Europeana [8]) suitable chunks of digital content related to the objects.

2) With the help of an authoring tool they assemble the chunks of content into threads of narration and select adaptive rules for their presentation. Similarly, they associate enabling (inter)actions over the objects to release the contents. The resulting structures for adaptive experiences are then downloaded onto the smart objects that will compose the exhibition.

3) Visitors approaching the smart objects at the exhibition site, either individually or in groups, will be able to experience the physical dimension of exhibits as well as the (social) engagement that may be favored by the tangible interaction coupled with the contextual delivery of mindful content such as a coherent story or an appropriate soundscape.

4) The logs of what visitors have experienced onsite, will translate into a digital souvenir that become the basis for a further personalized exploration online, after the visit.

2.1 A Quality Control Process for Personalization

meSch extends the typical architectures for personalization in the cultural heritage domain [2], by designing a multistage personalization process that decouples the inner personalization algorithms from content management issues, thus making it easier to plugin new digital resources and new domains, and implements a range of services able to respond differently according to the specific personalization stage, i.e. to be used within the smart objects vs. to be used online, offered to the authors or to the visitors. The personalization process has been designed to include human-supervised content structuring and experience design required to get high quality results, an essential requirement for the enjoyment of cultural heritage. Fig. 2 depicts, at a high level, the multilayer personalization process in meSch.

To address the issue of domain portability in a principled manner, a specific data access component will devise mechanisms for the effective integration of multiple digital repositories both public such as Europeana, or local, as museum archives, or Web2.0 such as DBpedia for the use of recommender functionalities. A uniform API will be made available to personalization services to grant data access, thus hiding the complexity of query and data mapping. This will facilitate the plug-in of new digital resources, also with different levels of content annotation richness. meSch puts a lot of emphasis on reusing existing digital resources: although this can limit the performance of the personalization mechanisms as the existing metadata may not hold useful details, we consider the exploitation of existing data a challenge worth tackling [2]. Indeed if the final interaction will be considered good enough this approach will open up many possibilities for the adoption of personalization by cultural heritage as there is no added costs for data preparation.

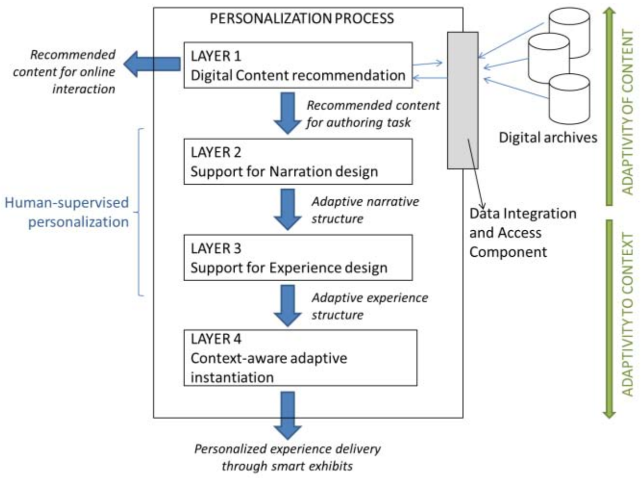

Fig. 2. Multilayered personalization process with human supervision

Four main stages of personalization compose the architecture:

1. the digital content recommendation service provides a flexible interface to different digital content archives (Layer 1);
2. the second stage supports curators in the effective use of the recommended content, its principled selection and the composition of alternative narrative threads according to educational and communicative goals (Layer 2);
3. in the third stage, the system supports the exhibition designer in describing how the narrative threads can be activated through tangible interactions with the objects and social behavior (Layer 3);
4. at the low-level, context-aware adaptive instantiation algorithms model decisions to be taken autonomously by the augmented objects at visit time according to the specific interaction context (Layer 4).

As shown in Fig. 2, Layer 1 and Layer 4 are performed by the system in an autonomous way, Layer 2 and Layer 3 are instead the two points where the human expertise comes into play. The figure also highlights a clear separation between the retrieval and composition of the content (layers 1 and 2) and the definition of the context (layers 3 and 4), as we distinguish two types of personalization: *personalization of content* pertains what (information or other) will be provided to visitors; *personalization in context* pertains the way in which the content is delivered. This split maps two distinct set of features, e.g. different content for different visitors groups or for different topics; different context for individual or social settings. It also maps two different types of cultural heritage professionals: *the curator*, often a scholar of the subject; and *the exhibition designer*, generally a creative mind, often an architect or a graphic designer. Curators will be able to specify which digital contents couple which real artifacts for which visitor group, while exhibition designers will decide the most appropriate effects to stimulate interaction among visitors and with the artifacts to unlock the digital content. Both professionals should be supported in their tasks of preparing the personalization structures that will deliver to visitors personalized content in context. This allows for multiplicity: the same structured content can be loaded on different smart exhibits so the content will be activated by different interactions (e.g. energetic children will have to work harder on the objects to release the content than older visitors); and the same smart exhibit can hold different structured content so that different content will be offered for the same interaction. A clear advantage of this approach is sustainability as the initial investment for the smart exhibits covers a number of different exhibitions each one with a different content. Moreover the same exhibition could travel to different institutions that will change the content to suite their visitors, e.g. translating it in the local language.

This approach poses the challenge of defining, implementing and evaluating models and tools for a human and a system that collaborate in building personalized tangible experiences for cultural sites and of determining the optimal division of responsibilities and roles between, for example, a curator authoring multiple narrative threads for schools and a recommender system that is actively finding new digital material for him to choose from.

3 Multipurpose Personalization Services

The meSch personalization architecture is currently under specification. An articulated research agenda has been put forward for each of the different personalization layers/tasks, as discussed in the following sections.

3.1 Content Recommendation Services

By definition, recommendation services are functions that exploit information about (i) users' personal characteristics, preferences and interaction history, (ii) semantic features of item descriptions or information articles and/or evaluation ratings, and (iii) possibly the behavior of a community of other users, to suggest items, pieces of information or services within a large information space, that best suit the user's needs in a given situation and context [21]. Recommenders have reached good maturity and robustness to be deployed in real setting applications in diverse domains and for diverse (context-dependent) user tasks [1][5]. For this reason, they represent a viable solution to introduce a first personalization layer over content extracted from knowledge sources to tune search results to the interaction context and task.

In meSch, recommendation techniques will work on top of meta-search results to additionally estimate the relevance of material (measured with various metrics, like similarity, divergence, novelty, serendipity), whenever this is made possible by the quality of the annotations in collection data. Several recommendation strategies will be evaluated with users/authors to find the most effective balance with respect to the quality of user experience. This layer serves two different requests for personalized content: (a) curators' content inspection through the authoring toolkit used for composing narrations and (b) visitors' information requests from the online interface.

Recommendation for the authoring task.

The author is responsible for the ultimate choice of the most appropriate content to be assembled within the ecology of smart objects, however he might not be familiar with all the contents available in external databases. We envisage a recommender system to run in the background and to support the author who is searching, selecting and composing content into narratives, by appropriately reordering or highlighting search results. We see this partnership between author and recommender system as an opportunity for the author to discover new content that can enrich his current effort, as well as a way to become aware of what is available in the repositories. Content-based recommendation strategies [17] can be used to reorder the results according to their similarity/novelty with respect to the other contents already selected for presentation (e.g. if the author is composing a narrative thread conceived for primary school children on soldiers' life in the trenches during WWI and some photographs have been selected by the curator from the local museum database, the system may suggest similar material extracted from Europeana [8], as in Fig. 3).

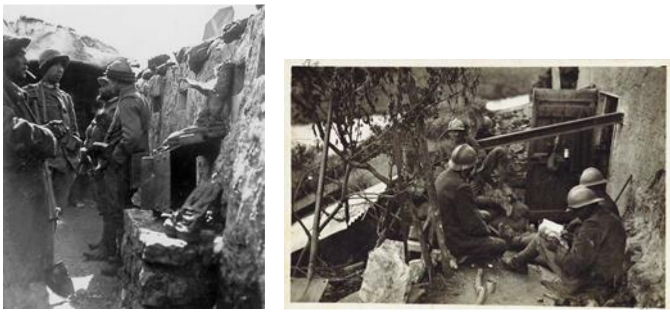

Fig. 3. Life in the trenches. On the left, picture from the local archives of the Museo Storico Italiano della Guerra. On the right, picture retrieved from Europeana (provider Museo Centrale del Risorgimento)

A recommender would also be able to automatically manage query relaxations [22] to suggest alternative content that may fit the narrative context as well (e.g., same searched content but in different media – image of a fort vs. 3D reconstruction; or different type of memento – letter vs. photography; or different period – photos of the same artifact in different periods).

When new smart objects have to be added to an existing exhibition, the recommendation services can take advantage, through collaborative-based recommendation strategies [13], of the actual behavior of users during onsite and online interaction to better tune the suggestion of additional information to be included in the new objects. For example, when adding to the exhibition a new smart object with audio/visual output capacity (e.g. a book with an embedded tablet), the system can suggest to the author texts from soldiers' letters or diaries, in case this content is similar to the type of information items that were more preferred by users, in previous onsite and online explorations.

Recommendation for online interaction.

In the meSch vision, individual logs of visitors' behavior during the visit to the cultural site will be used to dynamically compose data souvenirs and to bootstrap the online personalization of digital content without the need for the user to input explicit settings (thus solving the so called cold start problem). The data souvenir of the completed visit offers the starting point for online exploration, it gives hooks for the recommender to find additional related material. Users can save into their data souvenir the additional info objects explored online: this provides to the systems an explicit feedback about specific user interest and can be used as "rating" to enforce quality to collaborative reasoning and subsequent content suggestions. The recommender system will select (i.e., filter and re-order) the links proposed for further exploration by combining: (a) Content-based strategies that identify the information items, media types, subjects that most attracted the user interest during his visit (e.g. several interactions with biographic texts vs. historical descriptions; several pictures vs. audio files) and assign appropriate relevancy weights to additional information objects for exploration; (b) Collaborative-based strategies that exploit the online exploration behavior of users with similar onsite experience to better predict items of interest for the current user.

3.2 Adaptive Content Authoring

Content personalization for cultural heritage, for onsite use in particular, is a complex process that needs to take into account several components, e.g. the content, the description of the physical space, and the conditions (or rules) under which a certain content is delivered according to the contextual features. This unquestionable complexity calls for a collaboration between the author and the system to constrain the author to well-formed rules and schemas while at the same time gives space to creativity to generate engaging narrations. One example of a graphic tool purposefully designed to support non-technical users in composing the data structure needed for personalization was developed and used in the HyperAudio project [20]. Through drag-and-drop interaction, users were able to compose activation networks and run them in simulation mode to test if the personalization conformed to expectations. Templates were offered to speed up the editing and users only needed to load the desired content in the specific node [20]. The outcome of this activity was a set of data properly annotated into a micro-network and ready to be used in the onsite trial. In meSch we plan to exploit a similar approach. Through co-design activities with stakeholders we will first unpack the process curators and artists go through when creating a new exhibition or new educational material supporting visits (e.g., leaflets for school group activities when in the museum and back in class). This study will feed the identification of: best practices for content selection and composition, forms of personalization and presentation strategies used by museum staff, typical patterns of narrative. This will allow defining and testing various classes of adaptive composition rules and of pre-packaged schemas, like for example skeletons for composing narratives based on a temporal sequence (e.g., the life stages of a historical character), or reflecting a certain topic organization (e.g., comparison of different making techniques), or alternative object interpretations (e.g., historical vs. artistic description). The outcome will be a set of validated, ready-to-fill templates for adaptive narrative threads to facilitate authors' composition work, to be made visually available in the authoring tool. Mechanisms to create templates from scratch will be however offered to those users more interested in technology and inclined to experiment.

3.3 Adaptive Interaction Authoring

As meSch will bring personalization into smart objects augmented with interaction abilities, how the objects are going to deliver the adapted content according to the context of use needs consideration. The personalization of the interaction is complementary to the personalization of the content, as the (adaptive) narration can be delivered through alternative objects, by different activating actions and social interactions. For example, in a treasure hunt setting a smart object might reveal its contents just when all the members of the same competing group are close to it, whereas in an individual visit setting the content may be unlocked by simply picking the object up, and if the object is manipulated for a little longer additional information is presented. This specific stage of personalization authoring requires to define a vocabulary of actions and interactions that can be implemented by the hardware platform. In this way the editing is based on the type of interactions that can be actually built.

Through co-design activities with curators and technicians we will define and test various classes of interaction rules and pre-packaged schemas, like for example skeletons for fostering social situations (e.g., with extensive use of collaborative multi-user actions such as people marching in line), or object manipulation (e.g., with content disclosed by varied or prolonged manipulations, such as wind up the radio to play war bulletins), or objects search (e.g. in a sort of treasure hunt or to find your enemy to unlock the full story of the battle). The outcome will be a set of validated, ready-to-fill templates for interaction to be matched onto the narrative treads composed in layer 2 of the personalization process (Fig. 2). By decoupling the layer of content from the layer of interaction, it will be possible for an author to easily use multiple alternative narrative templates with the same interaction template and vice versa, the same narrative template can be applied to multiple interaction templates.

3.4 Onsite Adaptive Experience Delivery

The output of layers 1-3 of the personalization process depicted in Fig. 2 is the human-supervised creation of adaptive structures of experience (i.e., content + interaction) to be downloaded into the smart exhibits. The last layer of the personalization process instantiates these structures according to the actual behavior of visitors and the context they are in, i.e. it performs the automatic adaptation of the experience to a specific visitor, or a group of visitors within their specific context, at the very moment when the interaction with the smart exhibits takes place.

Layer 4 of the personalization process is based on algorithms for low-level adaptivity decisions to be taken autonomously by the augmented objects according to the specific interaction context. These algorithms are required to relieve authors from specifying the exact object behavior in all possible contextual situations. This includes resolving conflicts when alternative object behaviors are possible. For example, in case there are four people potentially affected by a presentation, the final decision of alternative audio is based on content mediation with respect to the narrative threads currently followed by the four people. Similarly, objects and locations that support different single-user and multi-user actions (like picking up, holding, collecting, sharing…) will need low-level, instant decisions on how to prioritize interaction events.

4 Conclusion

The meSch project aims at advancing the state of the art in personalization for cultural heritage by integrating principles of context-awareness related to onsite, tangible and socially situated interaction with principles of content adaptation and by embedding these different personalization mechanisms into physical objects thus transforming them into smart exhibits. The project envisages a tool to empower curators to make the most of their (digital and physical) collections and to create compelling visitor experiences. Complexity reduction is achieved through a human-supervised multi-layer personalization architecture that splits the adaptivity of the content from the adaptivity of the interaction, with the system relieving the author from the most complex tasks and assisting him to achieve a high quality result. The final aim is that of taking personalization of cultural heritage to the large scale offered by the existing repositories of digital content and to test a sustainable architecture (i) able to serve different personalization tasks, (ii) portable to different content and (inter)actions vocabularies, and (iii) easily reusable in different physical sites.

Acknowledgements. The research described in this paper is part of the meSch project, Material Encounters with Digital Cultural Heritage. The interdisciplinary team includes: Sheffield Hallam University; University of Limerick; Waag Society; University of Strathclyde; eCTRL; DEN; University of Stuttgart; University Carlos III Madrid; Museo Storico Italiano della Guerra; University of Amsterdam - Allard Pearson Museum; MUSEON; FBK. More at http://www.mesch-project.eu

meSch (2013-2017) receives funding from the European Community's Seventh Framework Programme 'ICT for access to cultural resources' (ICT Call 9: FP7-ICT-2011-9) under the Grant Agreement 600851.

5 References

1. Aldrich, S. E.: Recommender Systems in Commercial Use. AI Magazine 32(3), 28-34 (2011)
2. Ardissono, L., Kuflik, T., Petrelli, D.: Personalization in Cultural Heritage: The Road Travelled and the One Ahead. UMUAI The Journal of Personalization Research, special 20th anniversary issue 22 (1-2), 73-99 (2012)

3. Bonis, B., Stamos, J., Vosinakis, S., Andreou, I., Panayiotopoulos, T.: A platform for virtual museums with personalized content. Multimed. Tools Appl. 42(2), 139–159 (2009)
4. Bowen, J., Filippini-Fantoni, S.: Personalization and the Web from a museum perspective. Museum and the Web (2004)
5. Burke, R., Ramezani, M.: Matching Recommendation Technologies and Domains. In Ricci, F.; Rokach, L.; Shapira, B.; Kantor, P.B. (Eds.), Recommender Systems Handbook. 1st Edition, Springer, pp. 367-386 (2011)
6. Calder, J., Melengoglou, A.C., Callaway, C., Not, E., Pianesi, F., Androutsopoulos, I., Spyropoulos, C.D., Xydas, G., Kouroupetroglou, G., Roussou M.: Multilingual personalized information objects. In: Stock, O., Zancanaro, M. (eds.) Multimodal Intelligent Information Presentation, pp. 177–201, Springer, Netherlands (2005)
7. Ciravegna, F., Dingli, A., Wilks, Y., Petrelli, D.: Using adaptive information extraction for effective human-centred document annotation. In: Franke, J., Nakhaeizadeh, G. & Renz, I. (eds.) Text Mining, Theoretical Aspects and Applications, Physica-Verlag, pp. 153-164 (2003)
8. Europeana. http://www.europeana.eu.
9. Falk, J.: Identity and the museum visit experience. Walnut Creek, CA, Left Coast Press (2009)
10. Ghiani, G., Paternò, F., Spano, L.: Cicero Designer: An Environment for End-User Development of Multi-Device Museum Guides. Proceedings of IS-EUD 2009, LNCS 5435, Springer, pp. 265-274 (2009)
11. Hage, W.R. van, Stash, N., Wang, Y., Aroyo, L.M.: Finding your way through the Rijksmuseum with an adaptive mobile museum guide, In: Proc. the 7th Extended Semantic Web Conference (ESWC), pp. 46-59, Springer-Verlag Berlin, Heidelberg (2010)
12. Hatala, M., Wakkary, R.: Ontology-based user modeling in an augmented audio reality system for museums. UMUAI - User Modeling and User-Adapted Interaction 15(3–4), 339–380 (2005)
13. Herlocker, J.L., Konstan, J.A., Borchers, A., Riedl, J.: An algorithmic framework for performing collaborative filtering, In: SIGIR '99 Proceedings of the 22nd annual international ACM SIGIR conference on Research and development in information retrieval, pp. 230-237, ACM New York, NY, USA (1999)
14. Kuflik, T., Kay J., Kummerfeld, B.: Lifelong personalized museum experiences. In: Berkovsky, S., Bohnert, F., Carmagnola, F., Cheng, D., Heckmann, D., Kuflik, T., Nurmi, P., Partridge, K. (eds.) Proceedings of Workshop on Pervasive User Modeling and Personalization (PUMP'10), pp. 9–16 (2010)
15. Lanir, J., Kuflik, T., Dim, E., Wecker, A.J., Stock, O.: The Influence of a Location-Aware Mobile Guide on Museum Visitors' Behavior. Interacting with Computers. First published online: February 24, 2013 (2013)
16. Oberlander, J., O'Donnell, M.,Mellish, C., Knott, A.: Conversation in the museum: experiments in dynamic hypermedia with the intelligent labeling explorer. New Rev. Multimed. Hypermed. 4, 11–32 (1998)
17. Pazzani, M. J., Billsus, D.: Content-based recommendation systems. In The Adaptive Web, pp. 325-341, Springer-Verlag Berlin, Heidelberg (2007)
18. Petrelli, D., Bowen, S., Light, A., Dulake, N.: Digital Christmas: An Exploration of Festive Technology. In: Proceedings of DIS 2012 - Designing Interactive Systems, pp. 348-357, ACM New York, NY, USA (2012)
19. Petrelli, D., Lanfranchi, V., Ciravegna, F.: Working Out a Common Task: Design and Evaluation of User-Intelligent System Collaboration. In: Costabile, M.F., Paternò, F. (eds.)

Proceedings of the 10th IFIP TC13 International Conference on Human-Computer Interaction INTERACT 2005, LNCS, vol. 3585, pp 309-322 (2005)
20. Petrelli, D., Not, E.: User-centred design of flexible hypermedia for a mobile guide: reflections on the HyperAudio experience. UMUAI - User Modeling and User-Adapted Interaction 15(3–4), 303–338 (2005)
21. Resnick, P., Varian, H. R.: Recommender systems. Commun. ACM 40(3), 56–58 (1997)
22. Ricci, F., Blaas, D., Mirzadeh, N., Venturini, A., Werthner, H.: Intelligent Query Management for Travel Products Selection. In: Wöber, K., Frew, A. J., Hitz, M. (eds.) Information and communication technologies in tourism 2002: Proceedings of the International Conference in Innsbruck, Austria, pp. 448-457 (2002)
23. Schmidt, A., Bial, D.: Phones and MP3 Players as the Core Component in Future Appliances. IEEE Pervasive Computing 10(2), 8–11 (2011)
24. Stock, O., Carenini, G., Cecconi, F., Franconi, E., Lavelli, A., Magnini, B., Pianesi, F., Ponzi, M., Samek-Lodovici, V., Strapparava, C.: Alfresco: enjoying the combination of NLP and hypermedia for information exploration. In: Maybury, M.T. (ed.) Intelligent Multimedia Interfaces, pp. 197–224. AAAI Press, Menlo Park (1993)
25. Stock, O., Zancanaro, M., Busetta, P., Callaway, C., Krueger, A., Kruppa, M., Kuflik, T., Not E., Rocchi, C.: Adaptive, intelligent presentation of information for the museum visitor in PEACH, UMUAI - User Modeling and User-Adapted Interaction 17(3) 257-304 (2007)
26. Wilkening, S., Chung, J.: Life stages of the museum visitor: Building engagement over a lifetime. Washington DC: American Association of Museums Press. (2009)

Wolfsoniana Smart Museum. A Pilot Plant Installation of the PALM-Cities Project

Andrea Caridi[1], Mauro Coccoli[2], Valentina Volpi[3]

[1] Darts Engineering,
Piazza della Vittoria, 9/3 - 16121 Genoa, Italy
[2] DIBRIS - Department of Informatics, Bioengineering, Robotics, and Systems Engineering, University of Genoa, Via Opera Pia, 13 - 16145 Genova, Italy
[3] CORIS - Department of Communication and Social Research, Sapienza University of Rome, Via Salaria, 113 - 00198 Roma, Italy
andrea.caridi@darts.it, mauro.coccoli@unige.it, valentina.volpi84@gmail.com

Abstract. This demo paper describes the pilot plant of the PALM-Cities Project, installed at a modern art museum in Genoa. The objective of the system is offering visitors a novel and immersive experience when visiting the exhibit. In fact, owing to automatic identification technologies, such as Near Field Communication (NFC) and the Quick Response Code (QR Code), people can interact with the surrounding environment by means of their own devices, e.g., smartphones or tablets, receiving personalized contents as they move in a given location. Different resources for the same artwork in the museum are delivered to different users on the basis of their preferences, such as, e.g., language, age and device.

Keywords: Smart device, smart museum, NFC, QR Code, mobile applications, personalized multimedia contents, user experience, user-centered design.

1 Introduction

The wide diffusion of low-cost smart devices with network connectivity capabilities is giving rise to a new generation of always-connected people. Besides, the evolution of automatic identification technologies based on the Radio Frequency Identification (RFID) is pushing the development of the Internet of Things (IoT) [1] and *smart museum* applications [2]. In this scenario, we developed a prototype, within the framework of the PALM-Cities (Personal Assistant for Mobile Liguria Citizens) project [3], whose objective is studying systems and networks for the delivery of personalized touristic and/or cultural contents on mobile devices in an easy way, in order to offer users a better experience in visiting a museum (see [3], [4] and references therein for a list of related works). To demonstrate the effectiveness of the proposed paradigm, we have chosen a local cultural setting, specifically the Wolfsoniana museum, for which information about the exhibited items has been tailored on the users' needs according to their language (i.e., English/Italian), level (i.e., kids/adults)[1] and device characteristics.

The remainder of the paper is structured as follows: in Section 2 the pilot installation is illustrated; in Section 3 the information architecture lying behind the system is depicted. Finally, Section 4 concludes the paper outlining future developments and research directions.

2 The PALM-Cities Pilot Plant

The pilot installation is publicly available to the visitors of the Wolfsoniana from March, until July, 2013. In this period, a restricted number of works of art has been selected, for which additional information will be delivered to the interested visitors' devices, by means of contactless interaction (Fig. 1). Specifically, both the Near Field Communication (NFC) and the Quick Response Code (QR Code) technologies have been adopted, which are readily available in a large number of mobile devices and allow implementing non-invasive and low-cost solutions, suitable for indoor and outdoor activities.

Fig. 1. The labels used in the Museum and the start-page of the Wolfsoniana App.

The system is composed by a network infrastructure, enabling communications within the exhibition, and by a mobile application available for both Android and iOS on the Google Play Store and the Apple App Store respectively. At the entrance, a smart poster illustrates users how to connect to the local free Wi-Fi and to install the App on their own devices, to start enjoying the Smart Museum experience.

2.1 The Network Infrastructure

The installed systems is composed by some passive technologies (Mifare Ultralight passive tags - ISO/IEC 14443 13,56 MHz and QR Code) that, by interacting with an

[1] The level parameter is the age of the user: in particular the kids's cluster includes the range 0 - 14 and the adults's one includes all the other users. The selection is based on different educational levels and cognitive dimensions of the target.

active mobile device and a mobile application, trigger the retrieving of contents from the museum server on the visitors' mobile devices via Wi-Fi. In more details, the RFID passive tags and the QR Codes contain a unique code that triggers a server to provide information related to a specific tagged piece and according to specific users' data, in particular to his/her age (the adaptation logic is based on the rule: age 0-14 or >14), language and device, accessible only when the mobile application is used inside the Local Area Network (LAN) of the museum, to cope with copyright issues.

2.2 The Mobile Application

The mobile application is made available for Android NFC-enabled smartphones and tablets, such as, e.g., Samsung Nexus S and Galaxy S III, as well as for iOS devices such as, e.g., Apple iPhone, iPod Touch and iPad (Fig. 2).

The provided features are the following: (i) *users' profiles*. At the first launch, the profile is set by inserting nickname, age and language; (ii) *RFID tag reading*. By approaching the NFC smartphone to the RFID tag on the plates, the visitor accesses to multimedia contents and can leave feedback; (iii) *QR Code scanning*. For devices without NFC, a camera is sufficient to scan the QR Code on the plates; (iv) *social activities*. The visitor can "like" the museum pieces; (v) *interactive map*. The visitors are given a map of the museum showing the tagged works; (vi) *help menu*; (vii) *fast switch*. Users can read a new RFID tag without returning to the NFC mode screen.

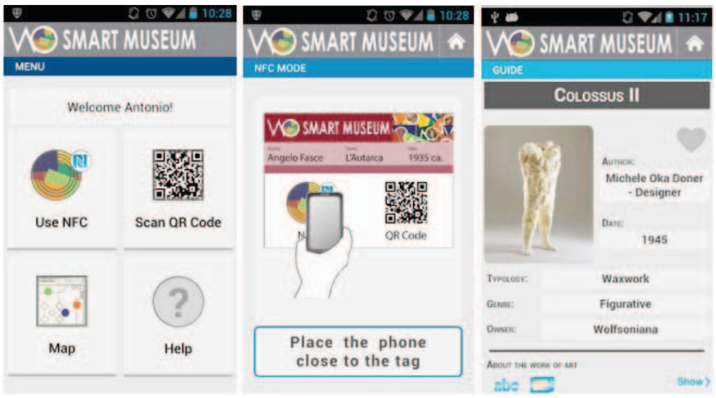

Fig. 2. Some screenshots of the Wolfsoniana App illustrating the NFC mode.

3 Information Architecture

To provide the system with the above listed capabilities, a suitable information architecture has been designed, reflecting users' profile and visualization issues. Summarizing, the system arranges the content to deliver according to *four* basic identifiers: (i) item; (ii) language; (iii) user level, and, finally; (iv) device type. Then, these four pieces of information are arranged in a tuple: *(item_id, lang_id, level,* *device)*, which is sent to the server through the wireless network via HTTP. Thus, the information structure is hierarchical and can be represented through a tree-shaped layout listing several resources for each item and, for each resource, a variety of implementations.

The flow of information happens as described in the above paragraphs describing the system and the App as well. Therefore, the device, by interacting with a given artwork, sends a parameterized URL to the underlying framework. We highlight that such a piece of information is composed by a *static* part embedded in the RFID tag, which is completed with a *variable* portion depending on users approaching the items, according to the schema of the tuple already introduced. This results in a query to be executed on the underlying data structure written in XML, which returns the selection of the best-suited contents for the user. Contents are transformed into a page through an XSLT transformation and they are further adapted to the specific capabilities of the device by means of Cascade Style Sheet (CSS) for performing the final rendering.

4 Conclusions

The described pilot plant offers the basic features of an adaptive system for the delivery of personalized contents. However, the whole project is based on a modular and flexible infrastructure, which may be the foundation of an integrated system of Tourism and Culture in the Liguria Region. In the future a holistic vision of Liguria territory, history, tradition and culture will be available to the visitors, who will be able to easily plan all the details about their visits to museums, theatres, historical sites, monuments, and tourist attractions in general.

Acknowledgments. This work is supported by the Regional Operational Program POR-CRO 2007-2013 of the Liguria Region, Italy. It is being developed by Darts Engineering, supported by DIBRIS, University of Genoa, and CATTID, Sapienza University of Rome, in collaboration with the Wolfsoniana Museum of Genoa and the Regional Foundation for Culture and the Performing Arts of the Liguria Region.

References

1. Gershenfeld, N., Krikorian, R., Cohen, D.: The Internet of things, Scientific American, vol. 291, no. 4, pp. 76-81 (2004)
2. Caviglione, L., Coccoli, M.: Design of a Software Framework for the Management and Personalization of Contents in Smart Museums, Int. J. of Software Engineering, vol. 5, no. 2, pp. 51-69 (2012)
3. Ceipidor, U.B., Medaglia, C.M., Volpi, V., Moroni, A., Sposato, S., Carboni, M., Caridi, A.: NFC technology applied to touristic-cultural field: a case study on an Italian museum, Proc. of the 5th Int. Workshop on Near Field Communication (NFC), pp. 1-6 (2013)
4. Roes, I., Stash, N., Wang, Y., Aroyo, L.: A personalized walk through the museum: the CHIP interactive tour guide. In CHI '09 Extended Abstracts on Human Factors in Computing Systems (2009)

A Proposal for an Open Local Movie Recommender

Federica Cena, Elisa Chiabrando, Andrea Crevola*, Martina Deplano, Cristina Gena, and Francesco Osborne

University of Torino, *3Juice - Italy {cena,chiabrando,deplano,gena,osborne}@di.unito.it,{andrea.crevola}@3juice.com

Abstract. This position paper describes the initial research assumptions and work carried on in the context of the ReAL CODE, a project aimed at realizing an intelligent and tailored distribution of digital creative contents towards end users. Contents belong to the movie domain and particular attention will be given to local contents, that will be promoted through recommendations in order to support local cultural heritage. The main research goals of the project are open data and interoperability of domain and user data. Our initial efforts on these key directions will be described.

Keywords: interoperability, open data, recommender, movies

1 Introduction

Movies belong to cultural heritage domain of a given place. On the one side the notion of cultural heritage links two concepts: heritage and culture. The term "heritage refers to local landscape, architecture and artefacts that make a particular place unique. The term cultural puts less emphasis on the place and, more generally, regards knowledge, and local traditions of a territory. We term cultural heritage resources both the physical heritage and the intangible activities related to it, like movies filmed and related to these places.
On the other end, it is not unusual to consider movies as belonging to cultural heritage and history of a given place. Increasingly, historians have moved away from a history that chronicles battles, treaties, and presidential elections to one that tries to provide an image of the way daily life unfolded for the mass of people. Film has an important role to play in these histories. [...] But when we focus on social and cultural history, especially the important role of leisure in the lives of ordinary people, film not only provides evidence and records but takes on a key role[1]. Movies filmed in a given place and in a given time are not only a priceless statement of the daily life and a picture of those places, but, when their authors and producers are somehow linked to that place, they also witness the cultural life and the cultural ferment of territory. The link to

[1] http://historymatters.gmu.edu/mse/film/socialhist.html

the evolving local movie resources produced in our region[2] is one of the key point of ReAL CODE[3] (Recommendation Agent for Local Contents in an Open Data Environment), a project aimed at designing an intelligent and tailored distribution of digital creative contents towards end users. Contents belong to the movie domain and particular attention will be given to local contents - namely movies produced and realized in the Piedmont region - which will be promoted through recommendations.

More in detail, one of the main goal of the project is to develop an open, interconnected and semantic database of local creative contents, in line with current trends of *Web 3.0* and *Open Data*. On the one hand the idea is to retrieve content information from different open datasets (Section 3) and automatically map it to the ReAL CODE own knowledge and format. On the other hand ReAL CODE own information will be available in a semantic and open format, so that everyone will be free to use, reuse and redistribute it. Another important goal of the project is to achieve user model interoperability (Section 3) in order to obtain a big amount of information about the user and her interest and to overcome the cold start problem in recommendations.

As for the user experience, the challenge is to realize an innovative and engaging application, that differs from traditional movie recommendation sites. In the ReAL CODE system, movies will be considered as an important part of people life, related to specific moments, places and people (Section 2). Based on user activity on the system and on imported user model, ReAL CODE will suggest the user movies that she could like, focusing on local contents. Movie operators could also join the system in order to insert and promote their own contents; in this way, the project will support the local cultural industry.

2 A Scenario

Similarly to music, movies are often related to people's special moments and they recall feelings and emotions. Independently from the feelings they trigger, movies are related to particular moments, to other people and to places.

Final User. *Marta enters ReAL CODE through Facebook login. Using information stored in the Facebook account, the system shows her some movies she liked or added to her profile. Among these movies there is "Roman Holiday", a movie with her favourite actress, Audrey Hepburn, that she used to see during her teenage years with her friends. She selects it and adds a few information about her sight experience: where she saw it (in her hometown, Rome), when (she can select a specific date) and with whom (her friend Sofia, who is also a*

[2] The movie industry plays a key role in the Piedmont Region, where we live. Piemonte Film Commission is a very active entity in producing movies and organizing movie-related events. The city of Torino hosts different film festivals (e.g. The International Turin Film Festival) and the National Cinema Museum, visited every years by thousand of national and international visitors.

[3] `http://www.ReALCODE.it/`. ReAL-CODE is a research project funded in the context of POR FESR 2007/2013 of the Piedmont Region, Italy.

Facebook's friend). She decides that she will add the other optional information (tag, vote, comments...) later. Immediately Sofia receives a Facebook notification about Marta's activity because ReAL CODE is in-and-out integrated with Facebook. Since she is pleased to remember this event, she decides to join ReAL CODE herself and she adds a comment on Marta's activity. After having added a minimum number of movies, the system can show Marta's "movie history" in multiple way: on a timeline, on a map or on a sort of diary. The system also shows her some recommendations about movies she could like. Beyond traditional suggestions (blockbusters movies, top rated movies, etc) the system recommends Marta local movies release.

Producer. *Francesco is a small movie producer that lives in a town near Turin. His company has a YouTube channel and a Facebook page. He decides to join ReAL CODE in order to increase its visibility, especially on niche and local audience. He imports into ReAL CODE information and videos from YouTube and Facebook and adds some other details about his company. Since next month he will participate with his last short movie to the Turin Film Festival, he accesses on ReAL CODE, adding specific information about his short movie, such as the movie categories, movie tags, etc. In this way, this movie could be recommended to potentially interested users.*

3 Data interoperability for User and Domain Models

User Model interoperability. The emergence of social networks on different domains (Facebook, Twitter, aNobii, LinkedIn) has made available an enormous amount of user data, such as demographic information, current location, friendship network, job and interests. Aggregating these data may be useful to solve the cold-start and sparsity problems, but it requires a deep knowledge of the user profiles in the various social networks and how combining data to obtain a complete and effective user model.

Some recent works [1][2] analyze and create user models from popular social networks. In Abel et al.[1] the authors capture user information of Twitter, Facebook, LinkedIn and social tagging activities in Flickr, Delicious and StumbleUpon to evaluate the performance of a strategy based on several cross-system user modeling. They improve recommendation quality in a significantly manner. Likewise, Shapira et al.[2] integrate Facebook data with the recommendation process and compare traditional different collaborative filtering methods with their cross-domain recommendation.

Differently from these previous works, we propose to combine content-based and collaborative filtering-based methods to suggest interesting movies according to a user model created from data extracted from different social networks. For our aims, the most interesting social networks are Youtube and Facebook.

Through the YouTube API[4], data we are considering to use are: *favorite videos* (videos flagged as favorite by the user), *watch history*(videos watched by

[4] Youtube API 2.0 `https://developers.google.com/youtube/`

the user), *uploads feed* (videos uploaded by the user), *video recommendations* (videos that may appeal to a user), *user subscription feed* (channels and people the user has subscribed to), *standard channel feeds*(channels that reflect user preferences). The two main problems we encountered with YouTube are: (1) most of the titles does not match exactly with official movies titles; (2) if the user's channel on Youtube has not been linked to her Google+ account, we can not find the corresponding user personal data.

On the other hand, personal data and information on users' interests can be collected through the Facebook Graph API and FQL interface[5]: user personal information (*i.e.* birthday, gender, geolocation); her friendship network; the lists of her favorite books, music and movies; the list of posts in her stream created by third application, such as Youtube. However, movies in Facebook are not categorized (*i.e.* comedy or horror), so it is necessary to retrieve this information from other external datasets, such as the ones described below. Moreover, recently, Facebook has introduced the notion of action on particular objects of the system (fitness, music, news, video and books), allowing users to tell sort of "stories" about what she did or what she would like to do. In particular, user may either tell that she watched a film and then evaluate it, or what she would like to see it in the future. Our idea is to import these data in the ReAL CODE, reason on them, and then use them both for the user timeline and for recommendations.

A recent innovation of Facebook allows users to **tell stories** on Facebook through a structured, strongly typed API[6] about what she did or what she would like to do. In particular, user may either tell that she watched a film and then evaluate it, or what she would like to see it in the future. Our idea is to import these data in the ReAL CODE reason on them, and then use them both for the user timeline and for recommendations.

Domain Interoperability. The Open Data initiative promotes the idea that the data should be freely available to everyone to use and re-publish. "Openness" of data enables the construction of a place in the web for global sharing, the "Web of Data". In this line, many information about the movie domain can be found in the Web and used in our project.

LinkedMDB[7] aims at being the first open semantic web database for movies, including a large number of interlinks to several datasets on the open data cloud. Data can be accessed using traditional Web browsers, Semantic Web browsers, SPARQL clients, but there are not APIs available.

TheMovieDB[8] provides a large movies database (118,000 titles) classified in 18 categories with other information (id, alternative title, cast, images, keywords, trailers, similar movies and so on) available through APIs.

[5] Facebook Graph API `http://developers.facebook.com/docs/reference/apis/`

[6] Recently, Facebook is introducing the notion of action on particular objects of the system (fitness, music, news, video and books) `http://developers.facebook.com/docs/reference/apis/`

[7] LinkedMDB `http://www.linkedmdb.org`

[8] TheMovieDB `http://www.themoviedb.org`

YouTube is an important data source, even if it contains only videos and clips, and not whole movies. However, information such as trailer and popular scenes can be very useful to enrich the domain model. Some of the available data from the API are: Videos related to a specified video; Title and author of the video (user ID, name, etc); Average rate and statistics (how many views, likes, bookmarks). There is not a semantic ontology describing categories of movies, but users can freely select some categories at the uploading time.

DBpedia[9] allows to extract structured data from Wikipedia and to link other data sets to it. 172,000 films instances are present, classified in a consistent ontology formed by a three-level taxonomy covering different general classes (genre, themes, year, location, nationality, etc) and 58 classes regarding genre of film, the most important category for us. These data are accessible via a SPARQL query endpoint. We chose DBpedia as a primary source for movie domain thanks to the advantages of its knowledge base: it covers many domains; it represents real community agreement; it automatically evolves as Wikipedia changes; and it is truly multilingual. Moreover, it is very informative in relation to movie domain, and this allow us to have a complete description of a film. For our purpose, it is particularly important the field "subjects", since it contains the ontological categories the movie belongs to, according to Wikipedia ontology, and we can use this information for user modeling purpose, making inference on user's interest in categories of films, starting from her interest in a specific movie. For the same reason, we will use also the other datasets when linked to DBpedia, since in this way it is possible to access Wikipedia categorization and integrate information with possibly missing ones.

Open Movie Recomendation Our idea is to create an open movie recommender, *i.e.* a recommender that uses open data from the Web as knowledge base and then combines content-based and collaborative filtering to suggest local interesting movies. Thus, after have retrieved the user personal information and the list of movies and friends in Facebook, ReAL CODE creates the user model of the specific user. Then, it finds meta information about movies (*i.e.* category, actors, register, and so on) from the external movies datasets seen above in order to create a domain model that can be used in the content-based recommendation. Moreover, ReAL CODE will use the information about user's social network in Facebook and in Youtube to suggest movies that are interesting for similar users belonging to the user's network.

4 Conclusion

Social Open Web (of Data) offers a lot of information we can get both about the user and the movie domain. Our next efforts will be concentrated on how to integrate, map, and reason on these data in order to have a knowledge base structured in such a way to provide new tailored user recommendation on local

[9] DBpedia `http://dbpedia.org/`

movies and to be open and available to other applications. Moreover, we are working in realizing an innovative user experience by means of game activities, that will be used also to infer new data about their preferences and thus enrich the user model.

Moreover, in order to improve the user model without just asking users to add explicit information and attract them to the ReAL CODE application, a set of Games With A Purpose [3] will be available. Through the gaming activity of the users and their mobile devices, the system will infer interesting data about their preferences and geolocalization. These data will be then used to provide customized recommendations. For example, one of the GWAPs provides a set of movies among which the player has to guess the correct ratings given by a random friend. Thanks to the confirmation by her friend, we can give points to the player as a reward, and implicitly update the user model of her friend. Another GWAP allows the player to add questions about movies and respond to ones posted by other users.

Thanks to this gaming activity, we can infer users' interests based on the selected movie with which they want to play. Furthermore, we can also take advantage of the human computation [4] approach for a community control of the correctiveness of the quizzes. Through the geolocalization of mobile devices, one typology of GWAP is connected to a treasure hunt approach: given a movie set in the current player city, she can collect points by checking-in in the places related to the plot of the movie. This game can be very useful also for the promotion of the tourism in a selected city, by incentivizing the users to visit specific places through gaming activities that can make the user experience more fun and appealing.

5 Acknowledgments

We would like to thank professor Giancarlo Ruffo for his helpful advices.

References

1. F. Abel, E. Herder, G.-J. Houben, N. Henze, and D. Krause. Cross-system user modeling and personalization on the social web. *User Modeling and User-Adapted Interaction*, 23, 2013.
2. B. Shapira, L. Rokach, and S. Freilikhman. Facebook single and cross domain data for recommendation systems. *User Modeling and User-Adapted Interaction*, 23, 2013.
3. L. Von Ahn. Games with a purpose. *Computer*, 39(6):92–94, 2006.
4. L. Von Ahn. Human computation. In *Design Automation Conference, 2009. DAC'09. 46th ACM/IEEE*, pages 418–419. IEEE, 2009.

Combining Composition Technologies and EUD to Enhance Visitors' Experience at Cultural Heritage Sites

Carmelo Ardito[1], Maria F. Costabile[1], Giuseppe Desolda[1], Rosa Lanzilotti[1], Maristella Matera[2]

[1] Dipartimento di Informatica, Università degli Studi di Bari Aldo Moro
Via Orabona, 4 – 70125 Bari, Italy
{carmelo.ardito, maria.costabile, giuseppe.desolda, rosa.lanzilotti}@uniba.it

[2] Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano
Piazza Leonardo da Vinci, 32 – 20134 Milano, Italy
matera@elet.polimi.it

Abstract. This paper illustrates our approach to enhance the visit experience of archeological parks. It exploits composition technologies, End-User Development and participatory design approaches, in order to allow different stakeholders to create, use and share Personal Information Spaces. Heterogeneous content can be combined and manipulated to satisfy different information needs, thus enabling personalized visits to Cultural Heritage sites.

Keywords: End-User Development, User-driven Service Composition, Personal Information Space.

1 Introduction and Motivation

Italy is one of the countries in the world with a rich patrimony of historical sites and cultural heritage. People are becoming increasingly aware of the importance of preserving and enhancing the value of cultural heritage, acknowledging these aspects as ways to help people to construct their cultural identities [1, 2].

Undoubtedly, information and communication technologies have a great potential "to enhance personal experience in cultural heritage sites". Therefore, for several years, we have been involved in projects related to Cultural Heritage (CH). The aim is to kindle people's interest in knowing more about history through guided analyses of the remains of ancient settlements, enhanced through the adoption of different types of tools. Our overall goal, as Human-Computer Interaction (HCI) researchers, is to contribute to fostering a wider appreciation of archaeology by offering tools able to inspire the general public and to increase awareness of the importance of CH. This is a highly complex challenge, which involves a wide variety of factors: who is going to use the tools, how the tools will be used, where the tools will be used. To cope with this challenge, experts from different disciplines have to share their knowledge, skills, practices and tools, in order to capitalize on the expertise and creativity that all such experts, and even end users, can bring in the design of the new tools.

In Italy, school pupils constitute a large proportion of the visitors to archaeological sites. They perform such visits with their teachers, as part of their school curricula. Thus, we have primarily worked to create educational games on different devices, which support school pupils learning about ancient history through their active involvement during visits to such sites and follow-up activities at school [3, 4, 5]. These games have been developed by setting up multidisciplinary teams that included experts in the CH domain, e.g., archaeologists, historians, directors and employees of archaeological parks, experts in technology and science, e.g., software developers and HCI experts, school teachers, as well as end users, i.e., children and other types of visitors. The discussions with the stakeholders involved in the multidisciplinary teams and various field studies provided inspiration about new ways of using current technology to support the work of the professional guides in organizing and conducting the visits, with the aim of enhancing the overall visitors' experience. To better inform the design of new applications, we carried out a contextual inquiry; the collected data were particularly useful to elicit important requirements, which we reported in [6] through a usage scenario. That scenario is revisited in this paper in order to describe further requirements about the use of composition technologies to enhance the personal experience of different stakeholders accessing CH sites.

Recently, we have developed a platform that, according to End-User Development (EUD) and participatory design approaches [7, 8, 9], supports end users, not technology skilled, to create personalized visits to CH sites. End users become both information consumers and producers, and they are actively involved in the management of CH information. The platform indeed implements a new composition paradigm to allow end users to extract contents from heterogeneous (personal or third-parties) sources, and compose Personal Information Spaces (PISs) that satisfy their situational information needs and can be ubiquitously executed on different devices. The platform provides people with the means to integrate data, services and tools, enabling them to play an active role in solving their every-day problems. The platform is general and is flexible enough to be adopted in different contexts of use [10, 11, 12]. In this paper, we show how it can be used to support the work of professional guides when accompanying visitors to archaeological parks and how guides and visitors can collaborate among them to create new visit experiences.

With respect to what we described in [6], we have refined our initial prototypes and improved their execution on multiple devices. In November 2012, we have also performed a field study to understand potentials and limitations of the composition and use of PISs by real users. The study is described in [13]; it involved 2 professional guides and 28 visitors at the archaeological park of Egnathia in Southern Italy. One of the results of study revealed the guides' need to communicate, both synchronously and asynchronously, with their peers during the PIS composition, for example, to ask advice about new services that can provide material they are not able to find through the services they have access to. Guides would also like to share their PISs with visitors to allow them to view and possibly add contents. Thus, we are currently extending our prototypes to support collaboration, to enable PIS co-creation by a group of end users and PIS sharing among different groups of stakeholders. The new prototypes will be discussed at the workshop. In this paper, in order to briefly describe our approach, we illustrate the scenario for PISs co-creation, usage and sharing (Section 2). We then describe the platform architecture (Section 3) and conclude the paper by indicating some future work (Section 4).

2 A Scenario of PISs Co-Creation, Usage and Sharing

The scenario illustrated in this section is articulated in 4 steps, which show how different people use a PIS to: a) be actively involved in the management of CH information; b) collaborate among them to take part to the enrichment and access to cultural heritage information, c) share information with others. The first step is already supported by our platform prototype described in [6], while the development of functionality supporting the other steps is in part recently accomplished work and in part still on-going work.

Step 1. Content retrieval from heterogeneous sources and PIS composition
Giuseppe, the main *persona* of our scenario, is a professional guide who accompanies people during visits to various archaeological parks in Apulia. Giuseppe is organizing the visit for a group of Italian tourists to the archaeological park of Egnazia. When he was informed about the visit by the booking office, Giuseppe also received information about the visitors, e.g., their profiles, the time that they want to spend for the visit and previous experience in other archaeological parks. The visitors, i.e., a group of 10 adults, are keen of Messapian history. They have planned about three hours for the visit. They previously visited the Messapian archaeological parks of Vaste and Roca. Now, they want to visit Egnazia, even if it doesn't present many remains of Messapi due the preeminence of the remains of the successive Roman Age.

Giuseppe decides to set up a "virtual" visit to show to people the Messapian remains of Egnazia by using the large multi-touch display, installed at the Egnazia museum, and his own tablet device. Thus, Giuseppe using his personal PC opens his own workspace in the Web platform that allows him to gather and organize in advance multimedia contents to be shown during the visit. Giuseppe starts composing his PIS by collecting the material he finds through the platform's web services, i.e., 3D reconstructions from Google Sketchup, photos of other archaeological parks from FlickR, videos of excavation campaigns from Youtube. Giuseppe adds and geo-localizes interesting content on virtual map of the Egnazia park.

During the composition Giuseppe would like to add new content, but he does not find services, among those provided by the platform, what might satisfy his needs. Since he knows that Andrea is more experienced about the Messapian history of Egnazia, Giuseppe asks his help. Andrea accesses the platform and looks at the list of services: he chooses some services and annotates them to specify the reasons for his choice and to give some suggestions to Giuseppe. Andrea shares the annotation with Giuseppe. At this point, Giuseppe has more elements to choose useful services, selects some of them and adds them to his PIS.

Step 2. Pervasive, multi-device access to PISs

The day after, the visit to the archaeological park begins with a briefing in front of a multi-touch display. Giuseppe puts his badge on the display and the system recognizes him and shows the PIS he created the day before. Giuseppe starts his narration interacting with the multimedia contents showed on the multi-touch display: every time he wants to show details of a specific content, he taps the icon on the virtual map where he had placed the multimedia material. After the initial explanation in front of the multi-touch display, the visit continues in the archaeological park, where the visitors can view the ruins of the ancient buildings. Giuseppe invites visitors holding a 3G smartphone to download an app to visualize multimedia contents Giuseppe could possibly share with them. During the tour park, Giuseppe interacts with his PIS showed on his personal tablet.

Step 3. Continuous dynamic PIS evolution

During the discussion with the group, a visitor suggests to Giuseppe a web site that tells about a legend of a man, who landed in Roca and came to Egnazia by travelling across the Messapian region. Giuseppe decides to search this Web site, finds it, and mashes the page up into his PIS. Giuseppe has noticed that, among the search results, there is a video about similarities and differences of Messapian archaeological parks. He selects it and the video is showed on the tablet. The video appears very interesting and stimulates a discussion between Giuseppe and the group of visitors. At the end of the video, Giuseppe decides to further modify his PIS by inserting the retrieved video. Giuseppe invites those tourists, who have previously downloaded the app, to look at the video on their device. In this way, Giuseppe can illustrate details about the video.

Step 4. PIS sharing and reuse

Cloe is a teacher, who was in the group that was accompanied by Giuseppe to visit the Egnazia archaeological park. She has to illustrate to her students the Aeneid epic poem. Since she remembers the legend that Giuseppe told her in Egnazia, she decides to use the interactive whiteboard to show the web site to the students in order to discuss the various hypotheses about Aeneas's wanderings. Thus, through the interactive whiteboard, she accesses the public area of Giuseppe's PIS and looks at the material used during the Egnazia park visit. Cloe shows the web page, and invites the students to investigate if that legend and the Aeneid poem can be correlated. During the lesson, her students can use their tablet or smartphones to search evidences in the Web. If they find interesting information they can add them to the PIS that Cloe had previously shared with them. The students can continue the search at home, and can update Cloe's PIS. The day after, Cloe uses again the interactive whiteboard to discuss with the students the new multimedia material.

3 Platform Organization

A platform, based on a general-purpose mashup environment [10, 11, 12], has been developed to allow people, who are not technology skilled, to retrieve contents from heterogeneous sources and use them to compose PISs in order to satisfy their needs in specific situations. The platform is flexible enough to be adopted in different contexts of use and ubiquitously executed on different devices. We briefly describe here the main components of its architecture, shown in Figure 1; with respect to what we reported in [6], to which the reader may refer for more details, the architecture includes new modules enabling PIS sharing and annotations

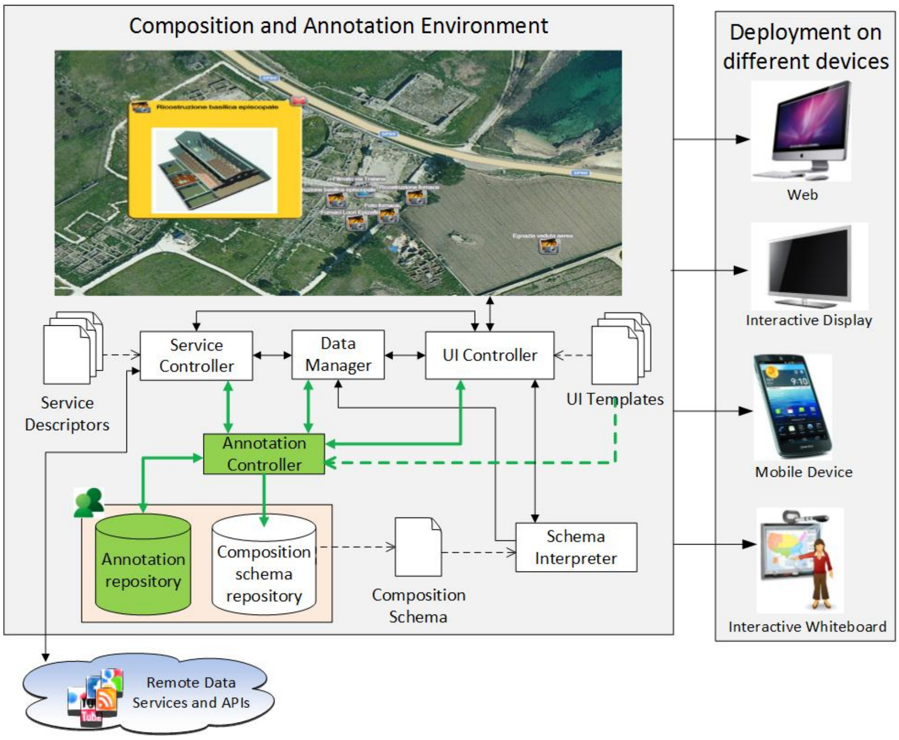

Fig. 1. Architecture of the platform for PIS composition, sharing and annotation.

The platform exploits a "lightweight" paradigm for the integration of heterogeneous resources, mainly adopting visual mechanisms through which end users, without any need to program or adopt complicated design notations, can express desiderata about the orchestration of different services. The accessible services have to be registered and described into the platform by means of *Service Descriptors*. Each descriptor specifies properties that the platform has to know for querying that service. Service registration is needed to prevent end users from dealing with technical properties when accessing a service. Different levels of service specification are possible: service descriptors can include only basic information, such as the service URI and the values of some parameters to build up simple queries, or they can specify multiple properties for more sophisticated service invocations. Service descriptors can be created by inserting the service properties into visual forms – no scripting or XML coding is needed; the registration module then translates the inserted data into XML-based descriptors. Therefore, even unskilled users would be able to add new services into the platform, provided that they know the basic required information. More complete descriptors, supporting sophisticated queries, could be instead created by technical users, e.g., platform administrators. Indeed, our platform especially suites meta-design scenarios, where different users with different skills cooperate to the creation of the composition environment for the end-user productivity.

Based on the registered services, end users, by means of a Web composition environment and through visual mechanisms suitable to their background and domain, compose contents, functions and visualizations. Figure 2 reports the customization of the composition environment for the guides of the archeological park of Egnathia [6]. The workspace composition proposes a map-based visual template to allow the guide to associate the content retrieved through services to specific locations of the park. The content is visually presented in different *resource windows*, each one associated to a specific service. The users select content items from such windows and, through drag&drop actions, associate them to points on the map. The result of such composition actions is immediately shown to the user, who can thus realize how the final application will work, and iteratively modify it, adding or dropping service items, until the desired version of the composite application is reached.

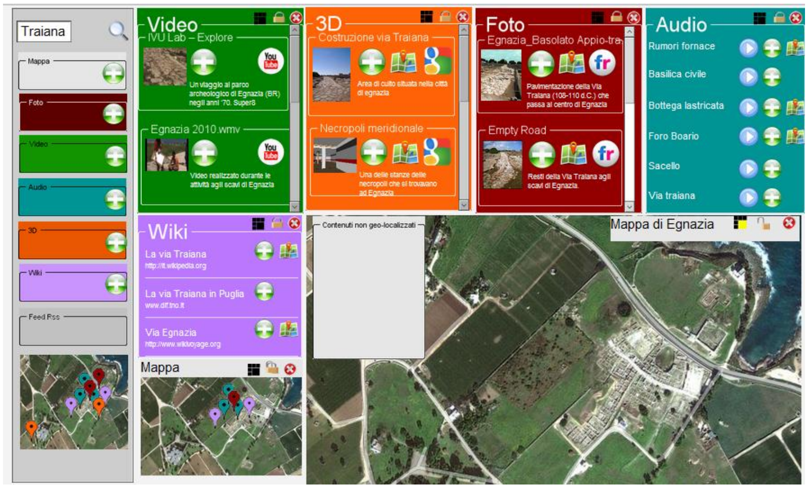

Fig. 2. The customization of the composition environment for supporting the activity of the professional guides at the archeological park of Egnathia [6].

This is only one example of visual composition that the platform is able to support, which mainly addresses the integration between the map-based service and all the other services providing contents. The platform indeed enables different kinds of integration (e.g., synchronization of different UI components [10] or combination of contents coming from different sources into an integrated data set [12]), which can then be properly exploited into the customized composition environments, depending on the emerging needs of the addressed user community. Also, the possibility to adopt specific visual templates, as the basis for the PIS composition, permits the definition of metaphors and composition actions that are meaningful for the target end-user community.

The visual composition actions performed by end users lead to the automatic creation of XML-based composition schemas, stored in the *Composition Schema Repository*. Some execution engines, developed for different client devices (e.g., Web browsers, large interactive displays, different types of mobile devices), interpret the created schema and dynamically generate the corresponding PIS. In particular, a *Schema Interpreter* parses the composition schema and then invokes the *UI (User Interface) Controller* that, based on the *UI Template* selected by the user during the PIS composition (e.g., the map template for the park of Egnathia park), dynamically generates the PIS user interface. The UI controller also invokes the *Data Manager* module, which in turn, based on the specification in the composition schema, queries the involved remote services through the *Service Controller*. The Data Manager is also in charge of storing (and managing the access to) possible user personal data stored in local repositories. The UI controller finally manages the rendering of the retrieved data through the visual elements of the adopted UI Template.

To accommodate the new requirements for sharing and communication emerged in the field study with professional guides and visitors, the platform has been extended to support the multi-user access to shared resources, and to facilitate the management and storage of annotations. The resulting annotation approach [14] allows the users to communicate with other stakeholders by adding comments on the different elements characterizing a PIS, namely the available services and the queries defined over them, specific content items retrieved through services, and the visual templates adopted for content visualization. The aim is to enrich the PIS with further information to be shared with others without corrupting the original resources.

To create annotations, the user visually selects the object to be annotated and adds the corresponding comment [14]. The *Annotation Controller* is in charge of interpreting the user annotation actions, identifying the annotation location and establishing whether it is related to services internal to the PIS composition, to services generally available in the platform -- and not necessarily included in a specific PIS -- or to specific content items or UI elements. The Annotation Controller communicates with the different modules managing the different levels the annotations can refer to. For example, when a user annotates a photo, the Annotation Controller retrieves the photo URI from the result set managed by the Data Manager (e.g., a Flickr result set with photos and their metadata), associates with this URI the note inserted by the user, and stores the created annotation to the *Annotation Repository*. The latter is used to store all the produced annotations and their associations with the original annotated PIS documents.

In order to contextualize the annotation within the specific situation of use in which it was created, the Annotation Controller receives a *state representation* from the different modules. The state representation makes it possible to present the annotations during later executions of the PIS, by reconstructing the original context where the annotation was created. For example, when an element of the UI is annotated, a set of properties of the template the UI is based on are also stored, such as the type of the template (e.g., a map) and the notable widgets that characterize the template (i.e., markers showing points on a map). Similarly, when the annotation refers to a service or to the service result set, the service settings and the specific query that was executed when the annotation was created are stored.

4 Conclusions

This paper has illustrated our approach to enhance the visit experience of archeological parks. Its distinguishing feature is the exploitation of composition technologies that allow different stakeholders to create, use and share Personal Information Spaces where heterogeneous content can be combined and manipulated to satisfy different information needs. The paper has in particular outlined how the use of composition technologies, adequately customized to the needs of CH stakeholders, leads to a personalized user experience while visiting CH sites, and to take advantage of that experience in different places and times, e.g., even after the visit and outside the archeological park, in a continuum where searching, composing and sharing increase the users' motivations to get actively involved and improve their knowledge.

It is worth remarking that several platforms, based on *mashup technologies*, have been proposed in last years, with the aim of being more oriented to end-users. However, they were not as successful as expected because they were based on composition languages that revealed inadequate for end users [15, 16, 17]. The adoption of a composition paradigm suitable for end users is instead a key goal of our approach to PIS composition.

Preliminary results of field studies that we are conducting with real users have helped us assessing the validity of the approach in the specific CH domain, especially in supporting the work of professional guides. However, we believe that the approach can be easily transposed to take advantage of other forms of cultural expression, for example in innovative contexts of "smart" urban cultural initiatives.

The field studies have also highlighted some desiderata of end users related to the possibility to collaborate with their peers and with the different stakeholders. For this reason, our current work is devoted to extending the composition approach to support further collaborative features to enable end users to co-create, also in live sessions, their artifacts.

Acknowledgements

This work is partially supported by MIUR and MISE under grants VINCENTE and LOGIN, respectively. The authors are grateful to Paolo Bottoni, Matteo Picozzi and Antonio Piccinno for their collaboration to part of the research work here described.

References

1. Merriman, T.: Public archaeology. Routledge, London, UK (2004)

2. Copeland, T.: Presenting archaeology to the public. In: Merriman, T. (ed.) Public archaeology, pp. 132-144. Routledge, London, UK (2004)
3. Costabile, M.F., De Angeli, A., Lanzilotti, R., Ardito, C., Buono, P., Pederson, T.: Explore! Possibilities and challenges of mobile learning. In: 26th annual SIGCHI conference on Human factors in computing systems, CHI'08, pp. 145-154. ACM, New York, NY, USA (2008)
4. Ardito, C., Costabile, M.F., De Angeli, A., Lanzilotti, R.: Enriching exploration of archaeological parks with mobile technology. ACM Transaction on Computer-Human Interaction 19, 1-30 (2012)
5. Ardito, C., Costabile, M.F., Lanzilotti, R., Desolda, G.: Integrating traditional learning and games on large displays: an experimental study. Educational Technology & Society 16, 44-56 (2013)
6. Ardito, C., Costabile, M.F., Desolda, G., Matera, M., Piccinno, A., Picozzi, M.: Composition of situational interactive spaces by end users: a case for cultural heritage. In: 7th Nordic Conference on Human-Computer Interaction: Making Sense Through Design, NordiCHI '12, pp. 79-88. ACM, New York, NY, USA (2012)
7. Schuler, D., Namioka, A.: Participatory design: principles and practices. Lawrence Erlbaum Associates, Inc. (1993)
8. Sutcliffe, A., Mehandjiev, N.: End-user development (Introduction to special issue), Commun. ACM 47, 31-32 (2004)
9. Lieberman, H., Paternò, F., Wulf, V. (Eds.): End user development, Springer, Dordrecht, The Netherlands, (2006)
10. Cappiello, C., Matera, M., Picozzi, M., Sprega, G., Barbagallo, D., Francalanci, C.: DashMash: a mashup environment for end user development. In: Auer, S., Díaz, O., Papadopoulos, G. (eds.) Web Engineering, vol. LNCS 6757, pp. 152-166. Springer Berlin Heidelberg (2011)
11. Cappiello, C., Matera, M., Picozzi, M., Caio, A., Guevara, M.T.: MobiMash: end user development for mobile mashups. In: 21st International Conference Companion on World Wide Web, WWW '12 Companion, pp. 473-474. ACM, New York, NY, USA (2012)
12. Cappiello, C., Matera, M., Picozzi, M.: End user development of mobile mashups, In: 15th International Conference on Human-Computer Interaction, HCII '13, Springer Verlag, in print
13. Ardito, C., Costabile, M.F., Desolda, G., Lanzilotti, R., Matera, M., Picozzi, M.: Personal information spaces in the context of visits to archaeological parks. Submitted for publication
14. Ardito, C., Bottoni, P., Costabile, M.F., Desolda, G., Matera, M., Piccinno, A., Picozzi, M.: Enabling end users to create, annotate and share personal information spaces, End-user development – IS-EUD 2013, LNCS, Springer, Berlin / Heidelberg, in print
15. Latzina, M., Beringer, J.: Transformative user experience: beyond packaged design, Interactions 19, 30-33 (2012)
16. Fischer, G., Giaccardi, E., Ye Y., Sutcliffe, A., Mehandjiev, N.: Meta-design: a manifesto for end-user development, Commun. ACM 47, 33-37 (2004)
17. Costabile, M.F., Fogli, D., Mussio, P., Piccinno, A.: Visual interactive systems for end-user development: a model-based design methodology, IEEE T. Syst. Man Cy. A 37, 1029-1046 (2007)

Chatting to Personalize and Plan Cultural Itineraries

Antonio Sorgente[1], Nadia Brancati[1], Cristina Giannone[2], Fabio Massimo Zanzotto[2], Francesco Mele[1], and Roberto Basili[2]

[1] Institute of Cybernetics "Eduardo Caianiello" of the National Research Council Via Campi Flegrei 34, 80078 Pozzuoli, Naples (Italy)
{a.sorgente,n.brancati,f.mele}@cib.na.cnr.it

[2] University of Rome Tor Vergata Department of Computer Science, Systems and Production, 00133 Roma (Italy)
{giannone,zanzotto,basili}@info.uniroma2.it

Abstract. In this paper, we present a system for the generation of cultural itineraries that exploits conversational agents to implicitly build formal user profiles. The key idea is that the preferences for user profiling are not obtained in a direct way, but acquired during a natural language conversation of the tourists with the system. When the user profile is ready, it becomes the input for the generation of the customized cultural itinerary. The proposed system, called DiEM System[3], is designed for dialogues in the domain of cultural heritage, but its flexible architecture allows to customize the dialogues in different application domains (cinema, finance, medicine, etc.).

1 Introduction

Offering personalized access to cultural artifacts is one of the most interesting challenges for promoting cultural heritage. Many artworks and cultural artifacts hardly reach the general public, even if someone can find such artifacts very interesting. Finding a way to help these people to reach the cultural treasures hidden in the sea of cultural offerings is an important mission (as in [1, 2]). Recommendation Systems (RS) are designed for this purpose: personalizing the experience of users (or tourists) by helping them to select cultural objects from a large set of possible alternatives. Thus, these systems are largely applied to suggest ways to visit cultural artifacts such as interesting places or artifacts in a town or in a museum [3–5].

Recommendation Systems need two basic capabilities: first, a way to *capture* user interests; second, a way to exploit this information to compute the best selection of items to suggest to the user. Generally, in tourism-oriented systems [3–5], the first activity is done by exploiting the history of the movements of tourists or asking to the users to fill out a form. Most systems allow only to

[3] The proposed system has been implemented within the DiEM project (Dialog-based Emotional Marketing) - Italian PON (Programma Operativo Pazionale) project.

suggest POIs (Points of interest) without generating customized itinerary for the tourist. Instead, the systems that suggest itineraries, propose only POIs near current location of the tourist. Generally, in the definition of the user profile, the time availability is not taken into account. This leads to plan itineraries that will probably be discarded by the user because they will not respect his time availability. As a rule, the systems that use the user profile for the generation of suggestions, allow the acquisition of user preferences by filling out forms. Instead, an interesting way to capture user interests is to talk with the users, see the Conversational Recommendation Systems (CRS) [6]. In this way, user can express preferences in a more natural way, i.e., chatting with a system able to acquire the user interests from the dialogue. But, these CRSs determine user interests by direct questions, e.g., *What type of food do you like?*.

In this paper, we then propose a novel approach to build Conversational Recommendation Systems for accessing cultural sites. The key ideas of our approach are: the preferences of the tourists are not asked directly but, rather, acquired during an information seeking phase performed in natural language with the system; the final output are cultural itineraries by connecting POIs taking into account time availability. Generally, during the activity of trip planning tourists search information about cultural attractions in the city and then identify the places to visit. Users implicitly declare their preferences by asking information about the cultural items and the touristic attractions of specific places. We exploit this natural interaction to gather these implicit preferences in user profiles. These profiles are then used by the planner that generates the cultural itineraries. Also, our idea, implemented in the proposed system, is to take into account the temporal availability of the user, and plan an itinerary based on the time that the tourist wants to dedicate to the visit.

The proposed system is designed for dialogues in the domain of cultural heritage, but the overall architecture of the system allows to propose dialogues in different application domains (cinema, finance, medicine, etc.). The customization of the dialogue is done by defining the ontology of the domain of interest.

The rest of the paper is organized as follows. Section 2 presents the related works in the area of recommendation systems. Section 3 introduces the DiEM system that aims to merge conversational recommendation systems with itinerary generators. Section 4 presents the application of the system to some running examples. Section 5 presents the planning for the evaluation. And, finally, Section 6 draws some conclusions.

2 Related Works

Recommendation systems have been largely used for applications in the tourism domain.

The first class of systems suggest points of interest (POIs) by only using the location of the user. Cyberguide [7] and GUIDE [8] are two examples of such systems. Cyberguide [8] was developed to provide maps and information about POIs in closed (e.g., a museum) and open environment. The GUIDE system, instead, provides information about the city of Lancaster. The system uses a WLAN access point to determine the position of the tourist and provides information through a web interface.

TIP (Tourist Information Provider) [9] introduces a personalization for the suggestions. It selects information about cultural sites to propose, based on user preferences, and suggests the ones that are close to the current position of the tourist. User preferences are acquired by asking to fill out forms.

Then, Varga & Groza and Chiu et al. [3, 10] introduced the use of knowledge bases for semantic searching, to enhance the way of detecting and selecting relevant information. For example, the recommendation system for travel planning presented in [3, 11], based on user preferences, uses DBpedia knowledge base for constructing a holiday trip, choosing a destination city.

Providing touristic itineraries is not new. Some systems that, in addition to make suggestions about POIs, provide itineraries were defined in [12–14, 5, 3]. Deep Map [12, 13] is a mobile system that generates customized itineraries for the city of Heidelberg. During the visit, the system analyzes the position of the tourist in the city and suggests information respect to his position, according to the objectives that the tourist wishes to achieve. The INTRIGUE system [14], instead, generates an itinerary to visit POIs selected by the user. The order of POIs to visit and how to move from a POIs to another one are specified in the generated plan. The planning takes into account the arrival and departure dates of the tourist. In the system presented in [5], the user gives his preferences about restaurants, accommodations, etc., and the system suggests nearby structures (spatially) to the place where the user is located, and suggests POIs that may be of interest for the user. DailyTRIP [4] is a web-based system which generates planning for several days, and for each day offers a different route. DailyTRIP takes into account the user position and preferences (time to visit the sites, the sites of interest for tourists, etc.). The system acquires the user preferences by filling out a form. Also, there is a mobile version of the system, called Mobile mTrip[4].

To the best of our knowledge, all these techniques have never been used in combination with a dialogue agent that implicitly extract user preferences by chatting about POIs for planning cultural itineraries in a city.

3 The DiEM System

This section describes our approach to generate cultural itineraries by exploiting conversational agents. Section 3.1 describes the general architecture of the system. Section 3.2 introduces our information seeking conversational agent that is used to derive user preferences. Section 3.3 describes the user profile generator. Finally, Section 3.4 describes the itinerary planner.

[4] http://www.mtrip.com/

3.1 The Architecture of DiEM System

The Fig. 1a shows the architecture of the DiEM system. The system is composed of three main modules: *Dialogue Agent*, *User Profile Generator* and *Itinerary Generator*. The Dialogue Agent manages the information seeking dialogue with the user. The User Profile Generator communicates, in a direct way, with the Dialogue Agent and acquires user preferences from the dialogue in order to generate the profile. The Itinerary Generator is the recommendation module: it generates cultural and personalized itineraries based on the profiles of the users.

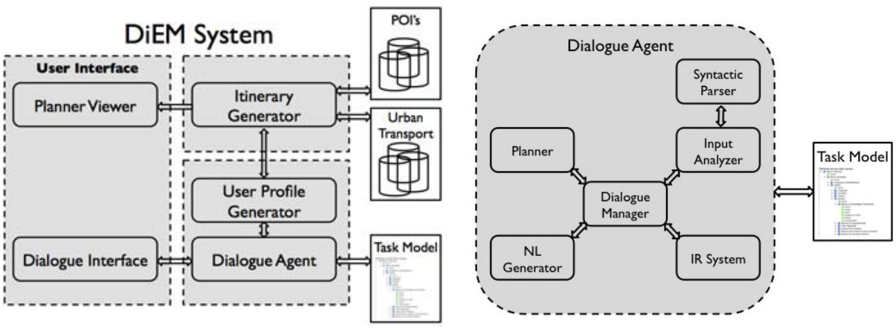

(a) Architecture (b) Dialogue Agent

Fig. 1: DiEM system

The system has a graphical interface that allows the user to converse with the system (Dialogue Interface) and to visualize the itinerary produced by Itinerary Generator (Planner Viewer). The proposed system has been designed for application in the cultural heritage, but its architecture can be adapted to different domains (Tasks), such as cinema, finance, medicine, and so on. In particular, it can be adapted by: (1) the personalization of dialogues, by defining the domain ontology, called Task Model (see section 3.2); (2) the description of user preferences for the generation of user profile, by defining the rules depending from Task Model; (3) the identification of recommendation services to use for the suggestion of items of domain.

3.2 Dialogue Agent

A Dialogue Agent is a system that communicates with users in spoken natural language in order to perform a specific task ranging from make travel arrangement to question answering in general telephone assistant.

The DiEM Dialogue Agent is designed to assist users to find information about specific cultural objects. The by-product of this interaction is the collection of their preferences. These are represented in a user profile as described in Section 3.3.

Here the processing of utterances and the management of the tourist have to be dynamically combined with the domain knowledge base (Task Model) that contains, in a taxonomy representation, the touristic information used by the dialogue agent during the dialogue session. A user can ask to the system question about cultural related objects (i.e., the ticket price of a museum) and tourist information in general (i.e., certified restaurants). The DiEM Dialogue Agent makes use of a frame-based model for information seeking in support of interactive Question Answering (QA) tasks. The aim is to limit the complexity for the user against cases where multiple and diverging answers are possible according to the retrieved information. The conversation helps the user to disambiguate the ambiguous interactions, to converge to the most useful answer(s) with the minimum number of interactions (turns). The dialogue agent here presented follows the architecture discussed in [15].

The internal modules are almost fully independent from the specific application and constitute a general framework for dialogue-based QA systems. Figure 1b shows the overall architecture of the Dialogue Agent. The modules do not interact directly with each other, but are governed by the Dialogue Manager module in a hub architecture fashion. The hub acts as a router which sends the requests to each specific module. A brief introduction of single module is hereafter discussed, for a more complete discussion see [15].

Dialogue Manager The central component in DiEM dialogue agent is the *Dialogue Manager*. It controls the activation and the flow of information across the individual dialogue components. During the coordination activities, the *Dialogue Manager* controls the dialogue evolution by storing the intermediate information (e.g. the focus of current turns) and activating the suitable state transitions according a finite state *dialogue model*.

The dialogue management task can thus be mapped to a form of navigation through a finite set of dialogue states. These states are directly mapped into a finite set of the speech acts, i.e., states that describe the user intention in the dialogue, such as *question, negation, clarification,..* originate from a limited set of dialogue states. They can be thus assimilated to a finite state (i.e. a Mealy) machine: every dialogue act is performed by the system according to a specific state and a user focus.

In a generic state $s \in S$, the automaton (1) detects the actual speech act characterizing the last user input, (2) changes its state according to the δ function and (3) outputs a symbol according to the output function λ. Output here includes the current topics and a dialogue act (e.g. clarification) used to trigger the Natural Language generation component. In the Dialog Agent, the transition function is also influenced by the current plan: given a state and the speech act characterizing the user reply, it produces the suitable next state. For example, clarifications may follow other clarifications in general, according to the next topics foreseen in the plan.

Natural Language Interaction In order to manage the user input in form of natural language utterances, the dialogue agent supports two levels of analysis: *Speech Act Recognition* and *Focus Detection*. Speech act recognition (SAR) supports the classification of the user utterances into the major Speech Act classes discussed in [16]. This reduces the huge variety of rules required to recognize acts from the user input. The SAR is accomplished by a data-driven supervised model called *Speech Act Classifier*. The analysis of input utterance, moreover, aims to detect the topic expressed by a user and it is implemented in the *Focus Extractor*. This component recognizes the *focus* suggested by a user, by analyzing the grammatical structure of his input utterance produced by the syntactic parser *CHAOS* [17, 18].

Task and Domain Model The interpretation of the user input depends on the available knowledge of the domain of application. An ontology infrastructure is available to the system where the domain and the task are represented. The *Task Model* represents the answers and their corresponding concepts organized into a hierarchy in which general concepts are on the top of the hierarchy, while the more specific ones are leaves. When a concept is selected as result of a user question it is marked as *interesting concept* for the user. This knowledge is accessed by an *Information Retrieval System*, that initially retrieves all the candidate answers to the user question. No constraint is posed to the IR engine except the requirement of deriving a relevance score associated to each candidate answer in the ranked list.

Planner Module The mechanisms used to find the correct response in the Task model, with respect to the user input, requires a plan of the interactions. This plan is based on the initial user question, if such question is ambiguous, may produce multiple answers. The *Planner* module, according to a probabilistic model, computes a set of interactions with the user in order to disambiguate the initial question.

Natural Language Generation The generation of the natural language output is under the responsibility of the *NL Generator* module. The NL generation is based on textual grammars, i.e. *NL Templates* [19]. Given the dialogue status, the proper system responses are detected and templates allow to efficiently compile the output according to the context (e.g. the user focus).

3.3 User Profile Generator

The DiEM User Profile Generator produces the user profile from the dialogue. It is important for the step of the cultural itinerary generation. In fact, the itinerary will be personalized depending on the profile of tourist who wants to visit the city. The system extracts the tourist features from the *interesting concepts* of the Task Model marked during the dialogue between the user and the system, i.e., preferences about what he wants to visit and behavioural aspects. This module transforms these features in constraints to be respected during the generation of the itinerary. In order to personalize the itinerary, the system takes into account: (1) tourist desires, the type of sites that the tourist wants to visit (church, museum, monument, etc.); (2) tourist interests (nature, food, exhibitions and special events that take place in the city during the period of the visit); and (3) tourist availability for visiting the city, total time that the tourist wants to spend in the visit and his consuetudes (for example the pleasure to walk). This profile is not selected within a pre-established list (for example a class of tourists interested in the church), but it is dynamically built depending on user preferences. The user profile adopted here is based on the user profile defined in [20].

As already mentioned, all the user information are acquired during the dialogue between the user and the conversational agent. The system extracts the information for generating the user profile, based on the questions that the tourist asks to the system. The user preferences are extracted with the application of a set of rules. For example, if the user asks the question "*Is there a botanical garden in Naples?*", it is possible to infer that the user is interested to visit sites like gardens, square, and in general, outdoor locations. Also, it infers that the tourist prefers to walk. In this way, the user preferences are enriched (or updated) for each turn of the dialogue.

The profile generation structure The Profile Generator Structure (PGS) is defined by the quadruple

$$\mathbf{PGS} = < F, P, I, R >,$$

where F is the set of features that characterizes the user profile. Each feature describes a property of the profile and it is expressed through a qualitative value. In the domain of cultural heritage presented, they are *pleasure_to_walk*, *pleasure_to_visit_castles*, etc. The value assigned to the features depends on the preferences P, which is the set of user preferences that the designer, who customizes the system in a particular domain, wants to evaluate. In according to the customization of the proposed system, the set of user preferences to evaluate are *church_interest*, *museum_interest*, *pleasure_to_walk*, etc. Also, the preferences are represented using quantitative values. The preferences have been introduced in order to analyze the user interests, which are fundamental for the generation of user profile. The correspondence between features F and preferences P is realized with the application of a set of rules R. For each feature f_i, a rule is defined as follows:

$$r_i = < f_i, c_i, v_i > \; with \; r_i \in R$$

where c_i is a condition and v_i is a value. Each condition is defined on one or more preferences and if it is true, the value v_i is assigned to the feature f_i. Finally, I is the set of influences. Each influence enables to define how an element extracted from the dialogue (such as focus, polarity, etc.) affects one or more preferences.

The influences are defined as follows:

$$i_i =< d_i, o_i > with\ i_i \in I$$

where d_i is a condition defined on the task model and o_i is a set of operations. During a dialogue when a sentence talks about a topic covered by the task model that matches with the condition d_i of the influence i_i, the set of the operations defined in o_i are executed. Each operation changes the value of a single preference.

An example. If we consider the previous sentence "*Is there a botanical garden in Naples?*" for retrieving information about user preferences and build his profile, we must define a PGS as follow:

$$
\begin{aligned}
P &= \{pleasure\_to\_visit\_garden, pleasure\_to\_visit\_square, pleasure\_to\_walk\}\\
F &= \{garden\_interest, square\_interest\}\\
R &= \{r_1, r_2\}\\
r_1 &= < pleasure\_to\_visit\_garden, garden\_interest > 5, high >\\
r_2 &= < pleasure\_to\_walk, garden\_interest > 5\ and\ square\_interest > 5, high >\\
I &= \{i_1\}\\
i_1 &= < topic = garden, \{inc(garden\_interest), inc(square\_interest)\} >
\end{aligned}
$$

From this structure, the Profile Generator System takes into account the elements extracted from each turn of the dialogue (such as focus and polarity) identified by the Dialogue Manager, and applies the influences defined in I. This step updates the values of the preferences defined in P. Updated the values of the preferences P, the Profile Generator System applies the rules defined in R that allow to define the values of features defined in F. The features F define the user profile.

3.4 Itinerary Generator

An helping system for visiting a city should suggest a cultural itinerary that satisfies the desires of the tourists related to what they want to visit, their behaviour and their time available for the visit. The Itinerary Generator proposed generates cultural itineraries that meet the need of specific users. The system plans the itinerary, performing the following steps (for a more complete discussion about the algorithms see [20]):

- *identification of the cities' areas.* The system identifies areas where are located cultural sites that tourist is interested in;
- *ranking of the areas founded in the previous step.* The system sorts the areas through an index that expresses the user interests of an area, defined in according to the user interests of the sites that belong to the area;
- *evaluation of the time required to visit each sites in the area.* For this evaluation, the system considers the user preferences (for example a user that goes quickly from a site to another or a user that lingers for a long time on some sites) and the time needed by the user to reach sites in the area on foot (in this case it takes into account the pleasure of user to walk);

- *identification of areas to visit in the days when tourists dwells in the city.* In this case, the system takes into account areas of greatest interest and with a visit time at most equal to the time that the tourist wants to dedicate in a day to visit. It provides the list of areas to visit for every single day (the visit of the areas in a day is organized to minimize the distance among areas);
- *identification of public transport to go from one area to another.* The system provides information about transportation to move among areas. The information indicates: the type of means of urban transport (bus, tram, metro), the line to use, the point of departure and arrival.

In addition, for each cultural site included in the itinerary displayed, the system shows (if available) descriptions and photos (they can be stored in multiple sources).

What happens if the system during the dialogue has not acquired enough information about the tourist preferences (*cold start*)? In these cases, the system for the generation of an itinerary uses the indices of notoriety regarding each cultural sites. These indices are used for the selection of the cultural sites to insert in the itinerary of a "generic" tourist. The indices of notoriety adopted are: (1) iconographic importance of a site (e.g. in Naples city, Piazza del Plebiscito and Palazzo Reale have an iconographic index higher than Cappella S. Severo); (2) number of annual visit for the site. The Itinerary Generator implemented is independent from the city to visit. As shown in Fig. 1a, the Itinerary Generator uses existent (external) resources to extract information about POIs and urban transport of the cities, that are different for each city. So, such resources will continue to work and to be updated by habitual users, without any change.

4 Running Examples

In this section, we present an example of dialogue (dialogue excerpt) between a tourist and the dialogue agent, and the related itinerary generated for the city of Naples. The Fig. 2 shows an interaction excerpt between user (U) and dialogue agent (A), and the type of user preferences extracted from it. This dialogue excerpt is a demo dialogue. In the boxes we explained the features that are activated. The first question of dialogue excerpt 1 in Fig. 2 suggests that the tourist is interested in the visit of museums, so the museums are inserted in the list of tourist preferences. Instead, in the second question, the tourist asserts that he doesn't have much time to visit museums: in this case the system acquires the information "little time to spend in the visit". This information will allow to plan the itinerary within the limits of temporal availability of the user. Continuing with the dialogue, in the dialogue excerpt 2 in Fig. 2, the first question suggests that the user is very interested in the visit of square, and using influences as presented in section 3.3, gardens and, in general, outdoor locations. Then, in the second question, the user expresses his pleasure to walk, and so the system, for planning the visit, takes into account this user preference. After the dialogue, the system processes the user preferences for the generation of profile and then

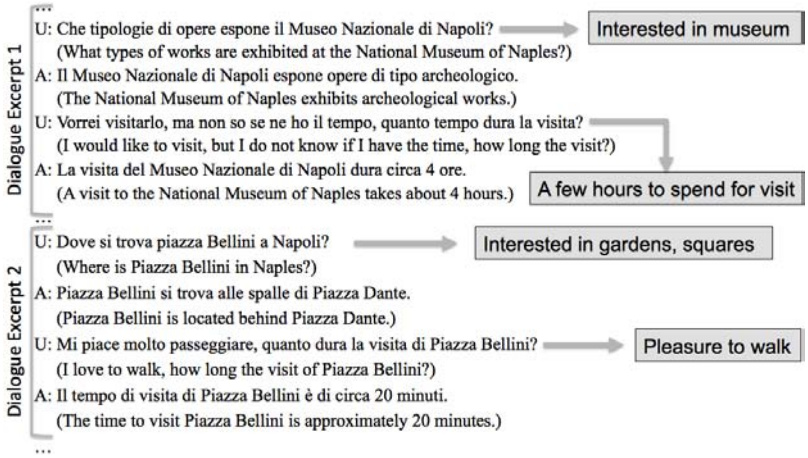

Fig. 2: Dialogue Excerpt

produces the customized itinerary that shows to the tourist through a layout similar to that shown in Fig. 3. The layout reports the map of Naples city. The

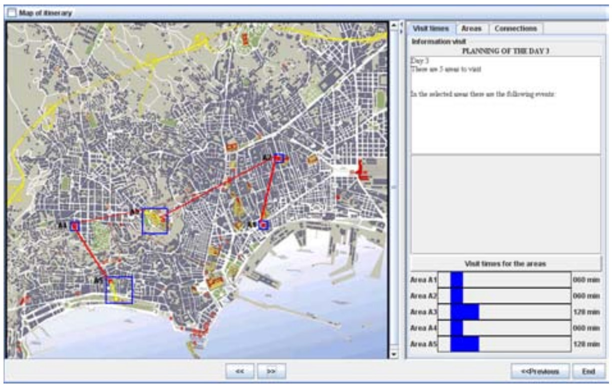

Fig. 3: Personalized Itinerary

map highlights the areas of interest for the tourist, the sites of interest for each area, and the path that connects the sites and areas (itinerary). Also, the layout shows information about urban transport to use for moving between the areas (point of departure, arrival point, and the number of the means of transport). Because the information about itineraries are different for each day planned, the system interface provides a layout for the tourist for each day planned. Also, if the system has information about cultural sites inserted in the planning, it allows tourist to visualize and to consult this information.

5 Evaluation Setup

The evaluation of dialogue systems, as well as for recommending systems, is problematic as it depends on a large set of factors, all of them highly depending on specific applications. Metrics that have been proposed for evaluate these type of systems aim to evaluate the *user satisfaction* in term of task success and to evaluate the implicit user profile acquisition versus the explicit one.

We plan to collect user satisfaction metrics via a questionnaire compiled by the users: they asked to fill out the questionnaire immediately after the completion of a system session. In order to make this evaluation more realistic, we plan to select users representative of a large variety of attitudes and criteria. The adopted criteria for evaluation suggested by the questionnaire were the following:

- *Topic understanding*: the ability of the system in recognizing the main focus of the originating question;
- *Meaningful interaction*: the quality of the system behavior according to the utility of each generated turn;
- *Topic coverage*: the user perception of how good is the system knowledge about the target topic;
- *Contextual Appropriateness*: the ability to produce clear turns consistent with the dialogue progress [21];
- *Interaction/Dialogue quality*: the overall quality of the system to generate sentences. It captures mainly the grammatical correctness of the NL generation;
- *Ease of use*: the usability of the system, that is the system friendliness perceived by the user;
- *User profiling*: The user preferences acquired by dialogue reflect the user profile;
- *Overall Effectiveness*: the user comprehensive judgment about the system usefulness for the planning of a tourist trip.

6 Conclusions

In this paper we have presented the DiEM System, a system for generating personalized cultural itineraries that can be customizable in different application domains such as cinema, finance, medicine, and so on. The main objective of this work has been to define a system for acquiring the user preferences and for generating of cultural itineraries in a natural way, through exploiting conversational agents. The DiEM System generates the cultural itineraries based on user profile that is acquired in implicit way by chatting with the tourist. The proposed approach takes into account user preferences such as tourist desires (type of cultural sites: church, museum, monument, etc.), tourist interests (nature, food, exhibitions and special events that take place in the city during the period of the visit), and tourist availability (temporal constraints, depending on total time for visiting and days of visit). Also, the system suggests information about urban transport to help the tourist to move from one area to another.

References

1. Beccaceci, R., Fallucchi, F., Giannone, C., Spagnoulo, F., Zanzotto, F.M.: Education with living artworks in museums. In: Proceedings of the CSEDU 2009 - ICCSE, Lisboa, Portugal (2009)
2. Zanzotto, F.M., Tudorache, A.G.: Travel with words: An innovative vision on travelling. IEEE/WIC/ACM Int. Conf. on WI and IAT **3** (2012) 107–111
3. Varga, B., Groza, A.: Integrating dbpedia and sentiwordnet for a tourism recommender system. In: ICCP, 2011 IEEE International Conference on. (2011) 133–136
4. Gavalas, D., et al, M.K.: Web application for recommending personalised mobile tourist routes. IET Software **6**(4) (2012) 313–322
5. Sharma, J., Singh, S.P., Singh, P.: Location based information delivery in tourism. IJCSCT **4(2)** (2012)
6. Thompson, C.A., Gker, M.H., Langley, P.: A personalized system for conversational recommendations. Journal of AI Research **21** (2004) 393–428
7. Abowd, G.D., Atkeson, C.G.e.a.: Cyberguide: a mobile context-aware tour guide. Wirel. Netw. **3**(5) (October 1997) 421–433
8. Cheverst, K., Davies, N., Mitchell, K.e.a.: Developing a context-aware electronic tourist guide: some issues and experiences. In: Proceedings of the SIGCHI conference on Human Factors in Computing Systems. CHI '00, New York, NY, USA, ACM (2000) 17–24
9. Hinze, A., Buchanan, G.: Context-awareness in mobile tourist information systems: Challenges for user interaction. In: International Workshop on Context in mobile HCI at the 7th Int. Conf. on HCI with Mobile Devices and Services. (2005)
10. Chiu, D.K., Yueh, Y.T.e.a.: Towards ubiquitous tourist service coordination and process integration: A collaborative travel agent system architecture with semantic web services. Information Systems Frontiers **11**(3) (July 2009) 241–256
11. Esuli, A., Sebastiani, F.: Sentiwordnet: A publicly available lexical resource for opinion mining. In: In Proceedings of the 5th Conference on Language Resources and Evaluation (LREC06. (2006) 417–422
12. Kray, C.: Situated interaction on spatial topics. PhD thesis (2003)
13. Malaka, R., Zipf, A.: Deep map: Challenging it research in the framework of a tourist information system. In Fesenmaier, D.e.a., ed.: ICT in Tourism 2000. Springer Vienna (2000) 15–27
14. Ardissono, L., et al., A.G.: Intrigue: Personalized recommendation of tourist attractions for desktop and handset devices. In: Applied AI, Special Issue on AI for Cultural Heritage and Digital Libraries. Volume 17. (2003) 687–714
15. Basili, R., et al.: Data-driven dialogue for interactive question answering. AIIA 2007 Artificial Intelligence and Human Oriented Computing (2007) 326–338
16. Bunt, H.C.: Context and dialogue control. THINK Quarterly **3** (1994) 19–31
17. Basili, R., Pazienza, M.T., Zanzotto, F.M.: Efficient parsing for information extraction. In: Proc. of the ECAI98, Brighton, UK (1998)
18. Basili, R., Zanzotto, F.M.: Parsing engineering and empirical robusteness. Journal of Language Engineering **8 (2/3)** (2002)
19. Stent, A., et al, J.D.: The commandtalk spoken dialogue system. In: Proceedings of ACL'99. (1999)
20. Gloria, M., Mele, F., Minei, G.: Generazione di itinerari culturali adattabili al turista. AI*IA 2003, Workshop AI for Cultural Heritage (2003)
21. Danieli, M., Gerbino, E.: Metrics for evaluating dialogue strategies in a spoken language system. In: In Working Notes of the AAAI Spring Symposium on Empirical Methods in Discourse Interpretation and Generation. (1995)

Creating Personalized City Tours Using the CHIP Prototype

Natalia Stash, Loes Veldpaus, Paul De Bra, and Ana Pereira Rodeirs

Eindhoven University of Technology,
Dpt. of Mathematics and Computer Science and Dpt. of the Built Environment,
Den Dolech 2, 5612 AZ Eindhoven, The Netherlands
{n.v.stash,l.veldpaus}@tue.nl
debra@win.tue.nl
a.r.pereira-roders@bwk.tue.nl
http://www.tue.nl

Abstract. This paper explores the applicability of the software prototype developed for personalized access to semantically enriched art collection of the Rijksmuseum in Amsterdam in a different environment – city rather than museum. As a case study we take Amsterdam, a World Heritage City, i.e. a city that includes urban areas designated as World Heritage (WH). This is the first step towards turning our prototype into a generic tool applicable for generating recommendations/personalized routes in indoor and outdoor environments based on semantically described data of the museum collection or points of interest in the city. Moreover we allow for user model information reuse between various domains/scenarios served by our Web-based application therefore addressing the cold start problem while starting to use a new application.

Keywords: Personalisation, cultural/world heritage, semantic enrichment, sightseeing, city guide, protected urban planet

1 Introduction

Personalized access to cultural heritage information attracts the attention of many researchers and practitioners. A variety of applications can be found for cultural places, such as museums ([1],[14]), cities ([3],[4]). The focus of this paper is more specifically, on World Heritage city guides to be used by locals and tourists to make them understand what makes the WH property outstanding. Some very interesting though non-adaptive mobile guides have been developed over the past years e.g. mTrip[1], Pocket Guide[2].

There exist databases for managing data about museum collections[3], and WH cities[4]. Moreover, there exist international standards for detailed descriptions of

[1] http://www.mtrip.com/
[2] http://pocketguideapp.com/en/city/map
[3] http://www.den.nl
[4] Protected Urban Planet (PUP), http://protectedurbanplanet.net

objects in a database such as Dublin Core[5], VRA[6], SKOS[7] and Spectrum[8]. Also for uniform content description terminology sources are being used. Terminology sources include thesauri, (controlled) keyword lists, taxonomies, classifications or ontologies such as the Getty Vocabularies[9] and the CIDOC Conceptual Reference Model[10]. The hierarchical and associative relationships, which are defined in such thesauri, help search engines in interpreting and grouping of heterogeneous sources. In this way the terminology sources can make a significant contribution to the effectiveness of semantic Web applications such as the MultimediaN ECulture demonstrator[11]. Although standards exist in the fields of museums and cultural objects[12], and combining those probably covers part of the heritage field, as for example CIDOC states "The term cultural heritage collections is intended to cover all types of material collected and displayed by museums and related institutions" and this "includes collections, sites and monuments relating to natural history, ethnography, archaeology, historic monuments" however the main aim remains Supporting Museum Documentation. Ontology creation for historic buildings has been explored only a little, particularly in relation to their conservation [3]. Only one attempt to build up an ontology around built heritage and WH was found[13]. This is a very timely and relevant recent attempt. However, such ontology is not yet used to build up nomination files for WH Properties. As such the descriptions found on the UNESCO WH List are not structured along this or any other ontology.

In a number of papers ([5],[6]) we presented the results of the CHIP project[14] (stands for Cultural Heritage Information Personalization/Presentation). As the project name says it dealt with providing personalised access to cultural heritage artefacts. The project was done between 2005 and 2009 in collaboration with the Rijksmuseum[15], Amsterdam, and Telematica Institute[16] in the Netherlands.

Within the project we performed the semantic enrichment of the Rijksmuseum collection data by connecting it to standard vocabularies and adding extra semantic relationships from these vocabularies; and we used it for developing our software prototype. With the help of this tool a museum visitor can prepare his/her museum tours in advance and follow them on a mobile device (if wifi or other location technology would be present inside the museum). The produced software is taking into account the specifics of the museum collection such as the way artworks are described, physical museum constraints, etc.. Data from

5 http://dublincore.org/
6 http://www.vraweb.org/
7 http://www.w3.org/TR/2005/WD-swbp-skos-core-guide-20051102/
8 http://www.collectionslink.org.uk/spectrum-standard
9 http://www.getty.edu/research/tools/vocabularies/
10 http://www.cidoc-crm.org/index.html
11 http://e-culture.multimedian.nl/
12 http://www.dlib.indiana.edu/ jenlrile/metadatamap/seeingstandards.pdf
13 http://www.cherplan.eu/cultureonto/
14 http://www.chip-project.org
15 http://www.rijksmuseum.nl
16 http://www.telin.nl/index.cfm?language=en

other museums can be described in a way similar to the way the Rijksmuseum collection is described, as will be shown in the following section. Therefore the software could be reused for the other museum collections as well.

Based on the knowledge and experience in developing tour guides for indoor environments gained within the CHIP project we aim at applying the same ideas in the outdoor environment. The same content-based recommendations based on semantically enriched data as discussed in [6] could be used for generating personalised routes along cultural heritage assets in indoor and outdoor environments. On the UNESCO website[17] one can find how those assets are described in official documents, per WH property. Analyzing those documents reveals a list of the attributes (WHAT is making them Outstanding) and values (WHY are they Outstanding). Those can be allocated in the city, so that a city guide can be prepared and the Outstanding Universal Value presented to the visitors by means of such a city guide. At the moment this analysis has to be done manually, and per case study as no official ontology has been agreed upon yet. If this would be the case, a more generic set of attributes could help on linking and matching sites to each other (e.g. show a WH Property with a *church building* as an attribute), which could then influence the recommendations within the application beyond the one city that is visited at the time.

The novelty of our approach is in the fact that the tool becomes more generic and applicable for entering either museum or city data (in a specific RDF format) e.g. by art experts or tour guides and that it allows for providing personalized access to this data for the museum or city visitors. Since our application is Web-based and the user models are being stored centrally on the server, they can be made available for all applications developed with our tool and placed on the same server. In this way the personalised museum and city guides can exchange and reuse/update information about the same user, e.g. if the user of our Rijksmuseum application indicates that (s)he likes *Baroque* style (the style of the major part of the artworks from the Rijksmuseum collection) then (s)he could be guided towards buildings in Baroque style in Amsterdam. In this way we can (a) connect applications for indoor and outdoor environments and (b) address the cold start problem. The CHIP software is open-source and platform-independent.

The rest of the paper is structured as follows. Section 2 discusses the specifics of WH Properties. Section 3 walks through an example scenario of generating a personalised tour through Amsterdam and describes the requirements of the CHIP software to include a museum or city data set. Section 4 provides conclusions for this paper and discusses some insights for future work.

2 World Heritage Cities, Specifics of Data Description

At the moment the WH List includes 962 heritage properties, 745 cultural, 188 natural and 29 mixed properties. This list is growing steadily, adding about 25

[17] http://www.unesco.org

properties annually. Nearly half of the current properties are located in urban contexts. Currently there are already over a thousand cities that have protected areas, inscribed on the UNESCO WH List, located in or at the outskirts of their urban areas, and a database of them is being built up on Protected Urban Planet site. A WH Property is listed for it being of Outstanding Universal Values (OUV). OUV is considered the highest level of significance, to be preserved as part of the WH of mankind as a whole [10] those outstanding values are conveyed by attributes, which can be tangible or intangible. The "qualities and characteristics seen in things - in particular the positive characteristics" [9] – embodying cultural values are defined as attributes, of which two types have been defined: tangible and intangible. The tangible attributes regard the legacy of physical artifacts such as "form and design; materials and substance and other internal factors". The intangible attributes regard non-physical aspects related to the cultural heritage properties, such as "use and function; traditions, techniques and management systems; location and setting; language, and other forms of intangible heritage; spirit and feeling; and other external factors" [11]. Every WH Property has such attributes and values, and those can be used as *characteristics* of the site to be described and mapped in the application.

Attributes, or their tangible results or representations can be mapped along the urban context to reveal the actual presence of the cultural significance as described in the official documents within its urban context. This way an overview of what is of value (attributes) and why they are outstanding (values) can be constructed per WH Property. Due to the lack of research in this field this paper takes a categorization based on what has been described, though with the remark that further ontology building would be very relevant . First of all the application considers the indicated difference between tangible and intangible attributes. Next, it uses a developed ontology of categorizing the attributes within eight cultural values[13]. Last, it categorizes on different levels of urban scale: building element, building, and urban. The building element scale could for example include signs and symbols on facades, or to the use of the same type of building material. The building scale could refer to specific building types or uses, or to urban objects e.g. bridges. On an urban scale the value could be found in attributes like the urban structure, a historic route, squares, or the roofscape. Such system would be improved or changed if research reveals a more adequate categorization.

To capture all levels of scale, this application takes the perspective of the urban settlement which includes the WH Property in contrast to existing applications on WH, which are mostly focused on the protected site only. As such, the application can be used to discover what is of value within the WH Property (attributes and values) but also the WH Property context (e.g. indicated key views, buffer zones, related conservation areas) and at the same time the user has the overview on the spread or concentration of sites, and specific attributes and values across the city. In addition, this way the application could in the future also include attributes and values that are listed at the national or local level.

As a case study for our city guide we take Amsterdam with its WH Property *Seventeenth-century canal ring area of Amsterdam* inside the Singelgracht. First, we have to decide upon the data description. Fig. 1 shows an example description of an artwork from the Rijksmuseum collection. Every artwork is provided with an image, textual description, information about creator(s), creation site, year of creation, material medium, material support, dimensions, exposition place (room number in the Rijksmuseum), a list of associated art topics (or themes).

Fig. 1. Artwork description

Fig. 2 describes the semantically enriched museum collection data. Connections were made to three Getty thesauri: AAT - Art and Architecture Thesaurus; TGN - Thesaurus of Geographic Names; ULAN - Union List of Artist Names, and Iconclass[18]. There are three kinds of relationships in semantically-enriched data about the Rijksmuseum collection:

- *Artwork feature* is an explicit relation between an artwork and a concept. E.g. the artwork in Fig. 1 is related to the concept “Gabriel, Paul Joseph Constantin” via the artwork feature “creator”, the concept “Scheveningen” via the artwork feature “creationSite” and the concept “Dutch landscapes” via the artwork feature “subject”.
- *Semantic relation* is a relation that links two concepts e.g. “teacherOf”, “style”, “broader/narrower”.

[18] http://www.iconclass.nl/home

– *Implicit relation* connects two concepts that do not have a direct link between each other but can be deduced if both (e.g. "Rembrandt van Rijn" and "Chiaroscuro") are used for annotating a large number of the same concepts, as discussed in [6].

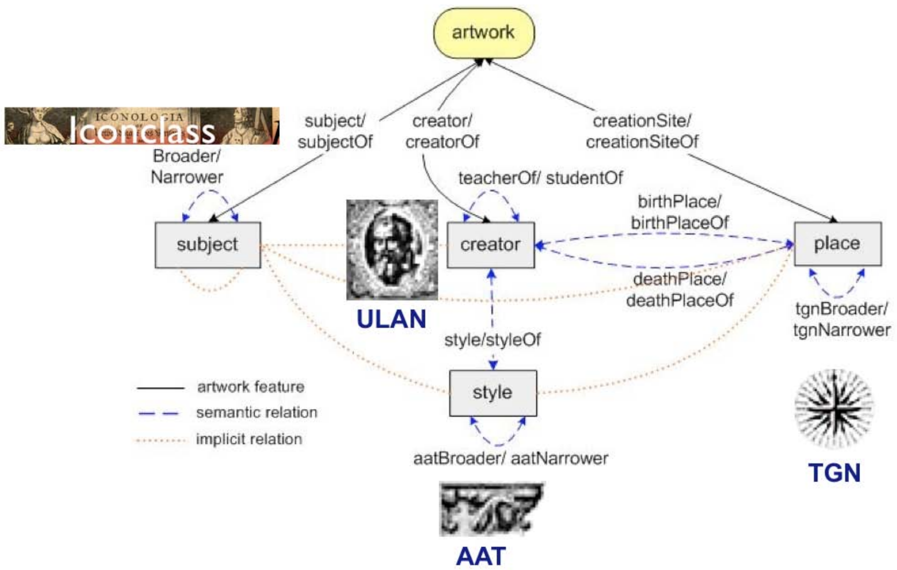

Fig. 2. Semantically enriched museum collection data

In a city guide we exclude the "Creation site" since it is obvious that this is the city that a person is visiting, "Dimensions" could refer to the number of floors or other descriptions indicating the size of the building, etc., "Exposition place" can be better renamed to "Location" and refer e.g. to the district of the city, etc.. For a point of interest (POI) in the city we need to add the following categories/attributes:

– *scale*, e.g. urban object element,
– *geographic location*: latitude, longitude,
– whether it is *tangible* or *intangible* attribute,
– *described value of attribute*, sentence context,
– *attribute value*: age, historic, scientific, aesthetic, social, political, economic.

The next section walks through an example scenario of generating a personalised tour through Amsterdam and explains what type of data has to be prepared from the side of the author of the application.

3 CHIP Prototype: From Personalized Museum Guide to Personalized City Guide

3.1 The Appearance of the City Tour

You can view the city guide at http://www.chip-project.org/cityguide and compare it to the walkthrough the CHIP demonstrator for the Rijksmuseum collection at http://www.chip-project.org/demo/chip_walkthrough/index.html.

Fig. 3 shows the *Sightseeing Recommender* page (after a number of interactions). It shows POIs in Amsterdam that you can rate (in the *Rate these points of interest* part), and once you do that it also shows a list of recommended POIs and topics that you can rate as well.

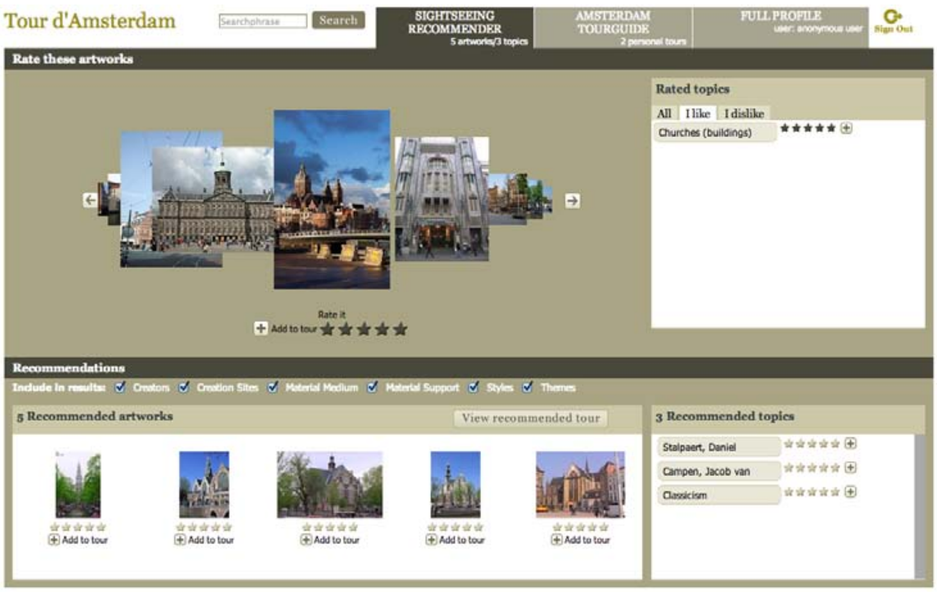

Fig. 3. Sightseeing Recommender - recommender for points of interest in the city

If the user wants to know more about some POI before rating it (s)he can click on its image and a full description will appear in the popup window (see Fig. 4). In this popup window the user can also see topics associated with the selected POI e.g. *Jacob van Campen* and *Stalpaert Daniel* – the architects of the *Royal palace*, the architectural style of the building – *Classicism* (currently shown in a list of themes and not separately), etc.. The user can toggle between options "Hide/Show namespace prefix" to see from which vocabularies/thesauri/specifications topics come from – *aat* prefix means that it comes from AAT thesauri, *ams* – base namespace for Amsterdam tour guide. The idea is that the initial set of POIs in the carrousel (Fig. 3) has a rich variety in topics. Rating some of these POIs thus gives a lot of information about person's preferences while visiting the city.

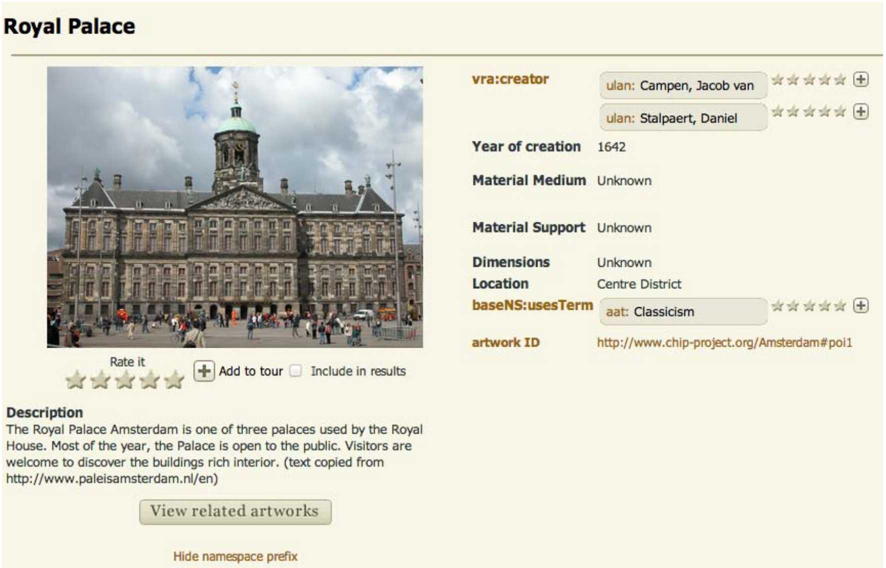

Fig. 4. Description of the point of interest

In the *Amsterdam Tourguide* tab the user can see the personalised tours' list:

- automatically generated *Tour of favourites* containing POIs rated with 5 and 4 stars, and *Tour of (20 top) recommended points of interest*,
- manually created tours.

In contrast to the museum guide there is no museum map view – only a visualisation of the tour on Google maps (Fig. 5) and a historical timeline (Fig. 6). In the current version of Google maps view only a set of POIs is displayed but the route between these points is not calculated yet.

After selecting the *Mobile Guide* tab the user can see what his/her tour will look like on a mobile device (see Fig. 7). In the first screen of the *PUP Sight Guide* on Fig. 7 the user logs in using his/her existing account or chooses the "Guest account" option, if (s)he hasn't worked with the demonstrator yet. (The second scenario is discussed in [14]). PUP stands for *Protected Urban Planet*. In the second screen the user can choose a tour to follow and adjust settings like number of POIs and the duration of the tour. In the third screen the user is presented with the carrousel of POIs in the selected tour. While following the tour the user can give ratings to POIs and related topics. Based on the ratings the rest of the tour can be adapted if the "Adapt tour?" option was selected

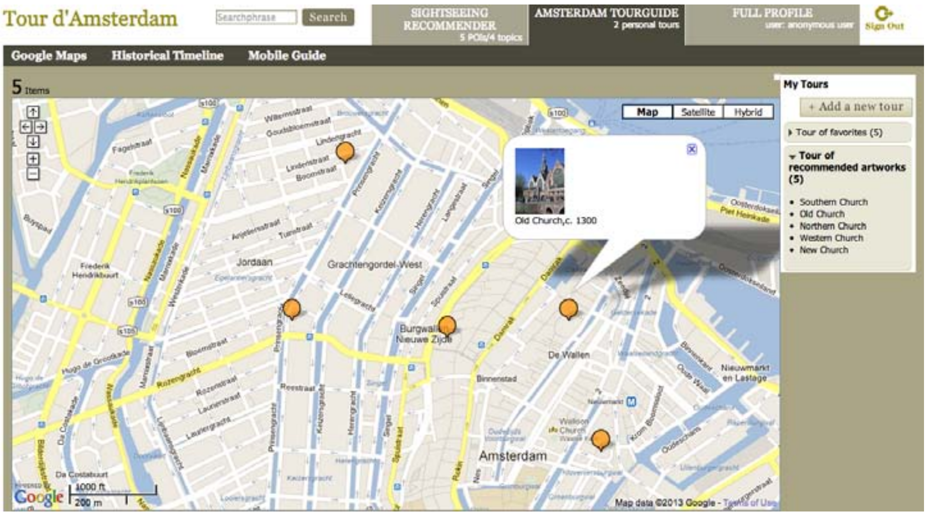

Fig. 5. Tour view on Google maps

Fig. 6. Tour view on historical timeline

Fig. 7. Mobile App *PUP Sight Guide*

in the second screen of the PUP Sight Guide. The last screen shows the tour on Google maps. An essential part of the GUI in the last three screens is the menu bar which aids the user in navigation. Icons in the menu bar link to tour configuration, tour overview, tour map, search function, help and logging out.

3.2 Initialization of the Demonstrator

If an art expert or city tour guide wants to use the CHIP software for providing visitors with the possibility of creating personalised tours through the city (s)he has to prepare the following information:

– RDF model describing the POIs in the city e.g. *Amsterdam.rdf*. See an *Example description of a point of interest* below:

```
<rdf:Description xmlns:vp="http://www.getty.edu/vocabularies/vp#"
   rdf:about="http://www.chip-project.org/Amsterdam#poi1">
   <rdf:type rdf:resource="http://www.vraweb.org/vracore/vracore3#Work"/>
   <vra:title>Royal Palace</vra:title>
   <vra:date>1642</vra:date>
   <vra:description>The Royal Palace Amsterdam is one of ...</vra:description>
   <vra:creator rdf:resource="http://www.getty.edu/vocabularies/ulan#500115589"/>
   <vra:creator rdf:resource="http://www.getty.edu/vocabularies/ulan#500065031"/>
   <ams:usesTerm rdf:resource="http://www.getty.edu/vocabularies/aat#300056513"/>
   <tgn:latlng>52.373, 4.891</tgn:latlng>
</rdf:Description>
<rdf:Description rdf:about="http://www.chip-project.org/cityguide/royal_palace_m.jpg">
   <rdf:type rdf:resource="http://www.vraweb.org/vracore/vracore3#Image"/>
   <vra:relation.depicts rdf:resource="http://www.chip-project.org/Amsterdam#poi1"/>
   ...
</rdf:Description>
```

– list of POIs' URLs that appear in the carrousel in the Recommender e.g.:

```
http://www.chip-project.org/Amsterdam#poi1
http://www.chip-project.org/Amsterdam#poi2
...
```

– (optional) weights for implicit relationships as explained in [6] e.g.:

```
0.229983829830428 http://www.chip-project.org/Rijksmuseum#encyclopedia47499
http://www.getty.edu/vocabularies/ulan#500011051
...
```

After starting up the server the author should go to the initialisation page to specify settings such as the names of the tabs in the demonstrator (e.g. *Tour de Rijks* or *Tour d'Amsterdam*), the location of the user profiles on the server, whether the demo is used for indoor or outdoor environment for making a decision upon showing/hiding tabs such as museum map, etc..

4 Conclusions and Future Work

Based on the knowledge and experience in developing tour guides for indoor environments gained within the CHIP project we aimed at applying the same ideas in the outdoor environment. Moreover we decided to investigate the usability of the existing software in a different domain – city rather than museum. By moving from museum to city we performed some testing of the existing CHIP software and improvement to make it a more generic tool. The presented *improved* version of the prototype only requires a proper description of the city data in RDF format, specifying the list of POIs to appear in the carrousel in the Recommender page and optionally the weights for implicit relationships, and an extra configuration step to choose the right data set (particular museum or city) to be used by the demonstrator. In fact, from the authoring perspective it is even easier to use the tool for preparing a city guide rather than a museum guide since a museum application would require the additional use of museum maps and specifying the coordinates (on the museum map image) for rooms, doors, hallways, artworks locations for constructing and visualising the route on the museum map.

All presented tools for generating/following personalised tour(s) – Recommender, Tour Guide, PUP Sight Guide – are Web-based, written using Java Servlets, HTML5/JSP, CSS, JavaScript, etc.. The idea is that Recommender and Tour Guide can be used e.g. for preparing a visit to the city in advance and next, the mobile PUP Sight Guide can be used for guiding the visitor through the city. The first next step is to turn the PUP Sight Guide into a GPS-enabled Web application. At the current state PUP Sight Guide has all the advantages of Web-based mobile apps such as platform-independence, easier and cheaper updates of a Web site than of a native app, independence from App stores, etc.. We should nevertheless take into account the advantages of the native apps such as targeting the specific limitations and abilities of the device in a much better than a Web app can while running inside a browser. We are planning to look into the existing frameworks such as PhoneGap[19] to see what kind of quality native app they can produce from our Web-based mobile app and whether this quality is sufficient.

We plan on improving PUP Sight Guide in a way that it could be used on the spot without preparing tour(s) in advance e.g. when the user selects the "Guest account" option (as shown in Fig. 7):

- calculating the optimal route on the fly given the POIs to visit, including starting point and end point of the tour,
- importing/using information about the visitor from social sites[20],
- information about what other visitors liked,
- taking into account opening hours of churches/museums to adjust the tour,
- making suggestions based on weather information, e.g. "It is going to rain in the coming hour, maybe it is best to visit a museum first", or time left "You should hurry, otherwise you'll miss your train".
- adapting the story content based on visitor's language, age, context e.g. time:
  - *evening/night* – "Look how beautifully and romantically the bridges over Amstel are lightened at night",
  - *day time* – "Come back here at night to see how beautifully and romantically the bridges are lightened".

[19] http://phonegap.com/

[20] e.g. Facebook, http://www.facebook.com; Twitter, http://www.twitter.com

Acknowledgments. This work is supported by EIT ICT Labs[21] and is based on the experience gained within the CHIP project. The CHIP project was funded by the Dutch Science Foundation program CATCH[22].

References

1. Stock, O., et al : Adaptive, intelligent presentation of information for the museum visitor in PEACH. User Modeling and User-Adapted Interaction, 17, 257–304 (2007)
2. Aoki, P. M., Grinter, R. E., Hurst, A., Szymanski, M. H., Thornton, J. D., Woodruff, A.: Sotto voce: exploring the interplay of conversation and mobile audio spaces. In: Proc. of the SIGCHI Conference on Human Factors in Computing Systems: Changing Our World, Changing Ourselves (2002)
3. Cheverst, K., Mitchell, K., and Davies, N: Gullivers Genie: The role of adaptive hypermedia in a context-aware tourist GUIDE. Community. ACM 45, 5 (2002)
4. G.M.P. O'Hare, M.J. O'Grady: Gullivers Genie: A Multi-Agent System for Ubiquitous Intelligent Content Delivery. Computer Communications, Vol.26, Issue 11, pp.1177–1187, 2003 Elsevier (2003)
5. Stash, N., Aroyo, L., Wang, Y., Rutledge, L., Gorgels, P.: Semantics-driven Recommendations in Cross-Media Museum Applications. In: Proc. the PATCH Workshop on Hypermedia and Adaptive Web-Based Systems, Germany (2008)
6. Wang, Y., Wang, S., Stash, N., Aroyo, L., Schreiber, G.: Enhancing content-based recommendation with the task model of classification. In: Proc. International Conference on Knowledge Engineering and Management by the Masses (EKAW), pp. 431–440 (2010)
7. Malcolm, T., While, A.: Ontology and the Conservation of Built Heritage. Environment and Planning D: Society and Space 27, no. 4, pp. 721-737 (2009)
8. Pereira Roders, A.R. and Hudson, J.: Change Management and Cultural Heritage. In E. Finch (Ed.), Facilities Change Management. Chichester: John Wiley & Sons, submitted / in press (2011)
9. Mason, R.: Values in Conservation Planning: Methodological Issues and Choices, in M. de la Torre (ed.) Assessing the Values of Cultural Heritage. Research Report, pp. 530. Los Angeles, CA: The Getty Conservation Institute (2002)
10. Convention concerning the Protection of World Cultural and Natural Heritage, Paris, available at: http://whc.unesco.org/en/conventiontext/, UNESCO (1972)
11. UNESCO, The Operational Guidelines for the Implementation of the World Heritage Convention, Paris: United Nations Educational, Scientific and Cultural Organization (UNESCO), (2011)
12. Veldpaus, L., Pereira Roders, A., Colenbrander, B.J.F. : Urban Heritage : putting the past into the future. The Historic Environment: Policy & Practice, 4(1), pp.18–33 (2013)
13. Pereira Roders, A.R.: Re-architecture: lifespan rehabilitation of built heritage – Book I, II, III. Eindhoven: Technische Universiteit Eindhoven (2007)
14. Stash, N., Veldpaus, L., De Bra, P., Pereira Roders, A.: The Protected Urban Planet App "PUP Sight Guide": Amsterdam as Case Study, To be published in the proceedings of KUI 2013, Conference Culture and Computer Science, Berlin, Germany (2013)

[21] http://www.eitictlabs.eu/, TIMS 12113 T1204 10015214, Emergent Social Mobility
[22] Continuous Access to Cultural Heritage, http://www.nwo.nl/catch

Interaction with Linked Digital Memories

Ilaria Torre

Department of Computer Science, Bioengineering, Robotics and Systems Engineering
University of Genoa, Italy
ilaria.torre@unige.it

Abstract. This paper presents an ongoing project on Web of Things and Linked Data. Physical objects are provided with a digital memory that enables them to store data about themselves and to link other resources in the Linked Data cloud. This possibility radically transforms the interaction with real objects and the user experience with the Web. In this scenario, personalization seems an opportunity and a requirement. The physical objects used in the project are museum artefacts. The paper describes architecture and main modules.

Keywords: digital object memories, Linked Data, personalization, museum

1 Introduction

The availability of technologies for the identification and connection of physical objects and the possibility to access them via HTTP are changing the modality to explore and experience the Web. The Web is going to become an extension of reality, making boundaries between physical and virtual objects more and more faint. As packaging is often mixed with products, web data associated to objects are going to fade into objects.

According to the paradigm of the Web of Things, real world objects can be connected in a network and accessed by using the web standards and protocols. The Semantic Web of Things (SWoT) uses web standards and protocols to expose physical objects as web resources and exploits the technologies of the semantic web to ensure semantic interoperability between systems that access and interact with heterogeneous physical objects.

Furthermore, the Linked Data (LD) principles provide a publishing paradigm in which structured data can be easily exposed, shared and linked by using HTTP URIs for identifying data and the RDF graph-based model to represent them. As a consequence, by using a URI that identifies the digital representation of physical objects, virtual and physical objects can be connected in a network of Linked Data [8]. Lower level identification systems enable the association of the URI to the physical object. The main technologies that are used for identification are RFID, QR-Code, Semacode and visual recognition.

Building a digital representation of a physical object makes objects digital memories which can be passive or active and can be used for different purposes, ranging from storing temporary data obtained by sensors to storing and representing complex information. For example, memories can store information about the provenance of the object, about its history, its use ant its position. Moreover they can also store the flow of messages generated by users while interacting with the object or while interacting with other users through the object. Barthel et al. [2] refer to them as Digital Object Memories.

The availability of information about an object while interacting with that object has many applications, such as job support, emergency handling, learning by doing, entertainment and edutainment. The linking with other pieces of information related to that object in the cloud of Linked Open Data (LOD) extends the possible applications and radically transforms the interaction with real objects and the user experience with the Web. In the domain of cultural heritage there are several examples of applications that experiment the Internet of Things technologies [1], [2] and examples of initiatives and institutions that follow the principles of LD to expose and connect artefacts with related data [5], [12]. However there are still few applications that combine them [6], [10].

This contribution presents an ongoing project which combines these approaches with the aim of enhancing the interaction with museum artworks by exploring the world "behind" the artwork. The key features of the project are: (i) search and exploration of the object memories and of the LOD cloud (ii) personalization and adaptation of content and user interface to the usage context and to the media used for interacting with the artwork.

The paper is structured as follows: the next section presents a background concerning LD and Object Memories and describes an architecture for representing them. Section 3 describes the main components of the prototype under development and presents an example scenario with two use cases. Finally, Section 4 sketches the next steps and concludes the paper.

2 LD and Object Memory: Background and Architecture

The so-called five-star rating scheme identifies increasing levels of structuring and linking of data towards a Web of Linked Data [3], with the requirement, for all levels, of an open licence to be Open Data.

There are several European projects and groups that work on LD. Organizations in the field of CH are particularly interested in the possibility to share and link knowledge about CH artefacts and repositories. Examples of cultural heritage libraries that make available their data according to the principles of LD are the British National Library[1] and the German National Library[2]. Both provide free access to their RDF/XML data, under a Creative Commons licence. Other active institutions are the museums. For example, the Amsterdam Museum has its entire collection represented as LD [5]. Several tools and services have been developed to support organizations to publish their data as RDF. Indeed, only a fraction of open data is currently repre-

[1] http://www.bl.uk/bibliographic/datafree.html
[2] http://www.dnb.de/EN/Service/DigitaleDienste/LinkedData/linkeddata_node.html

sented as RDF[3]. Examples of services aimed to facilitate the conversion of tabular data to semantically enriched data are Any23[4], Tabels[5], RDF Refine[6]. Moreover, D2R[7] server is a tool for accessing relational databases as RDF-graphs. They are particularly useful in scenarios when huge legacy data of an institution have to be extracted and transformed into RDF.

Linking data to other datasets or to more general vocabularies such as DBPedia is even a more costly task. The task of interlinking automation is faced in European and international projects. E.g., the 7FT LOD2 project[8] (2010-2014) dedicates a WP to the issue of reuse, interlinking and knowledge fusion, addressing the automatic and semi-automatic link creation with minimal human interaction. Example of frameworks and tools for supporting the search and interlink of data are Linked Media Framework[9] and the DBpedia Lookup Service[10].

Considering physical objects, additional requirements have to be satisfied in order to manage them as LD. A layered architecture makes easier the separation of functions by making layers autonomous.

In our project, we defined a three-layered architecture consisting of a *physical layer*, a *layer of the object memory* and a *LD layer*. The physical layer includes: the physical object (named entity in **Fig. 1**), the specific modality of interaction with the object, such as pointing, scanning, touching [11] or using a mediation device (e.g., a smartphone), and the modality of identification of the object, such as RFID, QR-Code, Semacode and techniques for visual object recognition. The layer of the object memory manages the description of the object and the way to access it, according to the Object Memory Model (OMM) [2]. The LD layer describes the physical object according to the LD principles.

OMM is a proposal of reference model for representing structure and content of object memories submitted to W3C [2]. In OMM, a digital object memory is described as a repository of digital data that is linked with a physical artefact, and may be populated with static or dynamic data from entities that interact virtually or physically with the artefact. This repository may exist at the artefact itself, outside, or both. The model includes an OMM header, a list of blocks and a set of metadata for describing blocks. The header contains the unique identifier for the object memory. In our architecture it is an HTTP URI. The model provides a list of blocks and a set of metadata that can be extended. Each block concerns a specific information fragment (a topic such as the object provenance, content and the level of protection required). A set of metadata is used to annotate each block, mostly based on Dublin Core: the

---

[3] For example, on PublicData portal, an aggregator of dataset descriptions from numerous other European data portals, only 459 out of more than 17.000 datasets are available as RDF [8].
[4] http://any23.apache.org/
[5] http://idi.fundacionctic.org/tabels/
[6] http://refine.deri.ie/
[7] http://d2rq.org/d2r-server
[8] http://lod2.eu/
[9] http://semanticweb.org/wiki/Linked_Media_Framework
[10] http://wiki.dbpedia.org/lookup/

block ID, Namespace, Format (and encoding), Title and Description of the block, Subject (free text tags, ontology concepts) and Link (used to split and distribute blocks in case of space constraints). The OMM layer is useful to manage aspects such as memory distribution, changing identifiers or physical relationships.

In our architecture, over this layer we have the LD layer. It is obtained from a subset of blocks of the OMM layer, automatically extracted and published as RDF (by LD wrappers). This layer is optional. It has to be implemented if the owner of the object wishes it to be accessed by RDF browsers and crawlers and if the owner wishes the object to be linked by other data in the LOD cloud. In order that wrappers publish data as LD, mappings between OMM metadata and RDF vocabularies have to be defined, together with links pointing to external and internal data sources.

Specific applications implementing the architecture can define how and where to store the object memory: locally in the object, or on a server, which can be easily accessed via HTTP. For example, Tales of Things[11], a platform using OMM, stores object memories on a server. Users can add stories about their objects and connect to other people who share similar experiences. Users have to print blank tags, stick them to the physical objects they want to have a digital memory and populate the memory by using an App which provides tools for accessing the object memory via HTTP.

Linking the object memory to other data in the LOD cloud presents opportunities for personalization, similarly to personalizing the access to open corpora, but with the advantage of managing structured data, which makes possible using filters in SPARQL queries and reasoning on data. This enables combining information from different sources and providing users with personalized mashup services, as in [12], or performing recommendations based on similarity measures between data as in [13].

3 Components of the System

To validate the approach, we are working on developing the modules in charge of identifying objects and accessing them. They include:

1) A configuration tool for the identification of the physical object and the association to its digital memory,
2) An authoring system that enables the writing, structuring, annotating and interlinking of the artefacts,
3) A system for the personalized access to the object memory and to the linked data, and for the adaptation of the interaction.
   Fig. 1 displays a diagram showing data flow and connections among architecture layers and components for the personalized access to the object memory. It will we described below, when presenting an example scenario.

1. Configuration tool. The layered architecture in Sec. 2 does not specify any particular identification system to be used. In our prototype we use the method of QR-code. To this end we have developed a tool that generates the QR-code and prints labels that can be attached to artworks. The QR-code contains the URI of the object memory.

[11] http://talesofthings.com/

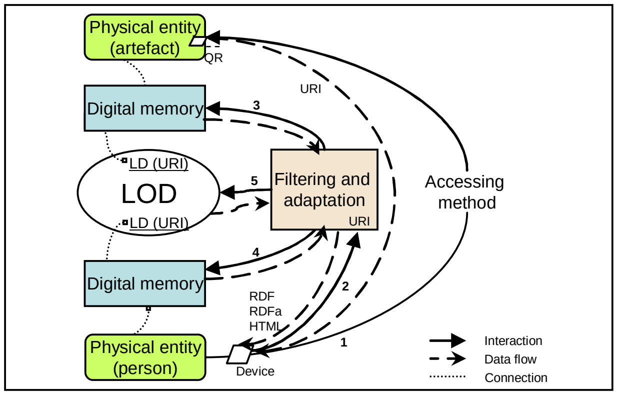

Fig. 1. Interaction diagram showing data flow and connections among architecture layers and components for personalized access to the object memory.

2. Authoring system for the knowledge definition.

The authoring system consists in a set of back-end modules made available to the administrator to populate the object memory associated to each artwork.

These modules provide support on three main tasks.

i) Querying the LOD cloud, in order to find vocabularies and reference ontologies that can be used to describe the physical object and to (inter)link the object with other resources. This task can be done manually by the administrator, using search tools on ontologies, such as DBPedia and Yago, or in a semi-automatic way, by using DBpedia Lookup APIs[12] and the Linked Media Framework[13]. Once they are collected, they are validated by the administrator who will use them in the next step.

ii) Editing the object memory upon the physical layer, by using the metadata specified in the OMM architecture, and defining mappings with vocabularies and links identified in the previous step. Related resources discovered on LOD will be used by LD wrappers in two ways: to connect to a similar resource (sameAs) and to connect to a resource that is used as a description (hasTopic).

iii) Adding optional materials about the object (e.g. movies, technical data, etc.). Additional materials can be provided in different versions, thus enabling the personalization of the returned data. Notice that also these materials are objects. Therefore they should be identified with an HTTP URI too.

3. Personalization module.

When interacting with physical objects, the interaction modality and the mediation device used to interact with the object has usually constraints which make very critical the task of personalization. A further reason that makes personalization important

[12] The Web service is based on a Lucene index providing a weighted label lookup, which combines string similarity with a relevance ranking (similar to PageRank) in order to find the most likely matches for a given term [4].

[13] The Linked Media Framework provides tools for querying LD and for integrating the returned results within the application.

is that when the user interacts with a physical object, s(he) has often time constraints and noise around him/her. In our project, this module manages two main tasks.

i) Personalized access to the object memory. The OMM model has been designed to support a basic level of personalization, used to filter content before downloading it. It is useful to adapt the content to the context domain, to the user profile an in case of limited bandwidth or connectivity. The set of metadata for each block provides information about data creator, data format, data content and blocks' location.

ii) Adaptation of the interaction. Physical objects can be of different type and can have different capabilities. Objects with processing capability can make recursive queries contacting other physical objects in the local network or in the LOD cloud.

Objects equipped with processing capabilities can also perform another type of adaptation, which concerns the adaptive composition of content to be displayed to the user. For this task, the project will exploit the features of the MM-Learning framework [9] that we developed to perform adaptive composition of materials and activities in order to suite the device constraints.

Example scenario. The idea at the basis of object memories is that objects can store all relevant information about them, accumulated along time. This concept is already widely accepted and applied to goods and products. However, it can be extended to all kinds of objects and to people too. Museum artefacts can have a long history. Several properties describe an object and each one may have an history. Typical properties are the artwork title, its author, date of creation and place, artistic current, the artistic phase of the author's life. Other information concern provenance, e.g., the origin, whether it was commissioned, purchased, moved from a place to another, stolen, retrieved. Other information concern whether it was repaired, which techniques have been used for restauration, its current quoted value, price fluctuation, etc. Other information can be about experts' reviews or visitors' comments.

A layered architecture, with information fragments organized in topics has two main advantages. First, it enables all these pieces of information to be managed in a flexible way. For example, it enables a LD wrapper to publish and make discoverable just a subset of the information stored in the object memory. Moreover it makes easier the filtering of content, based on the context of use and on the user profile. For example in the context of a museum, topics concerning restauration and quoted value can be filtered out, even though they could be re-filtered in if the visitor is an art merchant.

Considering **Fig. 1**, we sketch two use cases.

1) Context: *museum*; application used to access the object memory: *normal browser + QR reader*. Flow: (n1) the user/device interacts with the physical object (QR-code) and acquires the URI of the object. (n2) Http Request to the URI:

- If LD layer is *not set*, the URI is an active web page (public or private) which accesses the object memory, filters content and returns an HTML document.
- If the LD layer is *set*, different configuration options are available, according to LD principles: a 303 redirection with content negotiation is used to return the HTML page describing the artwork requested by the brwoser, as above; or the URI uses RDFa embedded in a HTML document (in this case the wrapper has to

use the mapping files to retrieve the RDF vocabularies and producing the RDF triples and links).

In both cases, the URI is a dynamic page which accesses the object memory (n3) and filters content based on the context of use (in this scenario, it filters out restauration and quoted vaule, that do not match the museum context).

2) Context: *museum*; application used to access the object memory: *specific App* which works as a RDF browser; LD layer: *set* . Flow: n1 and n2 are the same as in the first use case. Given that LD layer is set, the returned URI will be a RDFa document or a 303 URI. However, in this use case, the 303 redirect does not return a HTML document but RDF, since RDF browsers accept RDF+HTML. Moreover, in this case, before accessing the object memory (n3) and filtering content based on the context of use, the personalization module will access the user profile (n4) provided by the App. **Fig. 1** displays the case when the user profile is stored in an object memory representig the user. Given the user profile, the personalization module will decide whether to filter out other content or re-filter in content previously rejected. A further step is the exploration of the LOD cloud (n5) in order to follow RDF links and discover new links, as discussed in Sec. 2 and 3.

4 Next Steps and Conclusion

The project started on Jan 2013. Currently we have defined the architecture and the main modules that will perform the service. The modules for the identification of the physical object and for the association to its digital memory have already been developed. Modules 2 and 3 are under development.

An issue that has still to be faced in the workflow concerns the digital memory update, given changes in the physical object. Another issue that the current project does not address specifically but which is an interesting extension is enabling the object memory to be written by people and even to be updated, which requires defining roles and validation criteria.

The infrastructure will be useful to experiment new models of edutainment, based on using mobile devices, making education a pervasive process, which can be built and adapted to the context of interaction and to the specific user. Considering this underlying objective, another aspect of the project that requires to be detailed and improved is the provision of learning activities related to the visit of the museum or related to the artefact.

The experimental evaluation will have to test the validity of the approach and the performance of the prototype. Validating the approach can mean several features to be analyzed/compared and in particular: the user's satisfaction, her/his retention of what seen/experienced, the correctness of personalization. Concerning the performance, the evaluation will have to test in particular the time to return and visualize results. Emaldi et al [6] found that the time employed by the smartphone to visualize the information mainly depends on the number of triples to be processed and not on the number of triples retrieved. The evaluation should test the benefit of filtering content before download, based on the OMM features.

Acknowledgments. This research has been funded by the University of Genoa, within PRA 2012 projects, Prot. N. 2142.

References

1. Angelaccio, M., Basili, A., Buttarazzi A., Liguori ,W.: Smart and Mobile Access to Cultural Heritage Resources: A Case Study on Ancient Italian Renaissance Villas. In Proc. of the 2012 IEEE 21st International Workshop on Enabling Technologies: Infrastructure for Collaborative Enterprises (WETICE '12), Washington, DC, USA, 310-314 (2012)
2. Barthel, R., Kröner, A., Haupert, J.: Mobile interactions with digital object memories, Pervasive and Mobile Computing 9 281–294 (2013)
3. Berners-Lee, T.: Linked Data - Design Issues. http://www.w3.org/DesignIssues/LinkedData.html (2006)
4. Bizer, C., Lehmann, J., Kobilarov, G., Auer, S., Becker, B., Cyganiak, R., Hellmann, S.: DBpedia - A crystallization point for the Web of Data. Web Semantics: Science, Services and Agents on the World Wide Web 7(3), 154-165 (2009)
5. de Boer, V., Wielemaker, J., van Gent, J., Oosterbroek, M., Hildebrand, M, Isaac, A.:, van Ossenbruggen, J., Schreiber, G. (2013). Amsterdam Museum Linked Open Data, in Semantic Web Journal, In Press, DOI 10.3233/SW-2012-0074
6. Emaldi, M., Lázaro, J., Laiseca, X, López-de-Ipiña, D.: LinkedQR: Improving Tourism Experience through Linked Data and QR Codes. Proc. of Ubiquitous Computing and Ambient Intelligence (UCAmI 2012), LNCS 7656, 371-378, Springer, December (2012)
7. Ermilov, I., Auer, S., Stadler, C.: Crowd-Sourcing the Large-Scale Semantic Mapping of Tabular Data. Proceeding of the ACM Web Science http://svn.aksw.org/papers/2013/WebSci_CSV2RDF/public.pdf (2013)
8. Heath, T., Bizer, C.: Linked Data: Evolving the Web into a Global Data Space. Synthesis Lectures on the Semantic Web, 1:1, 1-136. Morgan & Claypool (2011)
9. Mercurio, M., Torre, I., Torsani, S.: Design of Adaptive Micro-Content in Second Language Acquisition. Journal of e-Learning and Knowledge Society, 7(3), 109-119 (2011)
10. Parra, G., Klerkx, J., Duval, E.: More! Mobile interaction with linked data. In Proc. of the Int. Workshop on Data-Centric Interactions on the Web, within the 13th IFIP TC13 Conference on Human-Computer-Interaction (INTERACT 2011) vol. 817, 37-47 (2011)
11. Rukzio, E., Broll, G., Leichtenstern, K., Schmidt, A.: Mobile interaction with the real world: an evaluation and comparison of physical mobile interaction techniques, Proc. of the Conf. on Ambient Intelligence, AmI-07, Darmstadt, Germany, 7–10 Nov. (2007)
12. Solomou, G., D., Kalou, A., K., Koutsomitropoulos, D., A., Papatheodorou, T., S.: A Mashup Personalization Service Based on Semantic Web Rules and Linked Data. In Proc. of the Conference on Signal Image Technology & Internet-Based Systems (SITIS '11). IEEE Computer Society, Washington, DC, USA, 89-96 (2011)
13. Thalhammer, A.: Leveraging linked data analysis for semantic recommender systems. In Proceedings of the 9th international conference on The Semantic Web: research and applications (ESWC'12) Springer-Verlag, Berlin, Heidelberg, 823-827 (2012)

Testaccio, A Digital Cultural Biography App

Maurice de Kleijn[1], Chris van Aart[2], Niels van Manen[1], Gert-Jan Burgers[3] and Henk Scholten[1]

[1]VU University Amsterdam, Spatial Information laboratory (SPINlab), The Netherlands
mtm.de.kleijn@vu.nl
[2]VU University Amsterdam, the network institute, The Netherlands
cj@vanaart.com
[3]Koninklijk Nederlands Instituut in Rome (KNIR) / VU University Amsterdam, The Research institute for the heritage and history of the Cultural Landscape and Urban Environment (CLUE), Italy/ The Netherlands
gj.burgers@knir.it

Abstract. Features from the past in present cities have been identified as valuable and as potential assets for future urban development. Transferring knowledge about the history and heritage of the urban landscape is considered to be a challenging matter. This paper discusses how the process of knowledge transfer between past and future oriented disciplines can be optimized by combining mobile geospatial app technologies with the methodological framework of the biography of the urban landscape. A digital cultural biography app for Testaccio, a district in Rome, is presented which aims to optimize the process of knowledge transfer, thus optimizing the use of urban cultural heritage. This best practice is considered to be highly innovative for the integration of heritage in current and future cities.

Keywords: mobile app, history, heritage, geospatial technologies, SDI

1 Introduction

This paper is about the design and deployment of mobile location based app technology to be used by spatial planners and architects who have to deal with physical and non physical remains from the past in the current urban landscape and incorporate these somehow in their designs for the future landscape. As identified by many scholars, the potential for heritage in urban development is significant [1-6]. However, the management of heritage in an urban context is considered a challenging matter. Especially in cities with a high density of physical and non-physical remains from the past, like Rome, the challenge is how to deal with and how to use these heritage features in the current and future city. One of the main challenges is to transfer the scientific knowledge about the history and heritage, generated by past oriented scholars like historians and archaeologists, to more future oriented disciplines like architecture and spatial planning. The landscape biographical approach developed by Kolen offers a methodological framework in which the gap between past oriented and future oriented scholars can be bridged, and thus help to optimize the use of physical and non physical heritage [7-9]. The biographical approach refers to the continual passing on of the landscape from one ´owner´ to the ´other´, shifting from one social context to another and influencing, and being shaped by, successive generations of inhabitants [7]. It combines different historical sources and past oriented studies, enabling scholars to identify and reconstruct dynamic processes of development and transformation. The information produced by these scholars, in turn acts as a source of inspiration for architects and urban planners. An important component of the biographical approach is therefore the ability to share data produced by the different disciplinary fields based on a location. In order to make this possible we have developed a Spatial Data Infrastructure (SDI) [10]. To make this data understandable, thus usable for future oriented scholars, the data needs to be elaborated and enriched, and functionality should be incorporated that facilitates the combining of information from different source. This paper focuses on the development of a *cultural biography mobile location based app* as an extension to the biographical research approach. It presents the development of the app for the neighborhood of Testaccio in the city centre of Rome, an area that is rich in heritage features and currently in flux. The app offers a tool to transfer the knowledge about the past dynamics of physical and non physical features to future oriented disciplines. Most map based apps are either for navigation or finding points of interest. Our app by contrast has been customized for professional usage, taking into account requirements for both past and future oriented disciplines. The app includes a cartographic storyboard of major transformation processes in Testaccio, derived by historians from historical maps of the area and presented in a series of maps with consistent font and key, designed to communicate the historical knowledge effectively to non-historians. It also makes the information accessible through a matrix on which time and theme (themes were defined by the intended end users) can be interactively combined, as the user sees fit. By thus integrating a pre-structured historical narrative with the matrix function that allows the user to create his or her own narratives, the digital biography aims to meet the variety of needs for historical and heritage information that the architect/ designer/ urban planner has in the course of the design process.

2 The Context: Testaccio

The district (*rione*) Testaccio, in the city centre of both modern and ancient Rome, is rich in archaeological and historical features. The area underwent numerous transformations, many traces of which remain in the current urban landscape. In the Roman period, Testaccio housed the river harbor of the *Urbs*. In the 3rd century AD it was incorporated within the Aurelian city walls. Major monuments are the Cestiam Pyramid, the Porticus Aemilia and the Monte Testaccio, a 35 meters high artificial hill consisting entirely of ancient pottery. In the late Roman period, Testaccio transformed drastically. The harbor structures were in decay and the area came to be used for agriculture, notably viticulture. Only in the late 19th/early 20th century Testaccio was *re-urbanized*, when it was built up with blocks of flats, accommodating the laborers so crucial to Rome's new urban and industrial development. The processes of the de-urbanization and the re-urbanization of Testaccio also produced physical evidence. The non-Catholic cemetery of Rome can be seen as a physical proof of this. According to 18th century rules of the Roman Catholic Church, non-Catholics were not allowed to be buried in the consecrated earth of Rome [11]. Instead, they had to be buried outside of the city. Whilst Testaccio was considered to be part of the *Urbs* in ancient times in the 19th century, it was peripheral indeed and considered to be exterior to the city. Only the subsequent re-urbanization re-integrated Testaccio into the urban tissue.

Currently, Testaccio is experiencing an urban regeneration process, partly through private initiatives, partly steered by municipal authorities. In the context of this process, initiatives are being developed to give credit to the district's history and heritage. The *Soprintendenza Speciale dei Beni Archeologici di Roma* (SSBAR) for instance formulated a strategy which aims to reconstruct the use of the district in the past two millennia and to develop strategies for re-using, redeveloping, thus *rediscovering* Testaccio's history and heritage; strategies that meet the modern needs of the district [3], [12].

3 The Conceptual Design of the Digital Cultural Biography App for Testaccio

The purpose of the digital biography app is to generate a shared view on the history and heritage of Testaccio for both past- and future oriented scholars.

Because the different disciplines and information sources –figuratively– meet at a specific spot, location is placed at the heart of app. Being able to access, add and share information on an interactive map while being at a specific location in a non rigorous way is identified as important functionality needed. This especially accounts for architects for whom fieldwork is an important part of their work progress and browsing associatively through the information enables them to generate a personalized narrative in their designs [8]. The digital cultural biography app for Testaccio is therefore developed as a mobile app with an interactive map and dynamic handling of information.

In order to transfer the disciplinary data on the history and heritage into information for other disciplines, three key components have been identified. Location, period and theme are seen as the key to transfer data into information, thus generating knowledge. The data to be shared geographically consists of points of interest and cartographical information. The points of interest link to stories, pictures, drawings and other documents from different sources. The cartographical data consists mainly of historical maps. Over 150 historical maps of Testaccio have been identified in different archives.[1] The historical maps form a valuable source of information. By overlaying the different historical maps geographically, the transformation dynamics in Testaccio from the 17th century until the present day can clearly be identified and studied. The digital cultural biography of Testaccio has therefore a number of georeferenced historical maps on top of the current topography.

For future oriented scholars to get a grip on the significance of the heritage features, having the information spatially accessible is not sufficient. Besides mapping the information about the history and heritage of Testaccio, the digital cultural biography app has therefore categorized the information chronologically and subdivided in different themes. A thorough analysis of the historical and heritage features has led to the classification of five different periods and eight different themes. Enriching the datasets, including the maps, with both attributes (era, theme) enables users to interactively explore the history and heritage of Testaccio. In addition to the historical maps, for every period a thematic map is produced in which the most important transfor-

[1] Thanks to a thorough survey by Joris Jehle.

mations of Testaccio are thematically summarized. Categorizing the information on the history and heritage of Testaccio as such enables future oriented scholars to browse through the information associatively. By offering functionalities in which they can spatially combine different periods of time and different themes enables them to arrive at a better understanding of the significance of the remains from the past, thus enabling them to generate personalized narratives in their designs.

Fig.1. The digital cultural biography app for Testaccio, showing the map of Nolli from 1748.

4 Technical Design

The core of the digital cultural biography Testaccio App is to enable users to access and add information based on a location, through an interactive map using a mobile device. To integrate these functionalities the app is developed for tablets and is developed by making use of geospatial technologies. The Testaccio app is developed as an app for iPads with 3G functionality. This type of tablet has an integrated GPS sensor, which enables it to communicate its location to the apps installed. For the development of the app Apple's iOS developer tool package `xcode` together with ArcGIS's runtime desktop Geographic Information System (GIS) software and the VU Amsterdam University library Geoplaza SDI are used [13-15]. These toolboxes enabled us to implement overlay functionality, drawing geometry functionalities and access information on the points of interest.

The thematic cartographical and thematic geographic information layers generated for the Testaccio app were prepared making use of desktop GIS software and integrated in the Geoplaza SDI, which provides mapping services that communicate with the app through the internet. Scans of the maps were georeferenced based on similarities with current topographical layers to a projected coordinate system. By analyzing these projected historical maps and analyzing archaeological and historical sources, thematic maps were generated in which the most significant features of Testaccio according to theme and period were created. Having this spatial information stored geographically, ArcGIS's runtime is able to apply overlay functionality and integrate the GPS signal on the map for an iOS app.

The Testaccio app is developed as an app that can also be used offline. In order to use the app when being in Testaccio, we have decided not to be dependent on mobile internet. The signal for the internet connection is too weak for the amount of data that the app needs. Especially the historical maps require too much data transfer and are therefore mostly cached from the Geoplaza SDI to the device. Although techniques to optimize this process are under development [16] and available in the Geoplaza SDI, the main bottleneck is the weak internet connection. The performance of the app in the short term had the priority. A disadvantage of this approach is that the app is very large – more than 1000 MB – and less interactive when data is updated.

The app is freely available in the Apple iTunes store [17].

5 Concluding Remarks and Future Research

The digital cultural biography app presented in this paper enables future oriented disciplines to access information on the history and heritage of Testaccio. By categorizing the information according to different themes and periods, past and future oriented disciplines create a common view which should enable them to gain a better insight in the significance of the heritage in the area. The interaction between the different disciplines takes place at a specific location. The role of geospatial technologies in this sense is therefore considered to be innovative.

The impact of the digital cultural biography app for Testaccio will be assessed in the next stage of research. A design concourse is organized in which architects are invited to produce a design for *Piazza Orazio Giustiniani*; half of the contestants will work with the app, the other half with a conventional textual history of the area. The assignment for the architects is to develop a design that takes the history and heritage of the square into account. Although a thorough analysis of the development of the app has been conducted, organizing this concourse will enable us to measure the actual impact of the app on the use of heritage in the designs, thus optimizing the potential of heritage in the urban landscape. The impact of the instrument will be evaluated by comparing the designs on different criteria, evaluating the design process with questionnaires, and interviewing members of the public about their views on designs created with and without the digital biography.

Integrating the spatial information on the history and heritage of urban areas on a conceptual level as presented above is considered to be a new way of stimulating and optimizing re-use of urban heritage in current and future cities. However, on a practical and technical level the app has to improve. We are currently at the turning point at which more data becomes available to the public and in which mobile internet networks can handle more data transfer. This app is therefore considered to be a best practice, which for future heritage issues can be used more generically by integrating different sources of open data and be part of a user centric Spatial Data Infrastructure (SDI) for the history and heritage of urban landscapes [10].

Acknowledgements

The research presented is part of research strategy *Rediscovering Landscape*, implemented jointly by the SPINlab and the CLUE Institute. The authors would like to thank Renato Sebastiani (SSBAR) for his active support and the invitation to join the Testaccio research program. Furthermore the authors like to thank Menje Almekinders and Joris Jehle who both did an internship at the KNIR and collected most of the content for the digital biography app.

References

1. Ruijgrok, E.: The three economic values of cultural heritage: a case study in the Netherlands. In: Journal of Cultural Heritage, 7:3, 206-213 (2006)

2. Nijkamp, P.: Economic Valuation of Cultural Heritage. In: Licciardo, G., Amirahmasebi, R., The Economics of Uniqueness. International Bank for Reconstruction and Development, Washington (2012)

3. Sebastiani, R., Serlorenzi, M. (a cura di): Il progetto del Nuovo Mercato di Testaccio. In: Workshop di Archeologia Classica 5 (2008)

4. Segarra Lagunes, M. (a cura di): Archeologia urbana e progetto di architettura. pp. 255 Gangemi, Roma (2002)

5. Segarra Lagunes, M. (a cura di): Progetto archeologico. Progetto architettonico. pp. 239 Gangemi, Roma (2007)

6. Travaglini M. C. (a cura di): Un patrimonio urbano tra memoria e progetti. In: Roma, l'area Ostiense-Testaccio, Catalogo della mostra, pp. 73-92, Roma (2004)

7. Roymans, N., Gerritsen, F., van der Heijden, C. and Kolen, J.: Landscape Biography as Research Strategy: The Case of the South Netherlands Project. In: Landscape Research 34: 3, pp. 337-359. (2009)

8. Bosma, K., Kolen, J. (Eds.): Geschiedenis en Ontwerp. Handboek voor de omgang met cultureel erfgoed, Historische Uitgeverij Vantilt, 400 pp., Nijmegen (2010)

9. Bloemers, T.H.J.F., (Eds.): The Cultural Landscape and Heritage Paradox: Protection and development of the Dutch archaeological-historical landscape and its European Dimension, Landscape and Heritage studies, Amsterdam University press, Amsterdam (2010)

10. De Kleijn, M., van Manen, N., Scholten, H., Kolen, J. A spatial data infrastructure for scientific research of the history and heritage of the Dutch urban and rural landscape. (in Press) (2013)

11. Beck-Friis,: The Protestant Cemetery. The Cemetery of Artists and Poets. Malmö Sweden (2010)

12. Contino A., D'Alessandro, L.: Nuovo Mercato Testaccio: prospettive di valorizzazione delle recenti indagini archeologiche. In: Analecta Romana Istituti Danici 32, pp. 51-65. (2006)

13. ESRI, http://www.esri.com/software/arcgis/runtime/

14. Apple, https://developer.apple.com/technologies/tools/

15. VU University Library and SPINlab, http://geoplaza.ubvu.vu.nl/

16. Loechel, A., Schmid, S.: Caching techniques for high-performance Web Map Services. (in Press) In: International Journal of Spatial Data Infrastructures Research (2013)

17. Apple, https://itunes.apple.com/en/app/testaccio/id605348629

Emotional Responses to Artworks in Online Collections

Federico Bertola and Viviana Patti

Università degli Studi di Torino
Dipartimento di Informatica
c.so Svizzera 185, I-10149, Torino, Italy
bertola@celi.it, patti@di.unito.it

Abstract. Artworks have a strong emotional impact on the visitors of an exhibition. Many museums and associations opened their collections for access on the web and have studied the potential of social tagging. User data collected by tagging in art social platforms are a precious information source about emotional responses to artworks and a feedback about the way in which users perceives collections. In this paper, we present our new achievements on this topic within the ArsEmotica framework. Our focus is on eliciting sharable emotional meanings from online collection visitors tags, by interactively involving the users of the virtual communities in the process of capturing the latent emotions behind the tags. We rely on methods and tools from a set of disciplines ranging from Semantic and Social Web to NLP. Such disciplines provide the building blocks for creating a semantic social space, where artworks can be dynamically organized according to a new ontology of emotions inspired by the well-known Plutchik's model of human emotions. The final aim is to involve users in the creation of such emotional space and, then, to offer them an emotion-driven access to the artworks.

Keywords: ontologies, social tagging, affective computing, online collections

1 Introduction

One of the emerging research fields, targeted at extracting information from the data supplied by the Social Web users, is emotion-oriented computing (a.k.a. affective computing), whose focus is to automatically recognize the users' emotions by analyzing their tagging or writing behavior, their conversations or their facial expressions [16]. In particular, the rise of social media has fueled interest in Sentiment Analysis and Opinion Mining: since emotions are often related to appreciation, knowing the feelings of the users towards target issues is important feedback that can support many decisional tasks and has found interesting applications in the business world. Affective computing is receiving increasing attention in many sectors. Its application to the Planet Art, however, is quite at its beginning [9, 5, 6]. In the last years, many museums and cultural heritage institutions opened their collections and archives for access on the web (see e.g. the American Guggenheim, Metropolitan Museum of Art, and San Francisco Museum of Modern Art, the European Rijksmuseum or, more recently, the Google Art project[1]), and the potential of social tagging, in this context, has been explored [19, 11, 5]. User data collected by art social platforms are a precious information source about trends and emotions. Collecting such emotional responses to artworks and collections can be important not only for artists and curators, but also for policymakers in the cultural heritage sector, that need advanced e-participation tools for being supported in their work, both at the decision-making stage, and in the ex-post evaluation of the impact of their policies (e.g. What is the sentiment of citizens about a publicly funded exhibition?).

In this paper we study how to apply Sentiment Analysis to the Planet Art, by exploiting, as information source, tags intended as textual traces that the visitors leave for commenting artworks on social platforms. In particular, we present our new achievements within the ArsEmotica framework [5, 6], where the focus is on methods to extract, from the floating individual meanings and reactions of the visitors of a collection, a *shared emotional semantics*. In previous work [5, 6], we have presented, and evaluated by user study, application software that analyzes tagged artworks from a social tagging platform and provides as output, for each artwork, a set of related emotions. Such emotions emerge as the most significant ones for capturing the users emotional responses toward that resource.

ArsEmotica creates a semantic social space where artworks can be dynamically organized according to an ontology of emotions. The sentiment analysis approach is, indeed, *ontology-driven*. Given a tagged resource, the correlation between tags and emotions is computed by referring to an ontology of emotional categories: from unstructured texts to concepts of an ontology. This is done by exploiting and combining Semantic Web tools and lexical resources. Intuitively, first, tags directly referring to ontological concepts are identified; then, users can give emotional feedback on potentially affective tags, by using the ontology; the final emotional output for the given resource is calculated based on identified emotional concepts, by exploiting automated reasoning on ontology relationships. Notice that there is a strong accent on the social dimension of the process: the key aim is to identify the emotions which better capture the affective meaning, that visitors *collectively* give to the artworks.

In the following, we will describe a new version of the ArsEmotica prototype (Arsemotica 2.0), and we will report about the most significant advancements. Our efforts were mainly devoted to improve the application framework, by devising proper graphical representations of the computed emotional outcomes, in order to effectively involve the user in the process of extracting the latent emotional semantics from tags. In other words, the goal was to setup a platform were the outcomes of our ontology-driven sentiment analysis could be *presented* in a graphical, intuitive way to users, and possibly *enriched* by a further user emotional feedback. To this aim we have designed a new ontology of emotions, which refers to the Robert Plutchik's circumplex model [15]. The ontology pro-

[1] http://www.googleartproject.com/.

vided a guidance in the design a new interface for ArsEmotica, where emotional responses to artworks are represented by means of a graphical representation inspired to the Plutchik's *emotion wheel*, a bi-dimensional representation of the model. Summarizing, the strong points of the new version are: a) the implementation in OWL of a new ontology of emotions inspired by the Plutchik's psycho-evolutionary theory of emotion, and its use in the computation of the emotional outcomes; b) the design of a new interactive user interface inspired by the model underlying the ontology of emotions; c) an extended and generalized architecture for the ArsEmotica application; d) the possibility to delivery the output of the emotional analysis in a machine-understandable format, compliant to emerging standards. The paper is organized as follows. Section 2 contains a brief overview of related work. Section 3 presents ArsEmotica 2.0, with a special focus on the new ontology of emotions. Section 4 presents the new ArsEmotica's interactive user interface. Final remarks end the paper.

2 Related Work

A high interest in monitoring the sentiment of the visitors in environments like museums and exhibitions is recently raised among art practitioners, curators and cultural heritage stakeholders, as witnessed for instance by projects like *e-motion*[2], where the experience of museum-goers are analyzed by means of visitor tracking and biometric measurements technologies. Moreover, many curators, cultural organizations, and museums explored the ways social technologies and principles of Web 2.0 can be applied in museums or exhibitions to foster social awareness and reflections about artworks. In this spirit, many cultural heritage institutions opened their collections for access on the web. Someone went one step beyond [17] by investigating the important role that artworks can play as "social objects", i.e. as a basis for an object-centered sociality, and in the foundation of a new generation of social applications, where the key aim is to stimulate participation and encourage visitors to share their experiences, both in case of virtual and physical collections [9]. Artworks, are a good example of social objects: they are objects that connect people, for instance by fostering conversation. They raise personal feelings and personal degrees of understanding and of acceptance. People can attach to them stories that are taken from their own experience. Recent works such as the visitor's guide in [10] or storytelling-based mobile museum guides [12, 7] can be viewed as contributions in this direction, especially [7] and [10], which pose an accent on social aspects, moving from the consideration that most people visit museums in *groups* [2].

The kind of work we are doing with ArsEmotica [5, 6] can be settled in this picture, but the emotional dimension moves to the foreground. ArsEmotica analyses tags and extracts the prevalent emotions associated to the artworks. It can be seen as an interpretative tool which produces new information by combining the individual contributions of the users. Having the emotional dimension in the foreground can be a key element for encouraging visitors of collections

[2] Mapping Museum Experience: `http://www.mapping-museum-experience.com/en`

to share their experiences about the artworks as social objects. The outcome of the collective experience can provide a means for integrating virtual and physical collections. A recent experiment in this direction is the project "A map of emotions #palazzomadama" for the Robert Wilson's video-portraits exhibition at Palazzo Madama Museum (Turin, Italy). Here visitors were invited to tag artwork's photos, taken during their visit of the physical collection, with the emotions evoked (e.g. #joy, #fear...), by relying on the use of Instagram mobile app running on their smart phones. The Followgram web application was, then, used as sharing platform for collection's photos and affective tags, with the final aim to offer new ways to visit Wilson's collection to online visitors.

Other relevant contributions on art and emotions can be found in a wide range of disciplines, ranging from art and aesthetics to psychology. In particular, recent investigations on psychology of art, which recognizes a specific role to emotions in the aesthetic experience, provide a cognitive grounding to the ArsEmotica perspective and aims. The role of the senses in the arts, both for what concerns creation and reception, has been studied in [3] by taking a multidisciplinary perspective, where computational approaches to emotion recognition can play a key role. Specifically, on this line, recently an interesting experiment has been conducted [20], aimed at analyzing abstract paintings and investigating why a specific painting is perceived as emotional, by combining the use of machine learning algorithms for automatic recognition and eye tracking technology. The approach has been tested on a dataset of abstract paintings (MART museum, Rovereto, Italy), which have been classified as being positive and negative. These results give us the opportunity to observe that, when dealing with the Art World, since affective information evoked by artworks is usually richer than a polarized appreciation, reducing emotional evaluations to positive (or negative) classifications is an over-simplification. On this respect, in ArsEmotica the reference to a rich cognitive-based ontological model of emotions allows us to classify artworks according to a congruous emotional space.

3 Arsemotica 2.0: a New Prototype

In this section, we describe the new version of ArsEmotica, the application software that we developed for testing our ideas. We will briefly recall the main characteristics of the application, by focussing on the original aspects and advancements with respect to previous work [5, 6]. In particular, we will present here: a new extended and generalized architecture (Sec. 3.1); a new OWL ontology, which refers to a state-of-the-art cognitive model of emotions (Sec. 3.2), and inspires the new interactive user interface discussed in Section 4.

3.1 Architecture

Figure 1 reports the four main steps that characterize the computation in ArsEmotica. They are briefly described in the following. The application can be interfaced with any resource sharing and tagging system which provides the data to be processed, i.e. digital artworks and their tags. Social tagging platform for art collections as steve.museum or armeteo.org, with active communities that frequently visit and comment online the collections, would be ideal data sources.

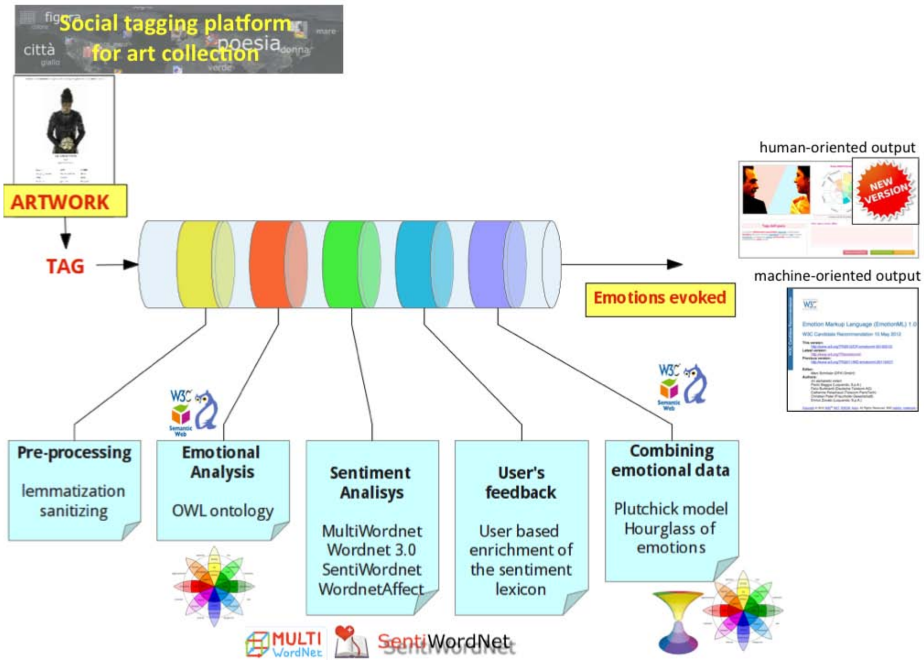

Fig. 1. ArsEmotica overall architecture.

- **Phase 1. Pre-processing: Lemmatization and string sanitizing.** In this step tags associated to a given artworks are filtered so as to eliminate flaws like spelling mistakes, badly accented characters, and so forth. Then, tags are converted into lemmas by applying a lemmatization algorithm, which builds upon *Morph-It!*, a corpus-based morphological resource for the Italian language.
- **Phase 2. Checking tags against the ontology of emotions.** This step checks whether a tag belongs to the ontology of emotions. In other words, it checks if the tags of a given resource are "emotion-denoting" words directly referring to some emotional categories of the ontology. Tags belonging to the ontology are immediately classified as "emotional".
- **Phase 3. Checking tags with SentiWordNet.** Tags that do not correspond to terms in the ontology are further analyzed by means of *SentiWordNet* [4], in order to distinguish *objective* tags, which do not bear an emotional meaning, from *subjective* and, therefore, affective tags. The latter

will be the only ones presented to the user in order to get a feedback on which emotional concept they deliver. The feedback is collected thanks to the interactive user interface described in Sec. 4, which has been designed in tune with the new ontological model of emotion presented below.

– **Phase 4. Combining emotional data and producing a set of emotions as output.** Based on data collected in the previous steps, the so-called emotional engine of the tool offers as output a set of emotions associated to the resource. We have implemented a new algorithm for accomplishing this task, where emotions collected in the previous steps are not simply ranked as in [5] but compared and combined. The algorithm implemented compare collected emotions, by exploiting ontological reasoning on the taxonomic structure of the new ontology of emotions. Moreover, it combines them by referring to the Hourglass Model [8], a reinterpretation of the Plutchik's model, where primary emotions are further organized around four independent but concomitant dimensions (*Pleasantness*, *Attention*, *Sensitivity* and *Aptitude*), whose different levels of activation can give birth to very wide space of different emotions. Shortly, in this model different emotions (basic or compound), result from different combinations of activation levels for the four dimensions. Dimensions are characterized by six levels of activation, which determine the intensity of the expressed/perceived emotion as a float $\in [-1, +1]$. This allows to classify affective information both in a categorical way (according to a number of emotion categories) and in a dimensional format (which facilitates comparison and aggregation), and provided us a powerful inspiration in implementing a new algorithm for combining emotional data in a final output.

  As an advancement w.r.t the previous version, the resulting output can be produced in different modalities. For what concerns human users, emotions evoked by artworks are visualized by a sort of *emotion wheel*, graphically inspired to the color wheel used by Plutchik for offering a bi-dimensional representation of his circumplex model of emotions [14] (Sec. 4). Moreover, in order to foster real interoperability and integration between ArsEmotica and other semantic applications, the system encodes the output in a machine-readable format, by using W3C standards such as RDF and EmotionML, which is currently an emerging standard for emotion annotation.

In the following, we introduce the ontology, which drives the emotional analysis and plays a crucial role in all core steps of the ArsEmotica computation.

3.2 A New Ontology of Emotions based on a Cognitive Model

We have designed for ArsEmotica 2.0 a new ontology of emotional categories based on Plutchik's circumplex model [15, 14], a well-founded psychological model of emotions. The ontology is written in OWL, structures emotional categories in a taxonomy, which includes 32 emotional concepts. Due to its role within the ArsEmotica architecture, the ontology has been conceived for categorizing

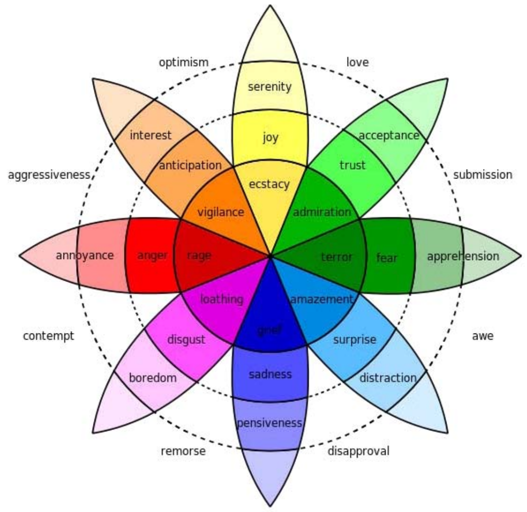

Fig. 2. Plutchik's circumplex model [15].

emotion-denoting words, as the one used in the previous version of the application. It, then, includes two root concepts: *Emotion* and *Word*.
Class Emotion For what concerns the class *Emotion*, the design of the emotional categories taxonomic structure, of the disjunction axioms and of the object and data properties mirrors the main features of Plutchik's circumplex model, (see the two-dimensional representation in Fig 2). In particular the *Emotion*'s hierarchy includes all the 32 emotional categories presented as distinguished labels in the model. Such model is represented as a *wheel of emotions*, which encodes the following elements and concepts:

– **Basic or primary emotions**: *joy*, *trust*, *fear*, *surprise*, *sadness*, *disgust*, *anger*, *anticipation* (i.e. *expectancy*); in the color wheel this is represented by differently colored sectors.
– **Opposites**: basic emotions can be conceptualized in terms of polar opposites: *joy* versus *sadness*, *anger* versus *fear*, *trust* versus *disgust*, *surprise* versus *anticipation*.
– **Intensity**: each emotion can exist in varying degrees of intensity; in the wheel this is represented by the vertical dimension.
– **Similarity**: emotions vary in their degree of similarity to one another; in the wheel this is represented by the radial dimension.
– **Complex emotions**: beside basic emotions, there are complex emotions, that are a mixtures of the primary emotions, just as some colors are primary, and others made by mixing the primary colors; in the model in Fig 2 emotions in the blank spaces are compositions of basic emotions called *primary dyads*.

Class Word For what concerns the class *Word*, it is the root for the emotion-denoting words, i.e. those words which each language provides for denoting emotions. Since we actually applied our application to use cases where tagging involved Italian communities, we actually defined and populated the subclass *ItalianWord*. However, the ontology is already designed to be extended with further subclasses of *Word*, for representing emotion-denoting words in different languages. Intuitively, each instance of the *Word* and *Emotion* concepts has two parents: one is a concept from the *Emotion* hierarchy (the emotion denoted by the word, e.g. *rage*), while the other is a concept from the *Word* hierarchy (e.g. Italian, the language the word belongs to). For instance, the following code excerpt corresponds to the description of the Italian affective word *rabbia*: it is both an instance of the concept *Rage* (an intense anger), and an instance of the concept *ItalianWord*, i.e. *rabbia* is an Italian word for denoting rage:

```
<rdf:Description rdf:about="http://www.arsemotica.unito.it/ontology/emotions.owl#rabbia">
  <rdf:type rdf:resource="http://www.w3.org/2002/07/owl#NamedIndividual"/>
  <rdf:type rdf:resource="http://www.arsemotica.unito.it/ontology/emotions.owl#ItalianWord"/>
  <rdf:type rdf:resource="http://www.arsemotica.unito.it/ontology/emotions.owl#Rage"/>
</rdf:Description>
```

Ontology Population We semi-automatically populated the ontology with Italian words by following the same methodology described in [5] for populating OntoEmotion. Shortly, we relied on the multilingual lexical database MultiWordNet [13] and its affective domain WordNet-Affect, a well-known lexical resource that contains information about the emotions that the words convey. A human expert checked the identified terms. WordNet is a lexical database, in which nouns, verbs, adjectives and adverbs (lemmas) are organized into sets of synonyms (synsets), representing lexical concepts. We have manually chosen an initial set of representative Italian emotional words for each concept. At the beginning we have chosen a set of nouns. Then, by making use of morpho-semantic relations between nouns and verbs or between nouns and adjectives specified in Italian dictionaries, we have expanded the initial set including new related verbs and adjectives[3]. Such words were used as entry lemmas for querying the lexical database. The result for a word is a synset, representing the "senses" of that word, which are labeled by MultiWordNet unique synset identifiers. Each synset was then processed by using WordNet-Affect [18]: when a synset is annotated as representing affective information, then, *all the synonyms belonging to that synset* are imported in the ontology as relevant Italian emotion-denoting words. This allowed us to automatically enrich the ontology with synonyms of the representative emotional words, but also to filter out synsets which do not convey affective information. Currently, the resulting ontology contains about 500 Italian words referring to the 32 emotional categories of the ontology.

[3] Unfortunately, this process has been carried on mainly manually: available lexical resources do not provide the information on such kind of relations for Italian.

4 An Interactive User Interface for Collecting and Presenting Emotional Responses

Our choice to design an ad hoc ontology for ArsEmotica was driven, on the one hand, by the need to find a interpretative graphical representations for presenting the outcomes of the automatic elaboration of the artworks' tags. On the other hand, we were looking for a simple interface to foster the users to annotate tags having an indirect affective meaning, by means of emotional concepts from the ontology. On this perspective, we have found the Plutchik's model very attractive for three main reasons:

- The reference to a graphical wheel is very intuitive and offers a *spacial representation of emotions* and their different relations (similarities, intensities, polar oppositions). Such kind of representation allows to convey to the user a rich information on the emotional model, without referring to tree-like visualization of the ontology hierarchy. For instance, looking at the wheel, it is natural to catch a similarity relation among *joy* and *trust*, while *joy* and *sadness* are polar opposites, and *love* emotion is composed of *joy* and *trust*.
- The use of *colors* for denoting different emotions provides a very intuitive communication code. Different color nuances for different emotions, transmit naturally the idea that primary emotions can blend to form a variety of compound emotions, analogous to the way colors combine to generate different color graduations. This aspect can play an important role in the development of an user interface for a cultural heritage application, and can be exploited for proposing to the user intuitive metaphors to browse the emotional space. In particular, it inspired us in the design of an interface to easy the task of emotionally classifying tags.
- The number of emotional categories distinguished in the wheel is *limited*. This aspect facilitates the user that is involved in an emotional evaluation.

Let us now present the interactive user interface that we have developed, and its sample application on a tagged artwork from the ArsMeteo online collection [1] (http://www.arsmeteo.org). The sequence of interactions offered to the user follows the flux of computation sketched in Fig 1.

An Emotional Wheel for Presenting Emotions Evoked by the Artworks (Phase 2). After the user selects an artwork from the collection, the application applies the emotional analysis on the artwork tags. The result of this computation, i.e. the *evoked emotions*, is presented to the user by a graphical representation called "La rosa delle emozioni", which strongly recalls the Plutchik's color wheel. For instance, by applying the emotional analysis to the artwork "Dove la Raffinata Ragazza Bionda guarda il Grosso Toro Malmorto" by Filippo Valente (Fig 3), the four red colored tags are identified as emotional according to the emotional ontology: 'orrore', 'infamia', 'cattiveria', 'tristezza'; the presence of emotional responses related to *sadness* and a strong disgust

Fig. 3. Showing the results of the automatic emotional analysis

(*loathing*) is highlighted by coloring the sectors of the emotion wheel corresponding to those emotions. Internal sectors of the ArsEmotica's wheel are intended as representing light intensity of emotions, while the external ones as representing high intensity. Underlined blue colored tags denotes tags that have been recognized by the sentiment analysis stage as possibly conveying some affective meaning. Then, they appear as active links for the user's emotional feedback: see e.g. 'sangue', 'sconfiggere', and so on.

User's Emotional Feedback (Phase 3). The Arsemotica's wheel offers an effective way for fostering users to annotate tags having an indirect affective meaning by means of emotional concepts from the ontology. After the user selects a tag to evaluate, the application activates a pop-up window, where an uncolored emotional wheel is shown. Users can express the emotional evaluation in terms of basic emotions with different intensities, and color the wheel accordingly, by clicking on one of the 24 sectors of the wheel; otherwise they can select compound emotions, by selecting the wedge-shaped triangles inserted between the basic emotions. In the example in Figure 4 the user associated to the tag 'sangue' (blood) the emotions *fear* and *disgust* (with high intensity, which corresponds to *loathing*). Notice that the tag evaluation is contextual to the vision of the artwork, which indeed remains visible in the background.

Final Emotional Evaluation of the Artwork (Phase 4). After combining the emotional data of phases 2 and 3, the resulting emotional evaluation is again presented to the user by using the ArsEmotica's wheel. If requested, the application can format the result in the standard markup language EmotionML.

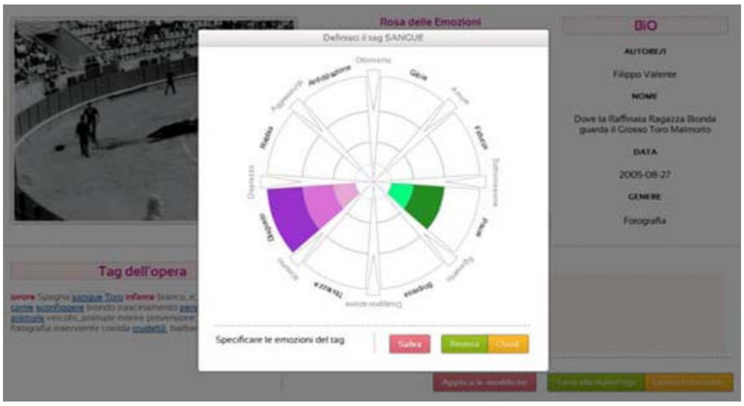

Fig. 4. Interaction with the user: collecting the tag-mediated emotional feedback

5 Conclusion and Future Work

In this paper we have presented our new achievements within the ArsEmotica framework. Our efforts were devoted to improve the application framework, by devising proper interpretative graphical representations to present the outcomes of the ArsEmotica analysis to the users. We described our new interface, that presents emotional responses to artworks by means of an emotion wheel inspired to the Plutchik's one. In general, designing engaging interfaces that allow an appropriate granularity of expression is not a trivial task. Therefore, the next step will be to evaluate the new prototype, and to carry on a user test. We plan to use tagged artworks from the ArsMeteo tagging platform as dataset and to involve users of the ArsMeteo community, which in the past have already actively participated to a user study on the first version of our prototype [6], and manifested great interest on the topic.

For what concerns the possible uses of ArsEmotica, we think that it can be exploited as a co-creation instrument for museums and virtual galleries. Moreover, the capability of extracting prevalent emotions can foster the development of emotion-aware search engines, emotional tag clouds or interactive map of emotions, which could enable new ways of exploring art collections.

References

1. Edoardo Acotto, Matteo Baldoni, Cristina Baroglio, Viviana Patti, Flavio Portis, and Giorgio Vaccarino. Arsmeteo: artworks and tags floating over the planet art. *In Proc. of ACM HT '09*, ACM:331–332, 2009.
2. Liliana Ardissono, Tsvi Kuflik, and Daniela Petrelli. Personalization in cultural heritage: the road travelled and the one ahead. *User Model. User-Adapt. Interact.*, 22(1-2):73–99, 2012.
3. F. Bacci and D. Melcher. *Art and the Senses.* OUP Oxford, 2011.

4. Andrea Esuli Stefano Baccianella and Fabrizio Sebastiani. SentiWordNet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In *Proc. of LREC'10*. ELRA, May 2010.
5. Matteo Baldoni, Cristina Baroglio, Viviana Patti, and Paolo Rena. From tags to emotions: Ontology-driven sentiment analysis in the social semantic web. *Intelligenza Artificiale*, 6(1):41–54, 2012.
6. Matteo Baldoni, Cristina Baroglio, Viviana Patti, and Claudio Schifanella. Sentiment analysis in the planet art: A case study in the social semantic web. In Cristian Lai, Giovanni Semeraro, and Eloisa Vargiu, editors, *New Challenges in Distributed Information Filtering and Retrieval*, volume 439 of *Studies in Computational Intelligence*, pages 131–149. Springer, 2013.
7. Charles Callaway, Oliviero Stock, Elyon Dekoven, Kinneret Noy, Yael Citron, and Yael Dobrin. Mobile drama in an instrumented museum: inducing group conversation via coordinated narratives. *New Review of Hypermedia and Multimedia*, 18(1-2):37–61, 2012.
8. Erik Cambria, Andrew Livingstone, and Amir Hussain. The hourglass of emotions. In Anna Esposito, Antonietta Maria Esposito, Alessandro Vinciarelli, Rüdiger Hoffmann, and Vincent C. Müller, editors, *COST 2102 Training School, Revised Selected Papers*, volume 7403 of *Lecture Notes in Computer Science*. Springer, 2012.
9. Gunho Chae, S. Joon Park, Robert Stein, Jungwha Kim, and Susan Wiedenbeck. Exploring affective computing for enhancing the museum experience with online collections. In *Proc. of the Museums and the Web*, 2012.
10. Tsvi Kuflik, Oliviero Stock, Massimo Zancanaro, Ariel Gorfinkel, Sadek Jbara, Shahar Kats, Julia Sheidin, and Nadav Kashtan. A visitor's guide in an active museum: Presentations, communications, and reflection. *J. Comput. Cult. Herit.*, 3(3):1–25, February 2011.
11. Yi-Ling Lin and Lora Aroyo. Interactive curating of user tags for audiovisual archives. In *Proceedings of the International Working Conference on Advanced Visual Interfaces*, AVI '12, pages 685–688, New York, NY, USA, 2012. ACM.
12. Vincenzo Lombardo and Rossana Damiano. Storytelling on mobile devices for cultural heritage. *New Review of Hypermedia and Multimedia*, 18(1-2):11–35, 2012.
13. Emanuele Pianta, Luisa Bentivogli, and C. Girardi. Multiwordnet: developing an aligned multilingual database. In *Proc. of Int. Conf. on Global WordNet*, 2002.
14. Robert Plutchik. The circumplex as a general model of the structure of emotions and personality. In R. Plutchik and H. R. Conte, editors, *Circumplex models of personality and emotions*, pages 17–47. American Psychological Association, 1997.
15. Robert Plutchik. The Nature of Emotions. *American Scientist*, 89(4), 2001.
16. Marc Schroeder, Hannes Pirker, Myriam Lamolle, Felix Burkhardt, Christian Peter, and Enrico Zovato. Representing emotions and related states in technological systems. In Roddy Cowie, Catherine Pelachaud, and Paolo Petta, editors, *Emotion-Oriented Systems*, Cognitive Technologies, pages 369–387. Springer, 2011.
17. Nina Simon. *Participatory Museum*. Museum 2.0, 2010.
18. C. Strapparava and A. Valitutti. WordNet-Affect: an affective extension of WordNet. In *Proc. of LREC'04*, volume 4, pages 1083–1086, 2004.
19. Jennifer Trant and Bruce Wyman. Investigating social tagging and folksonomy in art museums with steve.museum. In *Proc. of the Collaborative Web Tagging Workshop (WWW'06)*, 2006.
20. Victoria Yanulevskaya, Jasper Uijlings, Elia Bruni, Andreza Sartori, Elisa Zamboni, Francesca Bacci, David Melcher, and Nicu Sebe. In the eye of the beholder: employing statistical analysis and eye tracking for analyzing abstract paintings. In *Proc. of the ACM Int. Conf. on Multimedia*, MM '12, pages 349–358. ACM, 2012.

Personalized Nichesourcing: Acquisition of Qualitative Annotations from Niche Communities

Chris Dijkshoorn[1], Mieke H. R. Leyssen[2], Archana Nottamkandath[1], Jasper Oosterman[3], Myriam C. Traub[2], Lora Aroyo[1], Alessandro Bozzon[3], Wan Fokkink[1], Geert-Jan Houben[3], Henrike Hövelmann[4], Lizzy Jongma[4], Jacco van Ossenbruggen[2], Guus Schreiber[1], Jan Wielemaker[1]

[1] The Network Institute, VU University Amsterdam, the Netherlands
{c.r.dijkshoorn, a.nottamkandath, lora.aroyo, w.j.fokkink, guus.schreiber, j.wielemaker}@vu.nl

[2] Centrum Wiskunde en Informatica, Amsterdam, the Netherlands
{leyssen, traub, jacco.van.ossenbruggen}@cwi.nl

[3] Web Information Systems, Delft University of Technology, the Netherlands
{j.e.g.oosterman, a.bozzon, g.j.p.m.houben}@tudelft.nl

[4] Rijksmuseum Amsterdam, the Netherlands
{h.hovelmann, l.jongma}@rijksmuseum.nl

Abstract. Diversity and profundity of the topics in cultural heritage collections make experts from outside the institution indispensable for acquiring qualitative and comprehensive annotations. We define the concept of nichesourcing and present challenges in the process of obtaining qualitative annotations from people in these niches. We believe that experts provide better annotations if this process is personalized. We present a framework called Accurator, that allows to realize and evaluate strategies and applications for personalized nichesourcing.

Keywords: cultural heritage, nichesourcing, annotation framework, qualitative annotations

1 Introduction

Acquiring qualitative annotations for the enrichment of cultural heritage collections is a significant effort for museums and heritage institutions. In our research we confront the challenge of obtaining accurate annotations by involving expert communities within the crowd, i.e. people that are external to the institution. For this objective, we turn to personalized nichesourcing, where we aim to identify the niche communities and find ways to adapt the annotation task to them.

In this paper, we describe briefly the motivation behind the approach and the project in which the investigations take place. We present the four main research challenges that drive the detailed investigations in addition to the main aspects of the implementation.

2 Motivation

Access and retrieval mechanisms for archives and museums typically rely on a rich description of the collection. Most cultural heritage institutions therefore employ professional art historians to describe their collections by manually compiling metadata for each item. The subject matter of collection items can be very diverse, for example, it can consist of historic figures, animals, plants and buildings. Additionally, these aspects often carry a hidden symbolic meaning. To adequately describe items in large and diverse collections, the knowledge of experts from domains other than art history is indispensable. Cultural heritage institutions therefore seek to understand whether and how they can make use of external users to produce these annotations.

The work in this research aims at understanding which strategies and techniques lead to precise annotations by (crowds of) users that are external to the museum. For this, the detailed investigations are organized in terms of four connected challenges, that we will describe further in Section 3. The first challenge in the project is to identify and model the niche of relevant experts and to motivate them to contribute to the annotation of collection items. Next, personalization mechanisms must make sure that the annotation task is adapted to the experts such that they are shown items that correspond to their expertise. The quality of the annotations and the level of expertise of the annotators have to be evaluated using trust evaluation algorithms. As a final challenge, all these aspects must be presented in an appropriate interface.

In order to perform this research, we develop a framework to support crowd annotation processes, called Accurator. It is used to conduct studies within the SEALINCMedia research project, for example in a use case with Rijksmuseum Amsterdam, as we will see in the example later.

3 Research Challenges

The overall objective of understanding which strategies and techniques lead to high-quality annotations by (crowds of) external users, is approached through four connected challenges.

One of the four main challenges of nichesourcing is finding candidate annotators that are able to produce high quality annotations for collection items. We believe that people participating in a specialist community have an active interest in that topic and might be willing to help and share knowledge related to it. We refer to these specialist communities as *niches* and focus on their manifestation, among others, on the social web. We analyze social data and perform user studies using the Accurator tool to understand what identifies a niche community, what indicates that a person is part of such a community and which properties identify a good candidate to provide qualitative annotations.

The challenge for recommender strategies in Accurator is twofold: keep the expertise needed to annotate the item in the range of the experts' knowledge and yet diversify the suggestions to get high-quality annotations for as many distinct items as possible. Our aim is to develop recommender strategies that use content patterns from the Linked Data cloud, resulting in a list of recommendations consisting of diverse items. We hypothesize that encountering diverse items to annotate will help keep the expert motivated.

We address issues of determining trust in the expert users and their contributed annotations by modeling the user reputation and tracking their expertise across various topics over time. We believe subjective logic is suitable to model the reputation of users and semantic similarity measures can be used to track and update the users' expertise. Since there is no gold standard for evaluating the annotations, we must rely on a peer reviewing process and other mechanisms such as determining the provenance of the annotations.

Since external users are not familiar with professional classification schemes and (art-)historical expert knowledge, our fourth challenge is to break down the annotation process into facile tasks that can be solved with little effort and without professional knowledge. We believe that the interface for such a system has to present the task in a straightforward way while motivating the users to spend the time contributing their knowledge. We investigate which design aspects and underlying mechanisms are responsible for the quality and quantity of tags added by users and how to visualize trust and personalization aspects.

4 Accurator Framework

The Accurator framework is developed to support and implement strategies and techniques which confront the previously mentioned challenges. We explicitly design the framework to test different strategies on various collections of artworks. In this section, we present the main system aspects.

Our main assumption is that making use of personalized nichesourcing increases the quality of annotations. We believe that we can automatically identify niche candidate users and create relevant user profiles to support their annotation task. Based on this knowledge about the candidate experts, we can then recommend them relevant annotation tasks and apply trust mechanisms to improve the recommendation and annotation strategies. Figure 1 shows the corresponding Accurator workflow.

The process starts (see Figure 1a) with searching the social web for user-generated content that is relevant for a specific topic. We calculate the relevance of the content creators with respect to the topic and exploit social relations to identify a topical niche and candidate experts from that niche. When a person starts using Accurator, a user profile (see Figure 1b) is created based on available data.

The next step (see Figure 1c) is the recommendation of collection items for a user to annotate. The recommendation strategy is based on specific patterns in the data, the user profile, and the current annotation quality of an item. Accurator allows to easily switch between different strategies to cater for users' diversity. In the process of personalizing the recommendations, the user's choice of items to annotate will subsequently affect the calculated interest of that user.

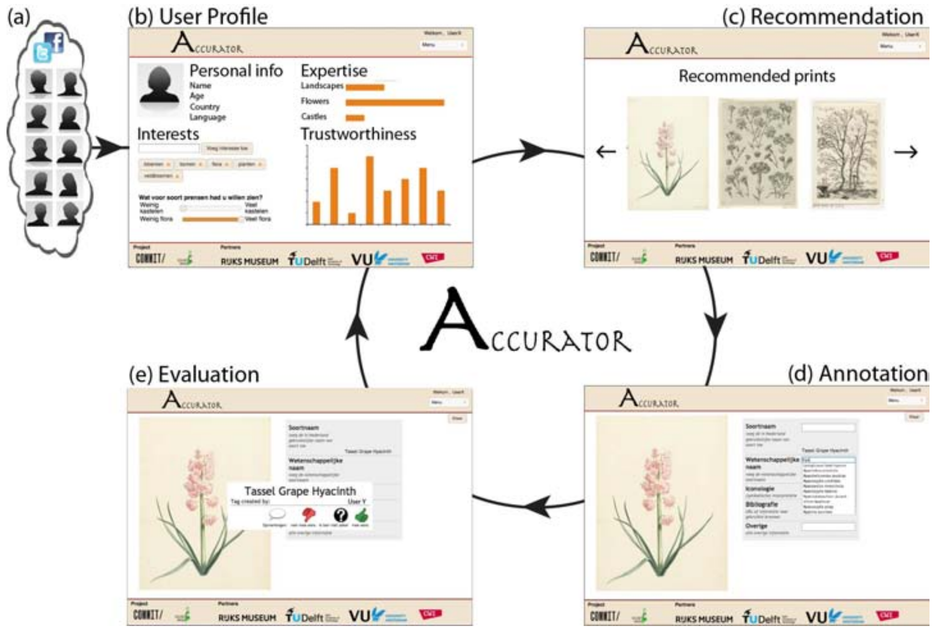

Fig. 1. Accurator personalized nichesourcing workflow

Figure 1d shows the interface where users add their annotations to an item. The presented fields depend on the topic and the user's expertise on that topic. Accurator can be configured to use domain vocabularies to support the user. Figure 1e shows the interface in which users can evaluate and review the annotations of other users. This task is only available to users who are considered trustworthy and have a certain level of expertise. The result of a review affects 1) the quality of an annotation, 2) the expertise level of the user, and 3) the trustworthiness of another user.

The Accurator prototype is built using Cliopatria[5] to store RDF, Google Web Toolkit[6] for the user interface, and Google App Engine[7] for hosting. Accurator is now used for experimentation with artwork data from the Rijksmuseum Amsterdam and a demo is available at `http://rma-accurator.appspot.com`.

Acknowledgements. This publication is supported by the Dutch national program COMMIT. We like to thank the members of the SEALINCMedia worktable and in particular the Rijksmuseum for their support.

[5] `http://cliopatria.swi-prolog.org/`
[6] `https://developers.google.com/web-toolkit/`
[7] `https://developers.google.com/appengine/`